Reasoning and Acting in Time

by

Haythem O. Ismail

August 23, 2001

A dissertation submitted to
the Faculty of the Graduate School of
State University of New York at Buffalo
in partial fulfillment of the requirements for
the degree of
Doctor of Philosophy

Department of Computer Science and Engineering

Dissertation Committee:
Dr. Stuart C. Shapiro
Dr. William J. Rapaport
Dr. Jean-Pierre A. Koenig

Acknowledgments

To be sure, acknowledgments are the hardest for me to write. For I believe that a Ph.D. dissertation is not just the result of a few years of graduate research; it is the culmination of a life-long process of learning, maturing, and interacting with others. Almost everyone whom I have known in some depth has somehow contributed to this dissertation, at least indirectly by contributing to who I have come to be. But, of course, for limitations of time and space, I cannot give each and every one their due. For this reason, I start by apologizing to those who are not mentioned below for having missed their names, and to those mentioned for words that can never do them justice.

The key to a successful Ph.D. episode is a good advisor. I thank my advisor, Stuart Shapiro, for everything that he taught me, for his patience and understanding, and for supporting me every step of the way. Never have I walked out of his office but motivated, inspired, and encouraged; how better can an advisor be? I will always be proud to be Stu's student.

Bill Rapaport and Jean-Pierre Koenig are my other committee members. Bill's keen philosopher's eye and impressive breadth of knowledge have certainly contributed positively to this dissertation. Both as his student and his TA, I have always been inspired by his precision and organization. JP's contribution to my graduate career is most notable; he always brought a different, insightful point of view to my research. Most of my knowledge of linguistics has been directly, or indirectly, through him. Thank you Bill and JP.

Special thanks to Don Perlis, my outside reader, for his helpful suggestions and impressively fast turn-around time.

To my office-mates, batch-mates, and class-mates Fran Johnson and John Santore: thank you for being great friends. Thanks to other SNeRG friends for their support and technical advice: Carl Alphonce, Bharat Bhushan, Debbie Burhans, Alistair Campbell, Rob Church, Tiago Maia, Tony Petre, and David Pierce. To SNeRG alumni: thank you for your efforts on SNePS that made my research career much easier. I am particularly grateful to Michael Almeida, Henry Hexmoor, and Deepak Kumar whom I have never met but on whose shoulders I have stood.

My sincere thanks to Bharat Jayaraman, Bina Ramamurthy, Ken Regan, Alan Selman, and Deborah Walters for their support and encouragement since my early days at UB. Many thanks to the wonderful CSE administrative staff, especially Eloise Benzel, Sally Elder, Joan Glinski, Gloria Koontz, and Lynne Terrana who were always there for help. Finishing my Ph.D. would have been much harder if it wasn't for the great technical support and system administration staff: Nina Banerjee, Harry DeLano, Davin Milun, Ken Smith, and Matt Stock.

I'd also like to thank members of the center for cognitive science for their support and for lively and inspiring discussions that have, no doubt, helped shape my character as a cognitive science researcher. Special thanks to Len Talmy who taught me cognitive linguistics and who has always been a source of encouragement and insightful comments, Roberto Casati for his valuable comments on earlier drafts of Chapters 3 and 4, and Gail Mauner and Erwin Segal for their continuous interest in my research.

My sincere thanks to my loved ones in Egypt: mom, dad, Rasha, Khaled, Yara, Hosaam, and relatives and friends. I deeply appreciate their support, encouragement, and patience during the five years of Ph.D. work when I was not there for them.

I am specially grateful to Dalia Fadel; if it wasn't for her support, the final months would have been much more hectic.

This work was supported in part by ONR under contract N00014-98-C-0062.

To Mom and Dalia

Contents

List of Figures

List of Tables

Abstract

Reasoning and Acting in Time
by Haythem O. Ismail
Major Professor: Stuart C. Shapiro, Ph.D.

This dissertation investigates aspects of the design of an embodied cognitive agent that interleaves reasoning, acting, and interacting with other agents while maintaining a record of what has happened and is happening in its environment. Among other things, such knowledge of its own history allows the agent to reason more effectively when faced with an emergency situation such as the failure of one of its acts. Crucial to distinguishing what *has* happened from what *is* happening is a notion of the present time, or "now". There has been much research in artificial intelligence on issues of time. However, it is typically the case that, although an agent reasons *about* time, it is not itself situated *in* time. Once the agent has a concept of "now", one that continuously changes to reflect the progression of time, a different kind of temporal reasoning problem emerges. For example, given that an agent could be told anything, how are formal representations of present states to be distinguished from those of past states, where the former, but not the latter, are pertinent to the agent's actions? And, given that the present must be distinguished, how is "now" to be modeled so that the mere passage of time does not result in computationally costly knowledge-base-editing routines? How can the agent reason about "now", when the very process of reasoning results in "now" changing? How can the agent be endowed with a feel for how much time has passed, which seems crucial for reasoning about persistence as time passes while the agent acts?

In this dissertation, the above issues are investigated in detail. A theory of subjective time is presented, accounting for a cognitive agent's *vague* concept of "now", its sense of temporal progression, and its feel for how much time has passed. An investigation of the impact of embodiment and time perception on issues of reasoning about persistence as the agent acts comes out as a natural by-product of the theory. The theory of subjective time is wrapped around a core logic of objec-

tive time that axiomatizes various aspectual phenomena needed for reasoning, acting, and natural language interaction. Unlike most theories of linguistic aspect, the logic of aspect presented here accounts for the agent's knowledge of states, events, and processes, not only from *static* natural language inputs, but also from the *dynamic* accumulation of perceptual and proprioceptual information as events unfold in time. Based on the logical analysis of the notion of telicity (the analysis goes beyond the standard telic/atelic distinction), a theory of how cognitive agents may employ reasoning to control the execution of sequences of acts is presented. The theory proposes a principled way by which agents may decide when it is time to move on to the next step in a sequence; an issue that has not been given much attention in the literature. Finally, the dissertation establishes a framework for interrupt-handling and error recovery based on a system of context-sensitive priorities among acts.

Chapter 1

Introduction

The passage of time is inevitable. It relentlessly moves forward, turning every moment of conscious experience into a link in a chain of past memories. You cannot stop the world, nor can you rewind it. Acting, talking, and even thinking take time. Everything exists in time, and time affects everything. Humans are particularly affected by the passage of time, and the study of embodied cognitive agents in artificial intelligence (AI) shows that the intimate relationship we have to time is surprisingly intricate.

In this dissertation, I investigate aspects of the design of a reasoning, acting, interacting agent that maintains a representation of a relentlessly moving "now". This research falls within what has come to be known as *cognitive robotics*. Cognitive robotics is that branch of AI concerned with "the study of the knowledge representation and reasoning problems faced by an autonomous robot (or agent) in a dynamic and incompletely known world" (Levesque and Reiter, 1998, p. 106). Needless to say, I will not be solving *all* the problems of cognitive robotics; rather, this research is about studying, fleshing out, and investigating solutions to a subset of them, in particular, a subset of those problems that face an agent *reasoning and acting in time*. These are problems that emerge when an agent interleaves reasoning, acting, and interacting with other agents while maintaing a sense of the passage of time. The issues to be covered include:

1. Representing "now".

2. Modeling the progression of time.

3. Reasoning about "now" as it continuously changes.

4. Employing embodiment and time perception to account for knowledge of the persistence of states.

5. Using perception and reasoning to control the execution of sequences of acts.

6. Handling interrupts.

7. Grounding notions of linguistic aspect in perception and action.

In this introductory chapter, the scope of the dissertation is demarcated, the framework and assumptions constraining the research methodologies adopted are outlined, and a brief discussion of the above-listed issues is presented.

1.1 The Two Extremes in Agent Design

Different designs of acting agents may be located at different points on a scale whose two extremes represent fully-deliberative agents and fully-reactive agents. A fully-deliberative acting agent is one whose every behavior is the result of its *reasoning* about relevant aspects of the world. By "reasoning" I refer to some process of manipulating a symbolic representation of the world and/or the agent's knowledge of it according to some sound system of logic.[1] A logic, as far as this work is concerned, is any system made up of a formal language, with clear syntax and semantics, along with a set of rules of inference. A symbolic representation of the world (and/or the knowledge of the agent thereof) is, essentially, a knowledge base—a set of sentences of the formal language of the logic—and manipulating this representation is done through applying the rules of inference in order to add (or sometimes retract) sentences from the knowledge base. On the other hand, a fully-reactive acting agent does not have/need any logical representation of the world. It has built-in reactive mechanisms that allows it to act in response to sensations of its immediate spatio-temporal surroundings. It should be clear that both deliberative and reactive agent designs have their pros and cons. On one hand, a fully-reactive agent cannot come up with complicated plans of action in the face of novel situations or in order to achieve new goals, especially if coming up with these plans requires knowledge of more than the immediate spatio-temporal surroundings—the *here* and *now*, if you will. A deliberative agent, with general-enough knowledge of the world, has a much better

[1] Actually, the soundness requirement may be relaxed, given that human reasoning is probably *not* sound.

chance of coming up with such plans. On the other hand, a fully-deliberative agent runs the risk of spending too much time reasoning when swift action is required, whereas a reactive agent with hard-wired behaviors would certainly be more prompt in its reaction.

Evidently, a reasonable agent design falls at neither extreme of the deliberative-reactive scale but somewhere in between. One may combine the best of the two worlds with a layered agent architecture. In such an architecture, time-critical behaviors could be hard-wired at a low-level that interacts directly with the immediate surroundings, while reasoning and planning at higher levels effect less time-critical, but reasoning-intensive, courses of action. This approach has the appeal of not only presenting a practically-viable alternative but also providing a frame-work within which *human*-like intelligent behavior may be modeled. In particular, humans exhibit both reactive behavior (in collision avoidance, for example) and deliberative action (for example, planning which route to take to school given knowledge of the weather, time, road reports, etc.)

This work assumes such an intermediate approach and indeed adopts a layered agent architecture (see Section 1.5.1). The main focus, however, is on the higher levels, where reasoning takes place. Hence, I will have nothing to say about reactive behaviors and how exactly they feature in the agent's actions. Deliberative action and interleaving it with reasoning is my primary concern.

1.2 The Two Approaches to Deliberative Action

Deliberative action involves two activities: deliberating, or reasoning about action, and actual action execution in the world. The different temporal and causal connections between these two activities give rise to two major approaches to action in logic-based AI. With *off-line* execution of action, the agent first reasons, and the result of its reasoning is a formal representation of a sequence of actions. As far as the agent knows, should such a sequence be executed, it would effect whatever goal the agent is trying to achieve. Given such a sequence, the agent proceeds to execute it without any further reasoning—*unconsciously* so to speak. It should be noted that the distinctive feature of off-line execution is *not* the generation of a complete action sequence before actually executing it. Rather, it is the inability of the agent to reason about what it is doing *as it is doing it*, or about how the world evolves *as it evolves*. That is, the process of reasoning is not situated in the temporal and causal environment where execution takes place.

There are obvious problems with off-line execution. The most obvious is that execution is neither monitored nor controlled by the reasoning system. Thus, should an unexpected event happen, something which is very common in the real world, the agent would not have a chance to reason about how to adjust its plan or react appropriately. The result is that, except for highly controlled environments, the agent's actions are often doomed to failure. In addition, even the simple, error-free execution of sequences of actions might require reasoning for the agent to know when to move on to the next step in the sequence.

With *on-line* execution, an agent causally and temporally interleaves the processes of reasoning and acting. Note that, as with off-line execution, the agent may plan ahead, generating a complete action sequence before actually executing it. However, the crucial difference is that, with on-line execution, the agent can monitor and reason about the very process of execution. This gives it more flexibility to attend to errors and unexpected situations. In addition, the on-line approach presents a more realistic view of human-like intelligent action, whose investigation and modeling is of primary concern to this work.

But even within the on-line execution paradigm, there are different degrees to which the agent is *involved* in what goes on during action execution. For example, the agent's knowledge of what happens during its course of action may be restricted to its awareness of the *present* state. This, together with general background knowledge, is probably sufficient for the agent to react appropriately to an emergency situation. Nevertheless, in some sense the agent, though deliberative, merely reacts to current states of its environment. A more sophisticated agent would also *remember* what went on during its course of action: what it did, what emergency situations it encountered, how it reacted to them, when they happened, and when its reaction took place. This higher level of *awareness* of what is going on is particularly valuable for an agent that is expected to interact with other agents, mainly humans, reporting on what it did and what took place as it did it.

For this to be at all possible, the agent should maintain a temporal model of its environment and of how it evolved. It is the interleaving of reasoning, acting, and interacting as time relentlessly goes by that is the focus of the current study. It is what I call *reasoning and acting in time*.

1.3 The Factor of Time

In order to maintain a record of what it did and how the world within which it is situated evolved over time, an agent needs to keep track of the temporal relations among events it experienced, including its own acts. For example, it should know that, while acting, it picked up a block before carrying it to the nearby table and putting it down. It should also remember that, for example, its battery ran down as it was moving to the table and that it, therefore, had to stop, recharge the battery, and then resume its trip to the table. Although such objective knowledge of time is indeed necessary, much more is actually needed.

Most acting systems in AI incorporate time and reasoning about it either implicitly or explicitly. However, in most of those systems, time happens to be yet another domain phenomenon that a cognitive agent occasionally needs to reason about. In particular, as represented in those systems, time is totally decoupled from the agent's experience as it reasons and acts; the agent does not have any personal sense of the real time. Two questions now arise. What is a *personal* sense of time, and why is it important for cognitive agents? Basically, a personal sense of time is based on a representation of the real-time present that (i) allows the agent to distinguish it from the past and the future, and (ii) is continuously changing to reflect the progression of time. This is important for various reasons. First, an agent reasoning about a dynamic domain, one that changes as it is reasoning and acting, needs to distinguish present facts from past ones. In addition, since the agent is itself acting in such an environment, distinguishing the present is crucial for it to identify those facts that are pertinent to its actions. Second, in addition to acting, the agent may also be *interacting* with other agents (humans, for example). Such interaction, whether in natural language or not, would only be effective and intelligible if the agent has a clear distinction between what is now the case, versus what used to be the case. There are at least two reasons why this distinction is important. First, if the agent expresses its beliefs in natural language, then a distinction between the present, the past, and the future is required for generating sentences with the correct tense. Second, and more important, present facts have a distinguished status for an acting agent. Suppose an agent believes that whenever the fire-alarm sounds it should leave the building. If we tell the agent that the alarm sounded yesterday, it merely needs to remember this fact, and maybe derive some inferences from it. However, if we tell it that the fire-alarm is *now* sounding, it also needs to act on this fact and leave the building. Thus, distinguishing the present is important for correct behavior, not just for intelligible interaction and sound reasoning.

But distinguishing the present and modeling the progression of time comes with a host of intricate issues. For human agents, there is a continuous sense of the passage of time and it evidently poses no problems as far as reasoning, acting, and interacting are concerned. However, once it is decided that an artificial agent should be endowed with a sense of time, the intimate phenomenon of temporal progression with which humans peacefully live ceases to be simple and many of our familiar time-related experiences turn out to be more involved than we think. A number of the complications raised by modeling an agent's sense of time are discussed below; the point: *the passage of time affects everything*.

1.3.1 Representing "now"

For a logic-based cognitive agent, how is the passage of time to be modeled? To start with, to give the present a distinguished status, the agent should have some notion of "now". Deciding how this notion is to be represented is a very crucial matter. For example, one might choose to represent "now" in the logic itself. As an indexical, the interpretation of "now" is context-sensitive, and including it in the logic introduces context as an index of interpretation of the agent's knowledge base. One problem with this approach is that the agent's sense of temporal progression means that the context is always changing. In general, to ensure that sentences in the knowledge base mean what they are intended to mean regardless of the context, extensive manipulation of the contents of the knowledge base is inevitable. Such manipulations are not the normal expansion or querying of the knowledge base, but they necessarily involve retracting beliefs the agent once held. For example, at some point the agent might believe that "it is now raining". As time passes, such a belief would eventually become false and the knowledge base would have to be adjusted in order to reflect that it is no longer raining while allowing the agent to remember that it once was. From a computational-efficiency stand, such adjustments are typically bad news. But even if the adjustments required are not computationally very costly, the very idea of the agent *changing its mind* as a result of the intimate and (almost) continuous passage of time is at best awkward.

Instead of introducing "now" as a an indexical term in the logic, one might choose to represent it only meta-logically. In particular, at any point, there would be a term in the knowledge base which, in the meta-theory, is distinguished as denoting the current time (the meta-theory may span the inference engine, the acting executive, and the natural language understanding/generation system). Thus, that this term stands for "now" in a particular context is not something that is represented by a belief of the agent in the knowledge base. Rather, it is more of a perceptual phenomenon providing the agent with a subjective sense of time while all its beliefs are, formally, of the objective, context-independent species. As time passes by, a *new* term is introduced into the knowledge base and is meta-theoretically stamped as the new "now". Thus, the passage of time in this case leaves the knowledge base intact, without the need for retracting any beliefs. Only the meta-theory needs to be updated as time passes by, and this may be as simple as changing the value of a variable. This, however, comes with a price; since "now" is unmentionable in the knowledge base, representing and interpreting certain rules and plans of action gets to be more complicated. I will call this *the problem of the unmentionable now*.

1.3.2 Reasoning in Time

Acting agents reason and act in the real world. In the real world, everything takes time. For an agent maintaining a sense of time and its progression, things can get quite entangled. Why? Because the term which, for the agent, denotes the current time at the beginning of a reasoning process may be different from the one denoting the current time at its end. This by itself need not be problematic, but it gets to be a problem when time is the object of reasoning—when the agent is reasoning about "now" and the very process of reasoning results in "now" changing. For example, suppose that the agent wants to cross the street, and it believes that "If the walk-light is now on, I can cross the street". To determine whether the walk-light is "now" on, the agent turns its head to the walk-light and indeed sees that it is on. But this takes time, however short, resulting in "now" changing. Should the agent cross the street? The answer is obviously "yes", but the more interesting question is: Why? Why should the agent cross the street given that, strictly speaking, the walk-light is observed to be on at a "now" that is different from the "now" at which the agent wondered whether the walk-light is on? Why is this observation relevant to the agent's interest? Why isn't it too late, for example? Human agents do not seem to have any problems in similar situations. They somehow manage to allow some degree of tolerance when interpreting "now". But how may such looseness be accommodated by an artificial agent following the strict rules of logic? I will call this *the problem* *of the fleeting now*.

1.3.3 Persistence of States

The notorious frame problem in logic-based AI is concerned with what does *not* change as events happen. Traditionally, the problem is one of coming up with the appropriate logic and set of axioms to efficiently represent and reason about persistence in a dynamic domain. What is not traditional, however, is the situatedness of the reasoning agent itself in time—its sense of temporal progression as time passes by. For such an agent, problems of persistence of world states are not merely encountered when the agent needs to reason, from the outside, about a changing world, but whenever "now" moves. For every time "now" moves, the agent should be able to tell whether what was just the case continues to be the case. Most of the time, human agents do not seem to have any problems here, and, fortunately, *embodied* artificial agents should not either. Interestingly, the reason is the very situatedness of those agents in time. In particular, most of the states an agent is aware of are perceived states of the environment or its own body. Knowledge of such states is continuously supplied through bodily feedback and the agent seldom needs to *reason* about whether they continue to persist as time goes by. Even if the agent ceases to perceive a state, usually its very sense of time allows it to reliably determine whether it continues to hold. For example, on my way to the mail-room, I pass by the open door of the secretary's room and see her in there busy typing something. Once I am past the open door, do I have any beliefs about whether she is still there typing? Of course, and the reason is primarily my sense of time, particularly my sense of how much time has passed and my knowledge of how often different states change. Therefore, in order to endow an artificial agent with a sense of time akin to that which humans have, not only should we provide it with a notion of "now", but also with a *feel* for how much time has passed along with knowledge of the typical durations of various relevant states and a sense of how long these durations feel. How is that to be exactly achieved? And how may such primarily perceptual phenomena be smoothly woven into the logical framework adopted here?

1.3.4 Integrating Experience over Time

Except for what it is explicitly told, an agent's knowledge of the world is conveyed through perception, in the case of external states, and proprioception, in the case of states of the agent itself.

In fact, our knowledge of the immediate surroundings is typically conveyed through perception and our awareness of what we are doing, not through communicating with other agents. But perception and proprioception only tell us what is the case "now"; they only provide a narrow window onto the world that keeps on sliding with the passage of time presenting us with one temporal frame after another. This way of experiencing the world allows an agent to form a belief, and report, that it is crossing the street, or *was* at some point crossing the street, that it is now on one side of the street but was recently on another. This is possible since these are all *states* that the agent can experience at some "now". But the agent should also be capable of forming a belief, and reporting, that it *crossed the street*, for example. How is this possible, given that, at no "now" does the agent experience something to that effect? Note that just knowing that it *was crossing* the street is not enough for the agent to know that it *crossed* the street. It also needs to know how the crossing unfolded in time and, crucially, how it culminated. Thus, the agent needs to integrate its experience over time, to generate, from its experience of states over consecutive "now"s, an overarching *event* that is not wholly located at any particular "now". These issues are certainly tied to issues of linguistic aspect. However, the theory of aspect required for an agent reasoning and acting in time is different from the common theories of *linguistic* aspect; for it needs to explain how the occurrence of an event is to be inferred given only knowledge of states over time.

1.4 Scope of the Dissertation

In order to present the scope of this dissertation, I have to distinguish between the *wide scope* and the *narrow scope*. The first defines the general concerns and long-term goals of the research presented here. Not all these concerns are addressed nor are all the long-term goals achieved by the work discussed in this dissertation. The second, on the other hand, defines the issues that the dissertation explicitly addresses and proposes solutions to.

1.4.1 The Wide Scope

The long-term goal of this research is to develop a precise theory of an embodied cognitive agent that can do all of the following:

1. Reason.

2. Act.

3. Reason about its actions and execute them on-line.

4. Recover from errors and handle interrupts in a robust and reliable manner using its knowledge of the world. The agent should be able to reason about how, when, and whether to attend to an emergency situation and to be able to resume whatever processes it might have to suspend if it decides to attend to the emergency situation.

5. Maintain a personal sense of time including a notion of a continuously moving "now" and a feel for how much time has passed.

6. Interact with humans in natural language. Interactions may involve instructions from a human operator; descriptions of how to act, when to act, current and past states of the environment, and general knowledge about the domain; and reports by the agent about what it did and what it has come to know. Such interactions can take place while the agent is actually acting, providing it with information that may help it achieve its goals or recover from failure.

Not any theory would be satisfactory, however; there are requirements that it should satisfy. First, the theory should be practical enough to be feasibly implemented in useful applications. Second, and more important, the whole research project leading to and embodied in the theory should be more than an engineering endeavor; it should have its roots deeply entrenched in the cognitive sciences. In particular, not only should the theory have useful applications, but it should also help us gain a deeper understanding of human-level cognitive processes and situated reasoning and use of language. As such, the theory should be faithful to results in related disciplines, particularly linguistics, philosophy, and psychology.

Although the theory presented in the dissertation does not completely cover all of the above issues, they are all kept firmly in mind, featuring as general background constraints on the methodologies adopted and the framework within which the theory is conducted. For example, the knowledge representation language will be designed parallel the structure of natural language (English, in particular) as closely as possible.

1.4.2 The Narrow Scope

The narrow scope of the dissertation is essentially the issues that it explicitly addresses and proposes solutions to. In very general terms, these may be summed up in the following listing of the contributions of the dissertation:

1. An interval-based logic of time and aspect that accounts for what an agent has in mind as events unfold in time. More precisely, the logic allows the agent to infer the occurrence of an event given only its knowledge of the states it experiences as time goes by. In this respect, an ontology of event categories is developed based on the notion of telicity. The dissertation goes beyond the traditional binary telic/atelic distinction and precisely characterizes five classes of event categories each exhibiting different telic features.

2. A theory of subjective time that features as a meta-logical theory on top of the logic of time and aspect mentioned above. The theory includes an account of "now", how it is represented and how and when it changes. In addition, a precise analysis of belief update as a result of perception, proprioception, communication, and inference culminates in a set of algorithms accounting for the interaction of these cognitive processes with the passage of time.

3. A precise analysis of the problem of the unmentionable now (see Section 1.3.1) and the problem of the fleeting now (see Section 1.3.2) culminating in solutions to the two problems. The solution of the problem of the unmentionable now involves building the agent's subjective sense of time into the very deep reasoning processes. The solution of the problem of the fleeting now involves endowing the agent with a feel for how much time has passed, knowledge of the typical durations of various states, and a richer representation of "now" as a lattice of intervals.

4. An account of the persistence of states as time flows that makes use of the agent's sense of time, its knowledge of the typical durations of states, and the continuous access to perceived and bodily states facilitated by its embodiment.

5. A theory of sequential acting that interleaves reasoning and acting to allow the agent to consciously decide when to move on to the next step in a sequence of acts. The theory makes use

of the analysis of telicity mentioned in 1 above to characterize the completion conditions of different types of acts.

6. The beginnings of a theory of interrupt handling and error recovery. The theory, again, interleaves reasoning and acting, allowing the agent to reason about what to do in face of an emergency situation. The agent's decision is based on beliefs it has about the context-sensitive relative priorities of acts.

1.5 The Framework

Before delving into technical details, a brief discussion of the framework within which the analysis is pursued is due.

1.5.1 Cassie

The theory of agents adopted here is based on the GLAIR agent architecture (Hexmoor et al., 1993; Hexmoor, 1995; Hexmoor and Shapiro, 1997). GLAIR is a layered architecture consisting of three levels:

1. **The Knowledge Level (KL):** The level at which conscious reasoning takes place. Our KL is implemented by the SNePS knowledge representation, reasoning, and acting system (Shapiro and Rapaport, 1987; Shapiro and Rapaport, 1992; Shapiro and the SNePS Implementation Group, 1999).

2. **The Perceptuo-Motor Level (PML):** The level at which routines for carrying out primitive (or basic) acts are located. This is also the location for other subconscious activities that allow for the agent's consciousness of its body, surroundings, and the passage of time.

3. **The Sensori-Actuator Level (SAL):** The level controlling the operation of sensors and actuators (being either hardware or simulated). The reactive mechanisms discussed in Section 1.1 are implemented at the SAL.

"Cassie" is used as the name of the GLAIR/SNePS-based agent. The KL corresponds to Cassie's "mind", and the SNePS knowledge base is to be interpreted as the contents of Cassie's memory.

At this point, I should explicitly state that I have nothing much to add to discussions of consciousness in cognitive science and philosophy. Throughout the dissertation, expressions like "consciousness" and "awareness" refer to a level of representation of information that includes entities that an agent can reason or talk about, namely the GLAIR KL. Thus, an agent may be aware (or conscious) of whether it is holding a block in as much as it can reason or talk about such a situation. Nevertheless, it might not be aware, and hence cannot discuss or reason about, any of the perceptual and/or proprioceptual processes involved in believing that it holds a block (which, in GLAIR, would be located at the PML).

1.5.2 SNePS

It is important to stress that SNePS structures do not represent "the world" (although they can). That is, the representations are not based on theories of physics or even (objectivist) metaphysics. For example, ontological questions about the (un)reality of time (McTaggart, 1908) or the status of temporal individuals (Chisholm, 1990; Pianesi and Varzi, 1996a, for instance) might be important insofar as they shed more light on the phenomena under investigation. Nevertheless, answers to those questions are not sought and should not be a factor to consider when making decisions regarding SNePS representations. SNePS structures represent the "mind" of a cognitive agent. That is, rather than representing the world, they represent *conceptualizations* of the world. There are two points to note about this conceptualization.

1. As pointed out above, the SNePS knowledge base is to be interpreted as the contents of Cassie's memory. Thus, whatever conceptualization of the world it represents, it is a conceptualization of the world from Cassie's first-person point of view. More precisely, the logical theory embodied in the SNePS knowledge base is not an objective theory *about* Cassie and her experience in the world; rather, it is Cassie's subjective theory of the world and her experience therein.

2. Not any conceptualization will do; after all, the theory of relativity (or any other theory of physics for that matter) is just one possible conceptualization. What I mean is the kind of naive conceptualization that forms the basis of our linguistic expressions, the understanding of the world that underlies the way we *talk* about it. In many cases, such an understanding is

in obvious contradiction with the laws of physics (Hayes, 1985; Talmy, 1988, for example). Evidently, we talk about *times*; they are part of our understanding of the world. Indeed, just arguing that, say, temporal individuals are not real is by itself evidence for the reality of their mental representation. Therefore, our representations will include SNePS terms corresponding to temporal individuals.

Cassie's conceptualization of the world is represented by a set of propositions that she believes. Propositions are first-class entities in the SNePS ontology (Shapiro and Rapaport, 1987; Shapiro, 1993). Thus, SNePS-based languages do not include predicate symbols in the traditional sense—symbols denoting tuples of individuals. Rather, there are function symbols denoting functions whose range is the set of propositions. Similarly, logical connectives are interpreted as functions over the domain of propositions rather than that of truth-values (Shapiro, 1993).

1.5.3 Intensional Semantics

Given this cognitive interpretation, what do SNePS terms denote? They simply denote mental entities, in particular, Meinongian objects of thought (Rapaport, 1978; Shapiro and Rapaport, 1987; Shapiro and Rapaport, 1991). Mental entities are intensional; that is, their "… identity conditions *do* depend on their manner of representation" (Shapiro and Rapaport, 1987, original emphasis). Accordingly, there is a one-to-one correspondence between the set of SNePS terms and the domain of interpretation: no two distinct terms denote the same mental entity, and no two distinct mental entities are denoted by the same term.

Note that this one-to-one correspondence is based on the assumption that SNePS terms represent mental entities in the mind of a *single* cognitive agent, namely Cassie. If I were developing a multi-agent theory, the semantics would have to be given relative to an agent. In that case, the one-to-one correspondence between terms and their denotations might fail, since the same term may mean different things to different agents.[2]. Throughout the rest of the dissertation, I will be assuming that all the representations are relative to the agent Cassie. This would result in predicates that would normally involve an agent-argument to appear in my theory without the need for such an

[2]Actually, in a multi-agent setting, one would assume that mental entities in the mind of one agent cannot appear in the mind of another. Thus, the issue would be much more involved than simply relativizing interpretation to an agent; for how can a term that is exclusive to one agent be interpreted with respect to another?

argument—the agent Cassie being implicit. For example, in Chapter 4, a binary predicate symbol, `Cat`,[3] is introduced to relate event tokens to their event categories. Since categorization is an agent-dependent process, an agent argument would need to be added to `Cat` if the theory did not include Cassie as an implicit agent.

On a different note, given the intensional semantics, SNePS terms may denote non-existent entities. They may denote entities that Cassie conceives of but that do not exist in the world in any traditional sense. Of particular importance, Cassie may conceive of an event before it happens and reason about it even though it may never actually occur.

1.5.4 Knowledge Acquisition

There are various ways by which new propositions may be added to Cassie's mind:

1. **Direct Assertion:** Using either natural language (English) or various interfaces to SNePS (Shapiro and the SNePS Implementation Group, 1999), one may communicate information to Cassie.

2. **Inference:** Using inference, Cassie's memory may be expanded by deriving new beliefs from old ones.

3. **Bodily-Feedback:** The PML may communicate information to the KL as a result of perception or proprioception.

Of particular interest to the current study are propositions about *current* states of Cassie and the environment. In what follows, I make the following reasonable assumption: If perception, proprioception, or direct assertion results in Cassie's acquiring a new belief, then forward inference is initiated. Why is this assumption reasonable? An acting agent should always be alert, paying attention to its immediate surroundings and the states thereof in order to direct its actions and decide what to do next. Forward inference provides one way of modeling this state of high-alert; not only does Cassie accept what she is told or directly perceives, but she must also reason about it to see if there is more to the new information than its explicit content. In particular, she should determine whether

[3] Actually, `Cat` is a function symbol; see Chapter 3.

the new information allows her to infer something about the *current* state. This is particularly important in the case of perception, since perceived states are usually so specific in character that they may, themselves, not be very helpful. Ultimately, however, I believe that full-fledged forward inference may not be needed. For example, although it is reasonable for Cassie, having acquired a new belief, to infer everything that she can about current states, it seems that inferring things about, for instance, the far past is not as motivated. In addition, forward inference should also be restricted to those propositions that seem pertinent to what Cassie is currently doing or interested in. These are all issues that certainly need to be considered. Nevertheless, I defer their investigation to the future and uphold the assumption that acquiring new information *always* initiates forward inference.

1.5.5 SNeRE

Cassie acts in three situations: when instructed to do so, when asked to achieve a goal and she knows about a plan to achieve it, and when she becomes aware of some situation to which she should react. Acts are represented by special SNePS terms that Cassie may reason about and may also execute. SNeRE uses a network activation mechanism (Kumar, 1994) where activating an act term corresponds to executing a procedure associated with it. This dual nature of acts requires a specialized system interleaving inference and acting; SNeRE (the SNePS Rational Engine) (Kumar, 1994; Kumar and Shapiro, 1994a; Kumar and Shapiro, 1994b; Kumar, 1996) is the sub-system of SNePS responsible for this. SNeRE primarily controls the relationship between acting and inference, allowing for "acting in the service of inference and inference in the service of acting" (Kumar and Shapiro, 1994a). In particular, SNeRE allows Cassie to act as a result of inferring some new piece of information representing a situation to which she should react. On the other hand, SNeRE provides the capability for Cassie to act in order to add a missing link to a chain of reasoning; for example, by performing a sensory act.

1.6 Organization of the Dissertation

The dissertation consists of 11 chapters (including the current chapter) and four appendices. Chapter 2 is a review of the extensive literature on issues related to the work presented in the dissertation. It covers the AI literature on action and temporal reasoning, the linguistics literature on the logic of aspect, and the different approaches to the representation of "now" in AI and philosophy.

Chapters 3 and 4 lay out the logic of time and aspect based on which the investigation proceeds. Chapter 3 covers the logic of time and states whereas Chapter 4 is dedicated to the logic of events and the analysis of telicity. The axioms and theorems of these logics are compiled in Appendix A. Appendices B and C include the proofs of theorems from Chapters 3 and 4, respectively.

In Chapter 5, the meta-logical theory of subjective time is presented. The theory provides an account of temporal progression: how the passage of time is represented and how an agent's beliefs are updated as time goes by. Chapters 6 and 7 are dedicated to the problem of the unmentionable now and the problem of the fleeting now, respectively. Issues of persistence of states are discussed in Chapter 8. The temporal progression algorithms developed in those chapters are compiled in Appendix D.

Chapter 9 introduces the theory of sequential acting and the cascading mechanism that allows the agent to interleave reasoning and acting in order for it to determine when an act is complete so that it can move on to the next step in a sequence. In Chapter 10, the beginnings of a theory of interrupts and error recovery is laid out.

Chapter 11 discusses a current implementation of the theories presented in the dissertation. Sample runs are provided to demonstrate the system. Finally, the conclusions and future directions research are discussed in Chapter 12.

Chapter 2

Related Work

The research presented in this dissertation strongly overlaps a number of areas in cognitive science. In particular, it fits within, supplements, and derives insights from previous efforts in AI, linguistics, and philosophy. The review of related work is divided below into reviews of previous work on time, aspect and action.

2.1 Time in AI

In this section, ontologies and representations of time in logic and AI are examined. We start by drawing a distinction between two views of time that motivate different approaches in the field.

2.1.1 The Two Views of Time

To begin with, there are two main views that a theory of time may adopt. These may be called the *subjective* and the *objective* views of time. Historically, the distinction was precisely formalized by (McTaggart, 1908) in his *A series* and *B series*. Basically, the A series is the series of positions in time "running from the far past to the near past to the present, and then from the present to the near future and the far future" (McTaggart, 1908, p. 458). The B series is the "series of positions which runs from earlier to later" (McTaggart, 1908, p. 458). The A series involves a deictic representation of time, relative to a cognitive being, hence the subjectivity. On that view, the agent is itself located *in* time (metaphorically, the agent is *on* the time line), and maintains a sense of the present time. The B series is a totally objective view of time, independent of the experience of any conscious mind. On that view, the agent exists *out of time* and merely reasons about a static configuration of events. (Lakoff and Johnson, 1999, Ch. 10) present evidence for two versions of the A series: one with a moving time and a stationary *observer*, whose location represents the present, and one with a stationary time and a moving observer. For example, compare *"The deadline is approaching"* to *"We are approaching the deadline"*.

In a logic-based theory of time, the subjective-objective distinction may appear in one of two guises. First the logic itself may be a logic of objective or subjective time. Second, a meta-logical theory of subjective time may be wrapped around an objective logic of time (this is the approach adopted in this dissertation, as pointed out in Chapter 1). The classical distinction between logics of objective and subjective time is illustrated by Galton's distinction between the *first-order* and *modal* approaches to temporal logic (Galton, 1987b). In the first-order approach, time is explicitly represented by terms in the logic; it typically features as an extra argument of most predicates and some functions. This approach draws its popularity from its exclusive use in the representation of time by physicists. The modal approach, on the other hand, does not include any time-denoting terms in the logic; time is represented through the temporal *interpretation* of the modal operators (Hughes and Cresswell, 1996, ch. 7). In this approach, a possible world essentially represents a moment of time, and constraints on the accessibility relation between worlds determine various properties of the structure of time (in particular, whether it is branching, and, if it is, whether it branches only into the future, or into both the future and the past). Typically, sentences that do not fall within the scope of a modal operator are interpreted to be assertions about the present; hence the subjectivity.

2.1.2 Objective Time

In this section, we review some of the major systems of objective time in AI. It should be noted that, since all of the following are logical formalisms, one can always embed them within a meta-theory of *subjective* time.

The Situation Calculus

Perhaps one should start by discussing the first AI system with an account of time, namely the famous situation calculus (McCarthy and Hayes, 1969). The situation calculus was introduced mainly to provide a framework for planning and reasoning about action. As such, it is not concerned with language and issues of linguistic aspect that seem to be central to a general theory of time for reasoning, acting, and interacting cognitive agents. Nevertheless, the situation calculus has proved to be one of the most influential systems of time in AI.

The formalism revolves around three main notions: actions, fluents, and situations. Traditionally, the domain of a situation calculus theory contains a single agent that is responsible for all the actions performed. In fact, the only events occurring in the domain are actions of that agent. All actions are assumed to be duration-less; they are characterized, not by any internal structure (since they have none), but by their preconditions and effects. Preconditions and effects are represented by *propositional fluents*, which, intuitively, are propositions whose truth varies with time. This variation is formally captured by the introduction of situations. A situation "is the complete state of the universe at an instant of time" (McCarthy and Hayes, 1969, p. 477) and a propositional fluent is said to hold, or not, in any given situation. Formally, propositional fluents are functions from situations to truth values.

Perhaps the main innovation of the situation calculus is its formal representation of situation terms. There is a distinguished term (typically labeled S_0) representing the *unique* initial situation; any other situation results from performing an action. An expression $result(a,s)$ denotes the situation resulting from performing action a in situation s. Thus, any situation uniquely corresponds to a sequence of actions starting in the initial situation. The situation calculus, therefore, provides a view of time that branches into the future. It should be noted, however, that, technically, situations are not times. In fact, (McCarthy and Hayes, 1969) discuss a function that assigns a time to each situation. Time itself has a standard linear structure, and two situations may have the same time, corresponding to different ways in which the world may evolve. As pointed out by (McCarthy and Hayes, 1969) (also see (Steedman, 2000)), situations correspond to reified possible worlds of a temporally-interpreted modal logic. The accessibility relation between situations is established by the performance of actions. In particular, a situation s_1 is accessible from a situation s_2 if $s_1 = result(a,s_2)$, where a is an action that could be performed in s_2, i.e., an action whose preconditions hold in s_2.

A fluent holds in a given non-initial situation, s, if it is an effect of the action resulting in s or if it is continuing from a previous situation. Inferring the persistence of fluents from one situation to the next is the essence of the notorious frame problem pointed out in (McCarthy and Hayes, 1969) and the motivation of much research thereafter (Lifschitz, 1986; Shoham, 1986; Brown, 1987; Hanks and McDermott, 1987; Reiter, 1991; Morgenstern, 1996; Shanahan, 1997, to mention a few). The frame problem is primarily the problem of determining what it is that does *not* change as result of performing an action. Since most fluents do not change as a result of performing some particular action (even within a simple domain theory with a limited number of fluents), the number of axioms required to state the non-effects of actions could be ridiculously large. The frame problem manifests itself in this dissertation in a rather non-traditional guise. Relevant discussions are deferred to Chapter 5.

McDermott's Temporal Logic

Following the situation calculus, one of the earliest elaborated representations of time in AI is that of (McDermott, 1982). The basic sort of entity in McDermott's logic is what he calls a *state*. States are "instantaneous snapshots of the universe". They are densely structured in a tree of branching *chronicles*. "A chronicle is a complete possible history of the universe, a totally-ordered convex set of states extending infinitely in time" (McDermott, 1982, p. 106). It is crucial in McDermott's system not to confuse states with time points. Every state is associated with a time, its date, represented by a real number. Different states in different chronicles may have the same date, reflecting the state, at a particular time, of various possible worlds. It should be clear that, as far as the basic infrastructure goes, McDermott's system presents a view of time that is very close to that of the situation calculus. The situations of the situation calculus are the states of McDermott but with a discrete rather than a dense structure, and the time of a situation in the situation calculus corresponds to the date of a state in McDermott's system. In addition, the branching model of time is characteristic of both formalisms.

The main power of McDermott's logic over the situation calculus lies in a finer analysis of events. Starting with states, McDermott defines intervals as segments of chronicles, that is, totally-ordered convex sets of states.[1] An event comes out in McDermott's logic as the set of intervals over which it happens—an idea traditionally attributed to Montague (Montague, 1960, for example). There are a number of things to note here:

[1]Note that, technically, these are not *time* intervals.

1. Unlike the situation calculus, McDermott's events are not durationless and are not necessarily the actions of a single agent.

2. By "event", McDermott refers to an event *type* and, as a set of intervals, each member of an event corresponds to a particular *token*.

3. This "external" characterization of events, although helpful in reasoning about event occurrence, does not characterize an event by any inherent features. I cannot help quoting (Galton, 1987b, p. 23):

   > [A]n event is what it is because of its internal characteristics. The external characteristics may suffice for certain purposes, but can never amount to a definition, and indeed, the complete set of chronicles required to retrieve the internal characterization from the external one could not, in principle, be obtained unless we already had at our disposal the very internal characterization we are seeking to retrieve.

   In particular, as pointed out in Section 1.3.4, we need an agent to be able to recognize an event occurrence from its knowledge of states holding ("states", by the way, are "facts" for McDermott and are the propositional fluents of the situation calculus). McDermott's external characterization of events clearly does not suffice for this purpose.

In addition to events, McDermott presents rather detailed accounts of causality and continuous change—the latter, in particular, is not traditionally accounted for in the situation calculus. However, recently, (Pinto, 1994) has explicitly introduced a *time line* into the situation calculus. Pinto's proposal renders the situation calculus even more similar to McDermott's logic. He distinguishes an "actual path" of situations—that determined by sequences of actions actually performed in the real world. He extends his system to allow it to reason about continuous change by interpreting the time line as a subset of the linearly-ordered field of reals.

Perhaps one of the more relevant insights of (McDermott, 1982) (to this dissertation) is his formalization of the notion of persistence. McDermott's discussion of persistence will be carefully examined in Chapter 8, where the context is more appropriate. For now, however, let me end this section by briefly pointing out the two main ideas about persistence presented by McDermott that are particularly relevant to this dissertation.

1. We do not merely perceive states holding; we perceive them persisting. That is, when we perceive some state, we typically assume that it will persist for a while.

2. We may defeasibly assume that a state continues to persist as long as it has not exceeded a certain typical duration—what McDermott calls its *lifetime*.

Allen's Theory of Time

Either explicitly or implicitly, both the situation calculus and McDermott's logic take points of time to be the set of building blocks in the temporal ontology. James Allen, on the other hand, argues for an interval-based semantics for time (Allen, 1983; Allen, 1984).[2] He uses thirteen exclusive relations that represent all possible ways in which two intervals may be related. Vagueness or lack of information is represented by allowing arbitrary disjunctions of the core relations to describe the relation between two intervals. Allen's system maintains a network of temporal constraints, where nodes correspond to intervals and arc labels are relations (or disjunctions thereof) between them. As more information is acquired by the system, constraint propagation is used to restrict the possible relations between pairs of intervals.

(Allen, 1984) presents a tripartite classification of *situation types*, or things that happen or obtain in time. The distinction is based on how each situation type distributes over time. To understand the basis of Allen's distinctions, keep three notions in mind: *all*, *none*, and *some*. For Allen, a *property* holds over all sub-intervals of any interval over which it holds. Allen's properties correspond to what I, following the linguistic literature, call states (note that, so far, we have seen four labels for the same concept). An *event*, on the other hand, does not occur over any sub-interval of an interval over which it occurs (think of "I walked to the store"). This distinction between states and events is one that I adopt here; I do not distinguish Allen's third situation type at any deep logical or ontological level. For Allen, a *process* occurs over some sub-interval of any interval over which it occurs. Allen's example of a process is "I was walking", where one does not have to be actually walking throughout the interval referred to (for more on this, see Section 3.5.2). Allen introduces three different predicates (HOLDS, OCCUR, and OCCURRING, respectively) to axiomatize the different features of temporal incidence exhibited by the three situation types. It should be noted

[2]Also see (Shoham, 1985).

that Allen's system does not make any type-token distinction within the domain of situations; only types are represented. An event token, which he needs in his analysis of causality in (Allen, 1984), has to be represented by both an event (type) and an interval over which it occurs. The absence of event tokens might be tolerable but it actually poses problems for Allen's theory. In particular, given his axiomatization of OCCUR, if an event occurs over an interval, then it does not occur over any sub-interval thereof. This is fine as long as we are restricting events to a certain type (*telic* events, see below). For example, "I walked toward the store", which is intuitively an event, obviously violates the axioms of OCCUR.

Although I do not adopt Allen's tripartite analysis of situation types, the core logic of objective time presented in this dissertation (see Section 3.4) follows the spirit of his interval calculus. I will have more to say about Allen's system later.

Criticisms and Revisions

Although they have proved very influential for AI research on time, the systems of McDermott and Allen were criticized in a celebrated article by Yoav Shoham (Shoham, 1987). Shoham makes two main points. First, he argues that neither author gives a clear semantics for his logic,[3] and, second, he contends that the classification of situation types (or "propositional types" as Shoham would prefer to call them) by both McDermott and Allen are both unnecessary and insufficient. Shoham presents a logic where sentences are pairs of propositions and intervals. The intervals are the classical model-theoretic temporal indices of evaluation, represented explicitly in the syntax. It should be noted that intervals are not basic domain individuals, but rather, pairs of points. Indeed, Shoham criticizes Allen's dismissal of points from his temporal ontology and introduces them as temporal primitives.[4] He then presents a rich ontology of situations, in which the classical events, states, and processes are special cases. This classification was later picked up and exploited by students of Allen (Koomen, 1991).

Shoham's logic, however, did not stay clear of criticism. (Bacchus et al., 1991) describe a non-reified temporal logic, BTK, that, provably, subsumes the reified logic of Shoham. Propositions

[3]"[A]lthough McDermott does give the semantics for what may be regarded as the propositional theory" (Shoham, 1987, p. 94).

[4]Although the reasons for his criticism are based on the complexity of some of Allen's axioms using intervals rather than on any serious inadequacies introduced by the absence of points.

in BTK are made up of atomic predicates taking both temporal and non-temporal arguments. The number of temporal arguments is arbitrary (possibly 0) and specifies the dependence of the predicate on time. The terms of the language are partitioned into two sorts, temporal and non-temporal, thereby granting time a special syntactic and semantic status without resorting to reification of propositions. In addition to subsuming Shoham's logic, BTK has three advantages over it. First, propositions may be dependent on an arbitrary number of temporal entities (not only 2 as in the case of Shoham), thereby allowing the representation of atemporal facts. Second, being a standard non-reified logic, it has the readily-available proof theory of first-order predicate calculus at its disposal. Third, BTK does not make any ontological commitments to whether temporal individuals are points or intervals; this is left for axiomatization by users of the logic.

A particularly insightful examination of Allen's interval-based theory of time (or any interval-based theory, for that matter) is that of (Galton, 1990). Galton convincingly argues that dismissing points from the the ontology of time presents problems for an accurate account of continuous change. Galton's arguments rest on realizing that there are two major types of states: *states of position* and *states of motion*. A state of position can hold at isolated instants and must hold at the limits of any interval through which it holds, "a state which consists of some continuously variable quantity's assuming a particular value" (Galton, 1990, p. 169), for example, a body's being in a particular position (while moving). A state of motion, on the other hand, cannot hold at isolated instants, "a state of affairs which consists of some continuously variable quantity's either remaining fixed or undergoing a change of value" (Galton, 1990, p. 169), for example, a body's being at rest or in motion. Starting with Allen's system, Galton proceeds to introduce instants (which are neither parts of intervals nor members of sets representing intervals) and a variety of predicates assigning temporal location to events and states. By extending Allen's axiomatization, Galton is capable of overcoming the difficulties encountered with continuous change.[5] Since I am not concerned with representing continuity, the interval-based logic presented in Chapters 3 and 4 should be sufficient for our purposes.

[5]It should be noted that Allen has actually introduced points as dependent temporal entities in later work (Allen and Hayes, 1985). Nevertheless, there is no mention of how properties may (or may not) hold at points of time.

The Event Calculus

Another highly influential system for temporal reasoning is Kowalski and Sergot's calculus of events (Kowalski and Sergot, 1986). Like the situation calculus and McDermott's logic, temporal individuals are given a subsidiary status in the ontology. However, instead of some notion of "an instantaneous snapshot of the world", the calculus of events takes events as the basic primitives of the theory. An event divides time into two parts—the part before its occurrence, and the part after. Thus, when it occurs, an event may initiate a number of time intervals and terminate a number of other intervals. In the event calculus, an interval is characterized by a unique state (or *relationship* according to (Kowalski and Sergot, 1986), adding yet another label for the same notion) that holds over it. States are, therefore, representable only as effects of events. An important feature of the event calculus is how it solves the frame problem. Like many other proposals, the event calculus solves the frame problem non-monotonically. However, since it originated in a logic programming culture, the non-monotonic solution to the frame problem is transparently implemented using negation-as-failure and a closed-world assumption.

One can think of the structure of time implicit in the event calculus as that of the situation calculus, where an event correspond to an act, and an interval extending from one event to the next corresponds to a situation. As with the situation calculus, events are assumed to be durationless. There are some important differences, however:

1. Events in the event calculus need not be acts performed by agents.

2. Simultaneous events are possible.

3. There need not be a total temporal ordering of events.

4. Events may alternately be taken as types or tokens. In the latter case, a neo-Davidsonian description of events is employed (see (Sadri, 1987)).

(Shanahan, 1995; Shanahan, 1997) present a richer version of the calculus of events that may be used to represent continuous change. For a comparison of the event calculus to Allen's theory, see (Sadri, 1987).

2.1.3 Subjective Time

The work reviewed above takes the objective view of time as a basic underlying methodology. On that view, time is yet another domain phenomenon that a reasoning system needs to reason about. Granted, the very introduction of time and change has important impacts on many syntactic, semantic, and ontological commitments. For example, as the above discussion shows, the introduction of time raises questions about reification and the nature of temporal individuals—issues that have effects on the overall structure of a logic (see (van Benthem, 1983) and (Galton, 1995) for a more elaborate discussion). As outlined in Chapter 1, this research is concerned with providing the logical infra-structure for an acting agent that may reason and talk about what is happening and has happened. In a sense, the agent would have a narrative-like structure of events as they unfold in time. A crucial feature of such an agent is a sense of the present that would put it on the time line, providing it with a subjective view of time. Representing the present requires a logical account of the concept of "now".

Kaplan's Semantics of Demonstratives

One of the most common methods of accounting for the semantics of indexicals (of which "now" is a special case) is that outlined by David Kaplan (Kaplan, 1979). In an intensional Montagovian setting, Kaplan argues for a two-step interpretation function. The *character* of an expression is a function from contexts (of utterances) to *contents*, which are the traditional intensions (i.e., functions from possible worlds to appropriate domain entities, truth values, or set-theoretical structures over the domain). A competent speaker of English would recognize that "now" is used (loosely speaking) to refer to the current time. This knowledge is actually knowledge of the character of "now". The content of "now", on the other hand, is a particular time that does not by itself convey any concept of the present. The semantics of "now" is even more involved. As pointed out by many authors (Prior, 1968; Kamp, 1971; Cresswell, 1990), "now" always refers to the time of utterance even when embedded within a nest of tense operators.[6] This proves to be a very technically involved issue in the semantics of "now". These features of "now" pose unique problems for any logical account of the present. The main concern is what a logic should represent: the character or the content of

[6] An exception would be the "now-point" of a narrative (Almeida, 1995; ter Meulen, 1997).

"now". Of course, logics of objective time do not even need to consider such an issue; only contents are represented, since characters of demonstrative-free expressions are constant functions. On the other hand, systems that need to have a notion of the present have (at least) four choices.

Tense Logic

First, either implicitly or explicitly, the present is directly represented in the logic. This is typical of classical tense logics pioneered by Arthur Prior (Prior, 1967; Rescher and Garson, 1968; Kamp, 1971). Classical tense logic is essentially a temporally-interpreted modal logic with two, rather than one, modal operators, together with their duals. If p is a proposition, "Pp" means that "It has been the case that p" and "Fp" means that "It will be the case that p". By itself, "p" refers to the *current* truth of p. Thus, syntactically, the present truth of a proposition is distinguished by having the proposition outside the scope of any tense operators. Semantically, expressions (which may be embedded within tense operators) are interpreted with respect to a particular temporal index representing the present.[7] Although these logics are capable of distinguishing among the past, the present, and the future, they do not represent the passage of time. In a sense, they endorse an implicit objective view of time where the present is a static pivotal point from which one views the temporal configuration of events. An agent reasoning in the world needs to maintain a notion of the present that continuously changes as events occur. Evidently, traditional tense logic, by itself, does not suffice for such an agent.

Lespérance and Levesque's Logic of Indexical Knowledge

The second approach, usually adopted in reasoning about actions and plans, is to represent the present using an indexical **now** term. The use of indexical terms, in general, was studied in depth by (Lespérance and Levesque, 1995) with special attention to the case of **now** in (Lespérance and Levesque, 1994).[8] The indexicality of such a term stems from its having a context-dependent interpretation, much in the same spirit of Kaplan's semantics discussed above. However, unlike the English "now", whose content depends on the context of utterance (or assertion), the semantics of

[7] See (Kamp, 1971; van Benthem, 1983; Cresswell, 1990; van Benthem, 1995) for other approaches that use multiple temporal indices. A discussion of the two-index case is presented in Chapter 6.

[8] Other authors have also used the same or a similar approach (Schmiedel, 1990; Artale and Franconi, 1998; Dix et al., 2001).

the indexical **now** depends on the *evaluation* context. In the context of acting and planning, it is the time of executing a particular instance of a plan that includes occurrences of **now** in its specification. Such an approach facilitates certain kinds of reasoning that seem to require an explicit representation of indexical time (Lespérance and Levesque, 1994). Along the lines of (Lespérance and Levesque, 1994) (and using the same syntax), the following is a possible representation of a plan to get to the other side of a street (probably for a rather despondent agent):

$$\mathbf{if}(\mathbf{At}(\mathbf{now}, \textsc{WalkLightOn}), \textsc{Cross}, \mathbf{noOp})$$

This roughly says that, if, at that the time of performing the action, the walk-light is on, then cross the street; otherwise do nothing. What should be noted is that **now** in the above form does not refer to the time of introducing the form into the knowledge base, or to any other *fixed* time for that matter. It is, in a sense, a place-holder for any time at which the plan is performed.[9]

However, as represented, and used, in (Lespérance and Levesque, 1994; Lespérance and Levesque, 1995), such a concept of the present is certainly not the same one represented by the English "now". First, because its interpretation is not dependent on the time of assertion, there are uses of **now** that do not correspond to the *real* present. Second, temporal progression cannot be modeled by such an abstract characterization of the present. For example, it is intuitive to express the English *It is now raining* as, for example, $\mathbf{At}(\mathbf{now}, \textsc{Raining})$. Since **now** is not interpreted at the time of the assertion, the expression $\mathbf{At}(\mathbf{now}, \textsc{Raining})$ would essentially mean that it is always raining, since **now** always refers to the mere notion of the current time.

As for the first point, (Lespérance and Levesque, 1995, p.82) explicitly state that their **now** is not intended to represent the English "now". As for the second point, they briefly discuss a solution which I will now consider in some detail. The obvious approach to modeling the passage of time within the theory of (Lespérance and Levesque, 1995) would be to appropriately edit the knowledge base every time "now" changes in order to preserve the truth of its sentences. Thus, $\mathbf{At}(\mathbf{now}, \textsc{Raining})$ should be replaced by something more appropriate once "now" changes. One problem,

[9]On a first pass, it seems to me that the use of the indexical **now** is only *required* in certain descriptions of such future actions. In particular, it seems that *imperative*, rather than *declarative*, representation of plans are the ones that are inherently indexical (see (Huntley, 1984) for a linguistic perspective.) By slightly modifying their language (basically, by adding an extra temporal argument to some of the predicates), I have managed to rewrite all formulas in (Lespérance and Levesque, 1994) without mentioning **now**. In fact, the discussion of (Perry, 1979), who is a major proponent of indexical beliefs (and is cited by (Lespérance and Levesque, 1994; Lespérance and Levesque, 1995)), is mainly about beliefs required for explaining behavior (see (Millikan, 1990) for a critical examination of Perry's position). More examination of when exactly indexicals are required is needed.

of course, is that such updates might be computationally expensive. To get around the problem, (Lespérance and Levesque, 1995, p. 101) suggest that "if all occurrences of 'now' are replaced by a new constant and the fact that this new constant is equal to 'now' is added, then only this single assertion need be updated as time passes." This indeed eliminates the problem of expensive belief update and provides a neat logical and computational account of "now". However, from a cognitive perspective, I find the very idea of erasing sentences from an agent's mind as time passes by far from natural. If such sentences represent beliefs that the agent once held, where do they go, and how come the agent would have no memory of them once time passes? Note that this cannot be explained away as a matter of forgetting, for forgetting is not that selective to always affect beliefs involving "now", nor is it vigorous enough to take effect with every tick of the clock. The only way to explain this mysterious disappearance of beliefs is by arguing that they exist at a lower level of consciousness with respect to other beliefs. If this were the case, why are such beliefs part of the logical theory (which I take to be representing conscious beliefs of the agent)? Thus, I do not see a way of reconciling the notion of "now" as a term in the agent's language of conscious thought with a cognitively plausible account of the passage of time. Note that I'm not trying to refute Perry's (or, for that matter, Lespérance and Levesque's) arguments for indexicality (Perry, 1979). The theory presented in this dissertation incorporates indexical *thinking*, but not indexical *thoughts*. That is, the reasoning process is indeed sensitive to indexicality, but the object language does not include any indexical terms.

Active Logic

The third approach to represent "now" is to do it indirectly, by means of a *Now* predicate, where the expression $Now(i)$ means that the current time is represented by the term i. This is exactly the method adopted in *active logic*, originally known as *step logic* (Elgot-Drapkin and Perlis, 1990; Perlis et al., 1991). Active logic is perhaps the only system that takes the issue of reasoning *in* time (in the sense of Chapter 1) seriously. Temporal individuals are represented by integers, with the usual numerical order implicitly representing chronological order. Thus, active logic generally adopts a discrete model of time points, though, in (Nirkhe et al., 1997), a point-based representation of intervals is used for reasoning about durative actions and events. Active logic is intended to model an agent's reasoning about its own reasoning processes in time. As such, it is not primarily interested in reasoning *about* time. Thus, one does not find any detailed analyses of different situation types (events, states, processes, etc.) within the active logic tradition. This is one of the differences between active logic and the system presented in this dissertation. Another difference is the finer level of granularity at which active logic represents the passage of time.

In active logic, time moves with every inference step. This movement of time is represented both logically and meta-logically. Logically, this is achieved by a special inference rule that essentially replaces $Now(i)$ by $Now(i+1)$. Meta-logically, assertions are associated with the step, i, of inference at which they were asserted. By having special "inheritance" (inference) rules, assertions made at step i may be rederived at step $i+1$, thereby allowing persistence of information. The system can reason about the present, since, at any step i, $Now(i)$ is asserted. In addition, beliefs about the past could be carried on to the present by yet another special rule of inference allowing the system to derive $Know(\alpha, i)$ from α being derived at step i.[10] Indeed, active logic proves very powerful in a number of applications (Gurney et al., 1997; Nirkhe et al., 1997).

However, the *Now* predicate, though well-suited for the kind of reasoning problems addressed within active logic, does not exactly reflect the behavior of the English "now". This is apparent in the special rules that are tailored just to get the appropriate inferences concerning the current time. Most notably, *Now(i)* is a special case that the general inheritance rule cannot apply to; that is, *Now(i)* cannot be carried on to step $i+1$. In addition, active logic would have to be supplemented with more special rules to correctly deal with occurrences of *Now(i)* within opaque contexts. Another problem with active logic is its use of integers to represent time. Apparently, the use of integers facilitates the expression of some crucial rules of inference (also the *counting* of reasoning steps (Nirkhe et al., 1997)) that depend on having a well-defined notion of the *next moment of time*, represented by the integer successor operator. However, such a representation forces a certain degree of rigidity on the kind of knowledge that may be entered into the system. For example, there is no way to assert at step $i+m$ ($m > 1$) that a certain event e_2 occurred between events e_1 and e_3 that happened at times i and $i+1$, respectively. In other words, once "now" moves, there is no way to go back and create arbitrary past temporal locations. This is definitely a big drawback if the system is to be used in

[10]No distinction is made between belief and knowledge, so $Know(\alpha, i)$ should actually be read as $Believe(\alpha, i)$. This is intended to mean that, at step i, the system believed α. This is a piece of knowledge that should always be believed by the system, even if, at some later step, α is no longer believed. However, it is not clear how the system may believe that, at some past time, some state held.

general interactions, where assertions need not only be about the present.

The SNePS Approach

The final approach for representing the present is that of (Almeida, 1987; Almeida, 1995) and (Shapiro, 1998). In this approach, only non-indexical terms exist in the logical language; the notion of "now" exists only in the meta-theory. This is achieved by having a meta-logical variable, NOW, that assumes values from amongst the time-denoting terms in the logic.[11] Use of the NOW variable is restricted to the natural-language interface and the reasoning and acting system; it is always replaced by its value (the content) when interacting with the knowledge base. The character of the English "now" is captured by the algorithms that relate "now"-tokens to the NOW variable. The passage of time is modeled by introducing a new term, representing the (new) current time, and making it the new value of NOW (see Chapter 5 for more details). Note that this is very similar to the suggestion of (Lespérance and Levesque, 1995) discussed above. The difference is that, instead of equating the new term to an indexical **now** with an assertion in the object language, it is meta-logically assigned to a variable.

This approach overcomes all the difficulties outlined above, since the logic is a standard demonstrative-free logic. On the other hand, it poses certain problems in cases where the general notion of the current time needs to be represented in the logic (see Chapter 6). In addition, a general difficulty with reasoning about "now" is its granular ambiguity—the fact that, at any time, there is a nest of "now"s representing the agent's notion of the current time at different levels of granularity. This is evident in the slightly different intuitions one may have regarding the exact meaning of "now" in sentences such as *"I am now in my office"*, *"I now exercise every day"*, and *"I am now working on my PhD"*. This issue has not been addressed in the research reviewed above, although it poses real problems for acting agents (see Chapter 7). Moreover, research on temporal granularity (Hobbs, 1985; Habel, 1994; Euzanet, 1995; Pianesi and Varzi, 1996b; Mani, 1998), though quite detailed in some cases, only provides insights into the issue; it does not present directly implementable computational solutions. Those models of granularity that are directly implementable (particularly in temporal databases) are concerned with representations of calendars, not the ambiguity of "now" within the context of an acting agent (Bettini et al., 1996; Cukierman and Delgrande, 1998; Bettini

[11] A similar approach was suggested in (Allen, 1983).

and De Sibi, 1999, for instance).

2.2 Aspectual Matters

Although the representation of time is, by itself, an interesting research issue, representing time is important only because agents need to reason about temporal creatures such as events and states. Such entities will henceforth be collectively referred to as *situations*. The study of situations investigates two main points: the classification of situation types and the reasoning and representational issues involved in a logical theory of situations. Major work in this area comes from investigations of linguistic aspect and philosophy of language (though we have touched upon these issues in our discussion of objective time above). A careful examination of the relevant literature and how other approaches relate to mine will be presented as I discuss my own proposal in Chapters 3 and 4. In this section, I briefly review the general concerns of research on linguistic aspect.

2.2.1 Classifications of Situations

In a very influential paper, (Vendler, 1957) pointed out that different verbs presuppose a certain temporal structure of the situations they describe. Although (Vendler, 1957) consistently refers to verbs, the examples that he cites obviously involve entire verb phrases. The effects of verb arguments, grammatical aspect, time adverbials, and, in fact, all constituents of a sentence on the situation type have been studied by a number of authors (Dowty, 1977; Declerck, 1979; Verkuyl, 1989; Krifka, 1989, for instance). Vendler classified situations into four classes:

1. **Activities:** run, push (a cart).
2. **Accomplishments:** run a mile, draw a circle.
3. **States:** know, love.
4. **Achievements:** find, recognize, spot.

Vendler provides various linguistic tests to distinguish among the four types. Intuitively, activities are situations extended indefinitely in time. Like activities, accomplishments occur over stretches of time. However, an accomplishment has a built-in *climax* marking the end of the situation—a property that has been referred to as *telicity* in the literature (Garey, 1957; Comrie, 1976; Dahl, 1981).

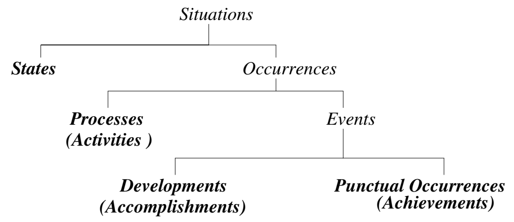

Figure 2.1: Mourelatos' classification of situation types.

States represent static situations, involving no change, that may persist over time periods. Finally, achievements are instantaneous situations involving no processes occurring over time. (Mourelatos, 1978) criticizes Vendler's system for its rather flat nature and argues for one with a richer internal structure (see Figure 2.1). What Mourelatos's account amounts to is the highlighting of generalizations that were overlooked by Vendler. Actually, however, Vendler's original analysis does not point to a flat structure, but rather a matrix, as pointed out by (Verkuyl, 1989, p. 44). (Bach, 1986) provides yet another more elaborate classification (see Figure 2.2). The overall structure of his classification is similar to that of Mourelatos, except that he distinguishes two types of states and achievements. A *culmination* is essentially the climax of an accomplishment (for example, *find, reach the top*), whereas a *happening* is a mere punctual event (*recognize, spot, notice*).[12] Vendler's and Mourelatos' typical examples of states are what Bach refers to as static states; dynamic states are represented by verbs such as *sit, stand,* and *lie*.

Given Bach's branching structure, there should be five dimensions along which distinctions among situation types are made. Some of these are obvious, some are less than obvious, and others are not unproblematic. The obvious ones are the durative-punctual distinction between protracted and momentaneous events, and the atelic-telic distinction (the absence/existence of a built-in ending state) between happenings and culminations (also between processes and protracted events). The basic distinction between states and non-states is that the former do not involve change, whereas the latter do. Finally, the difference between the two types of states (static and dynamic) is not very clear; one obvious difference is that verb phrases representing dynamic states can occur in the progressive whereas those representing static states cannot. What this exactly implies is not

[12] In the terminology of (Smith, 1997), culminations are "achievements" and happenings are "semelfactives".

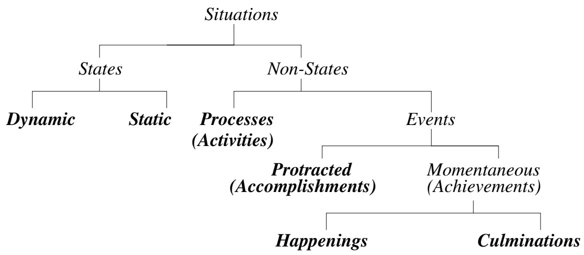

Figure 2.2: Bach's classification of situation types.

obvious.[13]

Another important classification of situations is that of (Moens and Steedman, 1988). Figure 2.3 depicts their entire ontology (this is adapted from Figure 1 of (Moens and Steedman, 1988)). The reader should note that Moens and Steedman recognize exactly the same distinctions as (Bach, 1986). The difference is that their matrix-, rather than tree-, structure mirrors how their situation types are related to one another. What is particularly insightful about Moens and Steedman's work is, not their ontology, but their analysis of the *internal* structure of events. They consider each event to be made up of three main components (note that this analysis does not apply to states): a preparatory process, a culmination, and a consequent state. This complex is what they call an *event nucleus*. To take an example from (Moens and Steedman, 1988, p. 19), the event of climbing Mount Everest consists of the preparatory process of climbing, the culmination of reaching the summit of Mount Everest, and the consequent state of *having* climbed Mount Everest. Using this tripartite analysis, Moens and Steedman are able to account for various phenomena in the temporal semantics of English, including the progressive, the perfect, time adverbials, and *when*-clause. The system presented in this dissertation includes an analysis of the structure of events that is somewhat similar (except at the level of details) to that of (Moens and Steedman, 1988). However, the discussion of linguistic phenomena in the dissertation is kept to a minimum, unlike the detailed analysis of Moens and Steedman. On the other hand, the system presented here makes distinctions that are not covered by their ontology and that do not seem to be captured by their nucleus-based analysis.

[13]Other authors provide similar classifications in terms of binary features (Gabbay and Moravcsik, 1979; Verkuyl, 1989; Smith, 1997). (Pustejovsky, 1991) presents a fine-grained structural characterization within what he calls the *event structure* of lexical verb entries. Also see (Talmy, 2000) for a different approach to the classification of situation types within the tradition of cognitive semantics.

	EVENTS		STATES
	atomic	extended	
+conseq	**CULMINATION**	**CULMINATED PROCESS**	understand
	recognize, spot, win the race	build a house, eat a sandwich	love, know
−conseq	**POINT**	**PROCESS**	resemble
	hiccup, tap, wink	run, swim, walk, play the piano	

Figure 2.3: Moens and Steedman's classification of situation types.

2.2.2 Aspectual Shifts and the "Two-Component Theory of Aspect"

Aspectual shifts among situation types have been studied by a number of authors (Moens and Steedman, 1988; Jackendoff, 1991; De Swart, 1998, for instance). For example, the English progressive shifts accomplishments into states. Thus, although 2.2 follows from 2.1, 2.1 does not follow from 2.2:

(2.1) John walked to school.

(2.2) John was walking to school.

This phenomenon has been dubbed "the imperfective paradox" by Dowty. Dowty reasons it is a paradox, since

> the meaning of an accomplishment verb phrase invariably involves the coming about of a particular state of affairs... Yet it is just this entailment that such a result-state comes about which fails when the accomplishment verb phrase appears in the progressive tense (Dowty, 1977, pp. 45–46, see fn. 2 therein for the use of "progressive tense").

However, a number of authors (Declerck, 1989; Depraetere, 1995; Smith, 1997, for instance) propose that nothing is paradoxical about the imperfective paradox. Those are proponents of what (Smith, 1997) calls the "two-component theory" of aspect. Basically, the theory calls for a distinction between telicity and mere boundedness (according to (Depraetere, 1995)), or the *situation type* (aktionsart, inner aspect) a sentence presents and the *viewpoint* it adopts (according to (Smith, 1997)). For example, (Depraetere, 1995) argues that although (1) is bounded and (2) is unbounded, they are both telic. According to this view, nothing is paradoxical about (1) and (2) since the overall aspectual character of a sentence is determined by two independent components, not just the core situation type. As it turns out, many of the important logical properties of situation descriptions originate from the notion of boundedness. In particular, unbounded expressions are (temporally) homogeneous, while bounded expressions are (temporally) heterogeneous. In fact, some authors (Galton, 1984; Herweg, 1991a) have only states (temporally unbounded situations) and events (temporally bounded situations) as the two main situation types (telicity and other properties may independently apply to states or events).[14] This reflects the long-noted symmetry between states (or processes) and events on one hand, and objects and matter on the other (Mourelatos, 1978; Bach, 1986; Jackendoff, 1991). Indeed, (Bach, 1986; Krifka, 1989; Schein, 1993; Krifka, 1998; Link, 1998) present a lattice-theoretic mereology of situations, where an event lattice and a process lattice are related by *grinding* and *packaging* operations.

2.2.3 The Logic of Situations

Careful examination of the above classifications is important insofar as it helps one understand linguistic data and inferences. To actually carry out inferences about situations, one needs to construct a logic. As (Herweg, 1991a) points out, there have been two main approaches to the logic of aspect: "the proposition-based approach" and "the eventuality-based approach".

The basic idea behind the proposition-based approach is that situations are properties of time intervals.[15] States and events correspond to propositions with different properties as to how their truth distributes over time. In particular, state propositions have the subinterval property (if true over some interval I, then they are true over all subintervals of I), whereas event propositions have the anti-subinterval property (they are false over all proper subintervals of an interval over which they are true). This corresponds to the homogeneity-heterogeneity distinction between states and events. As (Herweg, 1991a) (also (Galton, 1984, ch. 1)) points out, the proposition-based approach, though well-suited for the representation of states (which share a lot of features with propositions), is not suitable for representing events. In addition, it suffers from the problems that motivated Davidson's seminal work on the logical form of action sentences (Davidson, 1967).

The eventuality-based approach is based on the idea that the universe of discourse includes a

[14]Both Galton and Herweg view processes as forming a subcategory of states ("states of change"). Other authors call for a different analysis along the traditional Vendlerian distinction between states and activities (Taylor, 1977; Gabbay and Moravcsik, 1979; Smith, 1997; Smith, 1999, for instance).

[15]Typical examples of this approach in AI are (Bacchus et al., 1991) and the situation calculus with non-reified fluents (McCarthy and Hayes, 1969).

domain of situations (eventualities).[16] This probably dates back to (Reichenbach, 1947) and (Davidson, 1967). Typically, along the lines of (Bach, 1986; Link, 1998), the domain is partitioned into a set of events and a set of *quantities* of states. Those are accordingly mereologically structured in two separate lattices. The logical language includes eventuality predicates that are either stative or eventive (typically, complex lambda expressions). The homogeneity of the former, and heterogeneity of the latter, is captured by specifying how those predicates apply to an eventuality and its parts (descendents in the appropriate lattice). (Herweg, 1991a) argues that the main problem with this approach is that it treats states as individuals on a par with events.[17] (Galton, 1984; Galton, 1987a) (followed by (Herweg, 1991a)) present a hybrid system where states are represented following the proposition-based approach, whereas events are presented using the eventuality-based approach. In the system presented here, I will be, more or less, following the lead of Galton and Herweg.

2.2.4 Discussion

Primary results in the typology and logic of situations come mainly from the study of linguistic aspect and the philosophy of language. Research within AI mostly borrows readily available ideas from these two domains (with some exceptions (Shoham, 1987; Schubert and Hwang, 2000; Schubert, 2000, for instance)). These ideas, as should be expected, are developed with language as the only concern. An embodied agent reasoning and acting in time not only *talks* about situations, but also observes them as they unfold in time. The agent perceives states of its environment and its body, senses events happening in the world, and, most importantly, is often busy carrying out processes. What is needed is a two-pronged theory of situations, one that would (i) account for what the agent has in mind while acting or observing events as they unfold in time, and (ii) allow the agent to understand, and generate, natural-language utterances describing the state of the environment and its own actions.

[16] In AI, typical examples are (Allen, 1984; Schubert, 2000).

[17] This is not exactly precise: the problem is not treating states as individuals; the problem is assuming the existence of a state type-token distinction. One may safely assume the existence of state individuals, so long as they are conceived of as types; it is the concept of a state token that runs into problems.

2.3 Action

In this Section, a general review of the AI literature on action is presented. I will have more to say about the action literature in Sections 9.2 and 10.1 as particular problems are discussed. Research on action may be broadly divided into two categories reflecting two different concerns.

2.3.1 Reasoning about Action

The main concern of research on reasoning about action is the development of theories and systems for reasoning about actions and the changes they produce in the world. The basis of these systems is some logic that is used to describe actions, their preconditions, their effects, and states of the environment. Typically, the representation of actions is based on a kernel of primitive (basic, atomic) actions and various control structures to generate composite (complex, compound) actions from the primitive kernel. Control structures are often an extension of those found in programming languages, typically including sequential, conditional, and iterative structures. Different systems focus on different aspects of reasoning about action and hence employ different suites of control structures and different notions of primitiveness.

The situation calculus (McCarthy and Hayes, 1969) has probably been the most influential in shaping, motivating, and pointing out problems in reasoning about action. Based on the situation calculus, researchers at the University of Toronto have been developing a theoretical framework for cognitive robotics. One important outcome of these efforts is GOLOG, a high-level logic-programming language for reasoning about dynamic domains (Levesque et al., 1997).[18] The primary use of GOLOG is in reasoning about action; given a certain description of the world and an under-specified GOLOG program to achieve a particular goal, the interpreter produces a sequence of acts (or a set of alternative sequences) that, should they be performed, would achieve the goal. In addition to standard GOLOG, there are also CONGOLOG (De Giacomo et al., 2000) and temporal GOLOG (Reiter, 1998). CONGOLOG is equipped with control structures that are used to represent concurrent actions. The language may be used for reasoning about concurrent processes and exogenous events. Concurrency is modeled by the temporal interleaving of multiple processes according to some *fixed* system of priorities (more on this in Section 10.1). Temporal GOLOG is

[18] "GOLOG" stands for "alGOL in LOGic".

a slight mutation of standard GOLOG that is used for reasoning about continuous processes and actions with durational and chronological ordering constraints.

One of the characteristic features of the theory of GOLOG is its use of *successor state axioms* as a monotonic solution to the frame problem (Reiter, 1991). However, (Thielscher, 1999) argues that successor state axioms only solve the *representational frame problem*, and he proposes *state update axioms* (within his *Fluent Calculus*) as a solution to the *inferential frame problem*. According to (Thielscher, 1999, p. 280), the representational frame problem "is concerned with the proliferation of all the many frame axioms" about which fluents do not change as actions are performed. On the other hand, the inferential frame problem is concerned with the computational inefficiency with having to infer the persistence of fluents from one situation to the next using different instances of the, already many, frame axioms.

Based on Kowalski and Sergot's calculus of events (Kowalski and Sergot, 1986), Shanahan introduces a circumscriptive logic for reasoning about action (Shanahan, 1995; Shanahan, 1997). Shanahan's logic does not contain control structures to construct composite acts from primitive ones; rather, it depends on explicit temporal constraints for the representation of sequences of acts (see (Shanahan, 1998)). The logic can be used to represent and reason about concurrency, non-determinism, and continuous change. The frame problem is solved non-monotonically, using circumscription with prioritized minimization (McCarthy, 1980; Lifschitz, 1994). The temporal projection problem of (Hanks and McDermott, 1987) is avoided by designing the logic in such a way that minimization may be restricted to time-independent predicates.[19]

In a seminal paper, (Moore, 1988) directed the attention of the AI community to the inevitable interdependence of action and knowledge. Basically, Moore argues that performing an action, not only has physical preconditions, but also has *knowledge preconditions* (also see (Davis, 1994)). To take an example from (Moore, 1988), for an agent to be able to open a safe, it should be able to dial its correct combination. But for it to be able to do that, it should *know* this combination. Moore goes on to argue that *most* actions have knowledge preconditions, though these might not be as obvious as in the case of the safe. For example, for an agent to be able to put block A on

[19]Temporal projection has to do with projecting what facts will be true given a partial description of successive states of the world and a full description of events and their effects. Hanks and McDermott pointed out that using non-monotonic solutions to the frame problem may result in unintuitive models for temporal projection—those that are constructed starting with facts holding at later times and working back to the initial state.

block B (a standard task in AI), it should know which block is block A, which one is block B, and *where* the two blocks are located. On the other hand, (Moore, 1988) also argues that actions often have *effects* on an agent's knowledge. A typical example (mentioned by (Moore, 1988)) is the action of dipping a litmus paper in a solution which has the effect of the agent's knowing its pH (but note that, for that effect to be achieved, the agent must *know* the significance of the litmus paper test). The intimate interaction between knowledge and action is captured by a first-order logic with reified possible worlds representing knowledge states of the agent (i.e., it describes the semantics of a modal logic of knowledge). In this dissertation, I will consider some issues of the interaction of knowledge and action. This, however, is not the main theme of the dissertation and, hence, the analysis will not be as detailed as that of Moore. In particular, the system presented here does not represent an agent's beliefs about its own beliefs. It should be noted, however, that this may be readily accommodated within the SNePS framework and without the need to resort to modality or reified possible worlds. This is because, in SNePS, propositions (representing the agent's beliefs (Shapiro, 1993)) are individuals in the domain that may be reasoned *about* (see Chapter 3).

There have been several other logics of action, each tuned toward certain considerations. For example, (Davis, 1992) constructs an action language for representing interrupted and aborted acts (more on this in Sections 9.2 and 10.1). (Traverso and Spalazzi, 1995) present a logic for reasoning about actions, plans, and, especially, sensing. The control structures are tailored to facilitate explicit reasoning about failure. Within the description logic tradition, (Artale and Franconi, 1998) present a temporal logic for reasoning about acts and plans. As with Shanahan, there are no control structures for defining composite acts, but rather, explicit temporal constraints among acts (more on this below). (De Eugenio, 1998) also uses a description logic for representing actions and plans. Her main concern is interpreting complex natural language instructions (specifically, *purpose clauses*), a concern that lies behind her use of Jackendoff's theory of conceptual semantics (Jackendoff, 1990) to motivate the design of the representation language. Using a process algebra approach, (Chen and De Giacomo, 1999) describe a logic for reasoning about nondeterministic actions and concurrent processes performed by multiple agents. (Ortiz, 1999) develops a language for counterfactual reasoning about rational action, accidents and mistakes, attempts and failures, and various force-dynamic (Talmy, 1988) notions such as causation, enabling, letting, refraining, and hindering. In addition, there is the huge literature on planning which, though not the focus of this research, may be considered a typical example of reasoning about action. For a collection of classical work on planning, see (Allen et al., 1990).

2.3.2 Action Execution

Systems for reasoning about action might very well be used for pure reasoning about domains where actions are performed; the system itself does not include an acting component. Typically, however, the reasoning system is just part of a bigger system that includes an acting executive. In many cases, the interaction between the two components is minimal. In particular, reasoning takes place *off-line* (see Section 1.2), and its result is (typically) a sequence of acts to be executed by the acting system (Levesque et al., 1997; Artale and Franconi, 1998, for example).

On-line execution essentially involves temporal interleaving of reasoning and action. (Georgeff and Lansky, 1987) describe the Procedural Reasoning System (PRS), a planning system used for controlling a robot operating in a space station with navigation and emergency handling requirements. The knowledge operated upon by the reasoning system is divided into a data base of domain knowledge, representing the *beliefs* of the robot, and a set of declarative procedure specifications, called *knowledge areas* (KAs), representing knowledge about how to achieve goals or react to emergency situations.[20] In addition, PRS maintains two stacks: a goal stack and an intention stack. Goals represent desired behaviors of the system represented in a language combining ordinary logical connectives and some exotic temporal operators, and intentions are simply KAs. Whenever a goal is pushed onto the goal stack (either reflecting a *desire* of the robot or an emergency situation in the environment), PRS finds an appropriate KA and pushes it onto the intention stack. Executing a KA may result in beliefs being added to the data base or goals being pushed onto the goal stack. The robot expands its plans dynamically and incrementally, thus having the chance to interrupt the execution of a certain task in order to attend to an emergency situation.

Another system that interleaves acting and planning is IPEM (Integrated Planning, Execution and Monitoring) (Ambros-Ingerson and Steel, 1988). IPEM is based on a production-system architecture in which IF-THEN rules correspond to *flaws* and their corresponding *fixes*. Briefly, a flaw is a property or condition in a partial plan, and a fix is a plan transformation that would fix it. An execution monitoring mechanism makes use of the flaw-fix information to adjust the plan should

[20]Note that KAs are not considered part of the beliefs of the robot.

something go wrong during execution.

GOLOG-based systems come in two forms, those with off-line interpretation (Levesque et al., 1997; Reiter, 1998; De Giacomo et al., 2000, discussed in the previous section) and those with on-line interpretation (De Giacomo and Levesque, 1998; De Giacomo et al., 1998; Lespérance et al., 1998). Interpreting a GOLOG program off-line results in a sequence of primitive actions to be executed, unconsciously, by a (real or simulated) robot. More precisely, given a GOLOG program, the system "must find a sequence of actions constituting an entire legal execution of the program *before* it actually executes any of them in the world" (De Giacomo and Levesque, 1998, p. 1; original emphasis). With on-line interpretation, the system physically executes actions in sequence as dictated by a GOLOG program. Such programs may involve non-deterministic choice among acts, and the system needs to pick one of them and commit to it. Since, in general, one cannot undo an action, an on-line execution of a GOLOG program may turn out to be illegal. That is, it may not achieve its intended goal. However, this kind of incremental interpretation allows for conditional constructs depending on the execution of sensory acts whose outcome cannot be predicted off-line (De Giacomo and Levesque, 1998). It also allows execution monitoring so that, after executing a primitive act, the monitor checks whether an exogenous event has occurred and, if so, whether it needs to execute some recovery procedure to ensure achieving the goal (De Giacomo et al., 1998).

Within an abductive framework, (Shanahan, 1998) uses his calculus of events (Shanahan, 1995) for modeling high level cognitive processes required for the control of a robot. He uses a *sense-plan-act* cycle to interleave sensing, planning, and acting with the hope of achieving more reactivity. This kind of interleaving of acting, sensing, and planning allows the agent to use sensory data to guide its actions. In particular, Shanahan points out that, in case there is any discrepancy between the actual state of the world and the representation of it used to generate a plan, sensory input may help the agent revise the plan appropriately.

2.3.3 Discussion

The research presented above covers various topics related to acting systems. Although the issues addressed are diverse and important, none of the above systems totally covers all the areas that this dissertation is concerned with (as outlined in Chapter 1). In particular, all of the above systems fail to support two basic concerns. First, the knowledge of the agent in all of these systems is static. That is, the agent's knowledge base contains only background domain knowledge and plans for performing acts and/or achieving goals.[21] The agent's own actions in the world are not represented in its knowledge base. In other words, the agent has no memory of what it did, nor awareness of what it is doing, for that matter. Second, except for (De Eugenio, 1998), none of the above systems is concerned with natural language. The research programs underlying them concentrate on the problems associated with acting and reasoning about action, paying no attention to the issues of communicating domain knowledge and plans in natural language—an important concern for user-friendly practical robotics. This should be reflected on how actions and states are represented, for example. Even in the case of (De Eugenio, 1998), the research project outlined therein is not one of a general linguistically-competent cognitive agent (one that can understand general narrative, for instance), but rather of an acting system that can understand natural-language *instructions*. The goal of this proposal is to present a theory of an acting agent that has a record of what it did and observed in its world and that may engage in a dialogue about it. None of the above systems have such capabilities.

In addition, as will be discussed in detail in Chapter 9, there has been no general theory of how reasoning controls the sequencing of actions (see Section 1.4.2). It will also be argued in Chapter 10 that available accounts of interrupt handling based on domain knowledge are not general enough for our reasoning, acting, and interacting agent.

2.4 Conclusions

To sum up, it is clear from reviewing the literature that, although some of the issues that this dissertation purports to tackle are covered in previous work, there are two respects in which this coverage seems to be lacking. First, none of the reviewed contributions presents an integrated theory of reasoning and acting time addressing *all* the issues discussed in Chapter 1. Second, even when individual issues are addressed and successfully resolved, it is not obvious how the proposed systems could be extended in order to fit within an integrated theory as the one sought in this dissertation.

[21] This may not a be precise characterization of the system of (Georgeff and Lansky, 1987), which may acquire knowledge through its sensors while acting.

Chapter 3

The Logic of States

In this and the next chapter, I present a conceptual and formal framework for reasoning about situations. The ontology underlying the system is an epistemological one (eventually grounded in Meinongian semantics (Shapiro and Rapaport, 1987)), taking the perspective of a cognitive agent reasoning about the world as it unfolds in time. Because of the cognitive considerations underlying this work, I take the results of cognitive semantics seriously (Langacker, 1987; Jackendoff, 1991; Talmy, 2000), using *mereotopology* as a general framework motivating the formal system proposed (Eschenbach et al., 1994).[1] Explicit mereotopological concepts have been applied by a number of authors for the analysis of situations (Bach, 1986; Link, 1987; Krifka, 1989; Desclés, 1989; Pianesi and Varzi, 1996a; Krifka, 1998). However, these authors are primarily concerned with the study of *linguistic* aspect—the different means by which language structures situations. My concerns are slightly different, since I mainly need to account for what an acting cognitive agent *has in mind* as situations unfold in time. As shall be argued below, states play a central role in determining an agent's awareness of what goes on in its environment. In this chapter, I will lay out the foundations of a first-order calculus of states, $\mathcal{FOCS}$. The component of $\mathcal{FOCS}$ presented is a logic of time and states that would serve as the core of our theory of situations. In the next chapter, a logic of events will be erected on top of the foundations presented below.

[1] Basically, mereotopology (Varzi, 1996; Smith and Varzi, 2000) is topology based on mereology (the study of part-whole relations) (Leonard and Goodman, 1940) rather than set-theory.

Figure 3.1: A battery-operated robot transferring boxes from one side of the street into the store on the other side.

3.1 Aspect: An Acting Agent's Perspective

An acting agent reasons about various things. For example, it may reason about objects in the domain, the effects of different acts, and plans to achieve its goals. But if the agent is to interleave reasoning and acting, using reasoning to direct its actions, it would need to also reason about what *is* happening and what *has* happened. An example may help clarify the idea. Consider an agent whose sole job is to transfer boxes from one side, $S1$, of a street into a store on the other side, $S2$. Further, suppose that the agent is a battery-operated robot, and that a battery-recharging station is located on $S1$ (see Figure 3.1). Now consider the following situations.

1. The agent needs to cross the street. However, it cannot cross the street until it comes to believe that the walk-light is on. Thus, knowledge of what is happening in the environment (whether the walk-light is on) as it acts is necessary for the very choice of the agent's actions.

2. The agent's plan is to pick up a box, cross the street to $S2$, walk into the store, put down the box, and repeat the sequence until all the boxes have been transfered to the store. To correctly perform this sequence, the agent should, for example, put down a box when, and only when, it comes to believe that it *has* walked into the store. That is, to carry out a simple sequence of acts, the agent should be aware of what it has done at each step in the sequence.

3. The agent notices that its battery is low and, therefore, decides that it should recharge it. Now, whether the agent should immediately attend to the low-battery situation depends on what it happens to be doing. For example, if the agent is on $S1$, has just picked up a box, and is

ready to cross the street; it might be better to interrupt what it is doing, put the box down, and move to the recharging station. On the other hand, if it is in the store and about to put down a box, then it is more appropriate to continue what it is doing and then move to the recharging station. Things get more complicated if the battery goes low while the agent is crossing the street. If it is crossing from $S2$ to $S1$, then it should not interrupt such an activity, since it constitutes part of the plan to recharge the battery. However, if it is going from $S1$ towards $S2$, the agent may need to engage in some complex temporal reasoning to figure out whether there is enough time for it to reach $S2$, and then go back to $S1$ before the battery is totally dead. In any case, what seems to be a fundamental issue underlying any interrupt-handling-oriented reasoning is the agent's knowledge of what it is doing.

What the above example shows is that, for a theory of situated agents, one necessarily needs an account of, not only eternal facts about objects and actions, but also temporal creatures that occur in, and hold over, time: states, events, and processes; or, in general, *situations*. By "an account", I mean a logic, together with an appropriate ontology, of situations that would represent the agent's beliefs about different situations it may come to know of. Now, as just described, such an account can be any of those abundantly found in the literature on time and linguistic aspect (Taylor, 1977; Gabbay and Moravcsik, 1979; Allen, 1984; Galton, 1984; Parsons, 1989; Verkuyl, 1989; Smith, 1997, for instance). The task is then to merely choose one, and maybe argue for such a choice. In a way, this is what I shall be doing in this chapter. Nevertheless, there is more to the story.

What is the study of linguistic aspect concerned with? Primarily, language, in particular, the question of how it is that acoustic patterns, or symbols on paper, evoke particular beliefs about the temporal structure of situations. My concerns are slightly different in two respects. First, as described above, linguistic aspect covers much wider territory than I am willing to conquer. For example, I'm not concerned with the kinds of morpho-syntactic constructions that various languages employ to distinguish the different aspects (Smith, 1997). Nor am I willing to investigate the subtleties of particular aspects in particular languages (for example, the English perfect (Comrie, 1976; Michaelis, 1994)). My concerns are at the semantic end of the spectrum, where the logic of aspect resides, and where issues of representing and reasoning about different types of situations are at stake. Second, and more important, the task of the linguist is slightly different from mine. Given a sentence describing some situation, it is the linguist's job to provide a representation of its meaning that would allow us to make sound, intuitive conclusions about the reported situation. Now, note that the linguist is given "a sentence", itself a formal representation of the situation, annotated with the perspective and epistemic status of some speaker (or writer). For example, consider the following sentences.

(3.1) I was crossing the street.

(3.2) I crossed the street.

The two sentences may be describing a single actual situation, but certainly the descriptions are very different. The imperfective (1) describes the situation from an interior temporal point and, thus, does not convey any information about how, or whether, the situation has ended. The perfective (2) describes the situation as a unitary, temporally-bounded whole and, thus, reports that it has reached an end. It is the linguist's job to precisely account for how the various components of the above sentences interact to give rise to the aspectual properties exhibited by each. But note that the point where the linguist's role starts is not the beginning of the story; there is a speaker that has already carved a situation out of the continuum of perceptual experience, and that has selected a particular way of conceptualizing it. For the speaker, the world does not present the whole situation at once, but the speaker observes it, dynamically, as it unfolds in time.

More specifically, for a cognitive agent acting in the world, knowledge of situations does not only come from interacting with other agents. That is, it does not always come in linguistic form. In fact, most of the time, our knowledge of situations (in particular, of the immediate surroundings) is conveyed, not through language, but via perception and first-person knowledge of what we are doing. Information about situations presented to us through these non-linguistic means has a peculiar property that will be discussed in the next section; it comes in a particular, relatively rigid form that does not exhibit any of the fancy perspectival options available in linguistic reports (for example, the perfective/imperfective distinction).

The theory of situations developed here is two-pronged. On one hand, it needs to account for linguistic data, at least in a limited sense. The reason is that I believe that an important (or, at least, useful) feature of a cognitive agent is its ability to discuss what it has done and what it is doing, preferably in natural language. In this sense, the approach presented here is different from that of formal theories of actions in AI (McCarthy and Hayes, 1969; Kowalski and Sergot, 1986; Lifschitz, 1987; Levesque et al., 1997, for example), which are notoriously not concerned with language. On the other hand, the theory needs to account, within the same framework, for how a rigid representation of perceptual/proprioceptual information may give rise to various representations corresponding to the different ways language describes a situation. In developing the theory, I am closely following theories of linguistic aspect (in particular, (Galton, 1984; Parsons, 1989; Herweg, 1991b)) but occasionally diverting from them, both at the technical and the ontological levels, to account for phenomena that might not be relevant to language but are crucial for an agent reasoning and acting in time.

3.2 Of Time and Space

It is always said that there is a striking similarity between the ways we conceive of time and those in which we conceive of space, at least as revealed by language (Mourelatos, 1978; Bach, 1986; Krifka, 1989; Jackendoff, 1991; Talmy, 2000, for instance). In particular, objects (denoted by count terms) correspond to events, and matter (denoted by mass terms) corresponds to states. I will take this fairly acceptable view as a basis for my distinction between what I call "events" and what I call "states".

The very basic difference between objects and matter is topological: objects are conceived of as topologically-closed and matter as topologically-open. By alluding to topology here, I do not have in mind classical set-theoretical topology, but the more cognitively-relevant *mereotopology* (Varzi, 1996; Smith and Varzi, 2000): topology based on mereology (the study of part-whole relations) (Leonard and Goodman, 1940) rather than set-theory. What the distinction amounts to is that objects have their boundaries as parts whereas matter does not.[2] On this view, any bounded amount of matter constitutes an object. Thus, a pile of sand, a lake, and a beam of light are objects, but sand, water, and light are only matter.

Similarly, events are (temporally) closed situations, ones that have their boundaries as parts, and states are open situations (also see (Desclés, 1989; Pianesi and Varzi, 1996a)). This loosely corre-

[2] Without getting absorbed into philosophical details, note that I do not imply that boundaries are somehow determinant of the identity and individuality of objects. A suit, for example, might be considered a single object even though no self-connected (physical) boundary demarcates it. For more on this, see (Cartwright, 1975).

sponds to the linguistic distinction between bounded and unbounded sentences (Declerck, 1979). For example, the imperfective (unbounded) sentence (1) describes the street-crossing situation as a state, since the temporal boundary of the situation (i.e., its end) is not part of the description; as far as we can tell, the speaker might be still crossing the street. Thus, though pragmatically awkward, (3) is logically plausible:

(3.3) I was crossing the street, and I still am.

On the other hand, the perfective (bounded) (2) describes the situation as an event, a bounded whole. Because its temporal boundary is part of an event, event-sentences always imply that the reported situation has come to an end. Witness the awkwardness of (4):

(3.4) *I crossed the street and I still am (crossing it).

Now, as far as language is concerned, the space-time analogy (or metaphor (Lakoff and Johnson, 1999, ch. 10)) is almost perfect. However, there is a certain respect in which it seems not to hold. In particular, the analogy fails in the way we actually experience time and space. In our everyday experience, we encounter objects; we see them, touch them, and manipulate them in a, more or less, direct way. However, we rarely encounter matter per se. Matter always comes packaged, taking the form of objects that we directly interact with. Thus, we do not see “wood”, “glass”, or “paper”; we see chairs, bottles, and books. Of course, there are exceptions. For example, we directly encounter “water” when we’re totally submerged in a deep, indefinitely-extended ocean. Other examples, may be light, air, or sand on a beach. However, other than such large-scale natural extensions of matter, we always seem to get to matter only through objects that we encounter in our everyday getting-around in the world.

Our temporal experience, on the other hand, follows the exact opposite pattern (“opposite” with respect to the analogy discussed above). We never experience an event, a *whole* situation; no sooner have we reached the end of a situation, than its beginning has already moved into the past, beyond the reach of our conscious experience. Instead, the world continuously unfolds as “now” relentlessly moves, presenting us with a continuous flux of states. Consider, for example, the agent of Section 3.1. In any phase of its iterative act of transferring boxes from one side of the street into the store on the other side, the relevant aspects of the agent’s experience may be (roughly) plotted against time as in Figure 3.2. At any particular moment of experience, any “now”, the agent only has access to

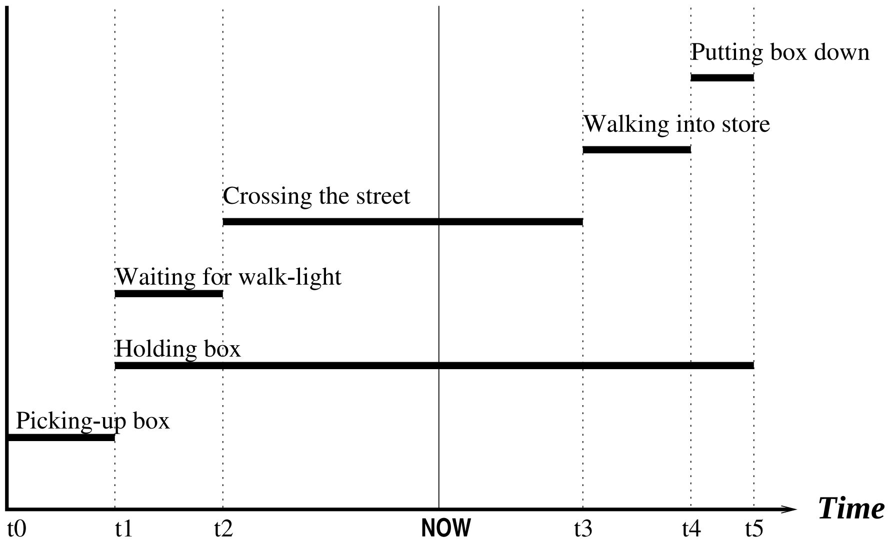

Figure 3.2: The experience of the agent of Figure 3.1 as the world unfolds.

how the world is at that moment, as reflected by the *state* of its sensors. For example, at the "now" shown in the figure, the agent believes that it is carrying a box and crossing the street; at any "now" between $t1$ and $t2$, it believes that it is holding the box and waiting for the walk-light; and between $t3$ and $t4$, it believes it is holding the box and walking into the store. Evidently, whatever is "now" the case is a state, never an event, for an event has its boundary as an essential part and, thus, can only exist in retrospect, when it has reached an end.[3]

But if experience consists of only a cascade of states starting, persisting, and ceasing, where do events come from? Events are purely conceptual beasts; we conjure them up by conceptualizing a whole out of some state's starting to hold, holding for a while, and then ceasing. Thus, whereas, in our spatial experience, we primarily encounter objects and then conceptually pass them through "the universal grinder" to get matter; in our temporal experience, we only encounter states which are then passed through "the universal packager" to yield whole events.[4]

The above discussion is not intended to present a mere amusing observation. It is intended to lay the ground for the kind of theory of situations that I shall outline below. The theory is intended to account for what an agent has in mind as situations unfold, revealing themselves piece

[3]An event can also be *imagined* in prospect (see Section 4.6.2).

[4]The terms "universal grinder" and "universal packager" are traditionally attributed to David Lewis. However, (Jackendoff, 1991, p. 24, fn. 11) reports that "David Lewis has informed [him] … that he has not used [these] term[s] in print: he in turn attributes the notion[s] to lectures or writings by Victor Yngve in the 1960s which he is now unable to trace."

by piece as "now" moves. In such a theory, states play a central role, everything else being, in some way or another, ontologically-dependent on, or -inferior to, them. In the following sections, I will develop a first-order calculus of states, $\mathcal{FOCS}$, that will serve as the formal framework within which discussions in the following chapters shall proceed.

3.3 A First-Order Calculus of States

Instead of laying out the syntax and semantics of $\mathcal{FOCS}$ all at once, the system will be introduced gradually, as the discussion flows. In this section, I will only state some general features of the logic. $\mathcal{FOCS}$ is a sorted first-order language. Using standard type-theoretical notation, the set of denotations of symbols of sort Σ is denoted $\mathcal{D}_\Sigma$.

3.3.1 Propositions

Before proceeding any further, let me point out a caveat in my use of the term "relation" below. Formally, a relation is a set of tuples. In standard model-theory for first-order logics, predicate symbols denote relations. A simple well-formed *ground* expression such as $p(a_1, a_2, \ldots, a_n)$, where p is an n-ary predicate symbol and $a_1, a_2, \ldots, a_n$ are terms of the appropriate sorts, is a true proposition of the logic if and only if $\langle [\![a_1]\!], [\![a_2]\!], \ldots, [\![a_n]\!] \rangle \in [\![p]\!]$.[5] But recall that I am not interested in truth; rather, the language being developed is to represent Cassie's beliefs. Now, among other things, Cassie may have beliefs *about* propositions; for example, she may hold a belief that Aristotle believed that the earth is flat. In other words, propositions seem to be the kind of entity that one may think and talk about (we are doing this right now). This is arguably a good enough reason to grant propositions first-class status in the ontology (we also stay within the familiar confines of first-order logic).[6] On that view, an expression such as $p(a_1, a_2, \ldots, a_n)$ is not a *sentence* in $\mathcal{FOCS}$, but rather, a term, one that denotes a proposition. What then is the status of p? Naturally, p can no longer be a predicate (at least in the traditional sense); it can only be a function symbol, denoting a function whose range is the set of propositions. Thus, strictly speaking, a symbol such as $\sqsubseteq$ (which will be used to represent temporal parthood) denotes a function, not a relation.[7] Nevertheless, for convenience, I shall

[5]For a term τ, $[\![\tau]\!]$ represents its denotation.

[6]In doing so, I am following (Shapiro, 1993). A more elaborate discussion of this issue may be found there.

[7]Though one can always view an n-ary function as an $n+1$-ary relation.

allow myself to be a little sloppy in informal discussions, and overlook the distinction. Thus, I shall engage in some "relation"-talk, saying things such as "$\sqsubseteq$ is transitive", although such terminology is mathematically incorrect given that $[\![\sqsubseteq]\!]$ is a function.

Given the above discussion, the first $\mathcal{FOCS}$ sort to be introduced is one for propositions. Technically, there is a sort, $\mathcal{P}$, where $\mathcal{D}_{\mathcal{P}}$ is a set of propositions. A subset, β, of $\mathcal{D}_{\mathcal{P}}$ represents Cassie's belief space—the set of propositions that she believes. p, p', and p_n, where $n \in \mathbb{N}$ is a natural number, are meta-variables ranging over members of $\mathcal{P}$. Now, instead of denoting functions over truth values, standard logical connectives in $\mathcal{FOCS}$ should denote functions over propositions. For the moment, I will assume this to be the case. For example, assuming standard semantics, $\neg$ is a function from $\mathcal{P}$ into $\mathcal{P}$ and $\wedge$ is a function from $\mathcal{P} \times \mathcal{P}$ into $\mathcal{P}$. However, this shall be revised in subsequent sections, where it will be argued that propositions, for all practical purposes, may be replaced with the notion of *eternal states* (see Section 3.6 below).

3.3.2 Terms

$\mathcal{FOCS}$ terms are of three possible forms.

1. Constants and variables are terms of the appropriate sort.

2. If f is a function symbol with domain A and range B, and if x is a tuple in A, then $f(x)$ is a functional term in B. Specific function symbols will be presented below as the discussion flows.

3. If x is a variable of sort Σ, τ is a term in Σ, and p is a term in $\mathcal{P}$ with one or more occurrences of τ and no occurrences of x, then $\forall x[p\{x/\tau\}]$ and $\exists x[p\{x/\tau\}]$ are terms in $\mathcal{P}$.[8] $[\![\forall x[p]]\!]$ is the proposition that, for all $\tau' \in \Sigma$, it is the case that $[\![p\{\tau'/x\}]\!]$. Similarly, $[\![\exists x[p]]\!]$ is the proposition that, for some $\tau' \in \Sigma$, it is the case that $[\![p\{\tau'/x\}]\!]$.

3.3.3 Rules of Inference

I assume the existence of standard inference rules (for example, introduction and elimination rules for connectives and quantifiers). However, because of assumptions underlying our theory, a couple

[8]For any term τ and any two terms τ_1 and τ_2 of the same sort, $\tau\{\tau_1/\tau_2\}$ is the result of replacing all instances of τ_2 in τ by τ_1.

of points need to be made explicit. In standard first-order predicate calculus, the interpretation of a rule of inference is that if the premises are true, then so is the conclusion. In our theory, truth is replaced with belief by Cassie. That is, if Cassie believes the premises (i.e., if they are in β), then she *may* also believe the conclusion. As (Shapiro, 1993) pointed out, it is important to note that rules of inference should not be taken as saying that if the premises are in β, then so is the conclusion. In particular, I do not assume β to be closed under deduction (Johnson and Shapiro, 2000b).

3.3.4 Equality

$\mathcal{FOCS}$ does not include an equality operator. However, there is an *equivalence* function, Equiv, that serves a broader, but needed, purpose.

- $\mathsf{Equiv} : \Sigma \times \Sigma \longrightarrow \mathcal{P}$, where Σ is any $\mathcal{FOCS}$ sort. $[\![\mathsf{Equiv}(\tau, \tau')]\!]$ is the proposition that $[\![\tau]\!]$ and $[\![\tau']\!]$ are co-referential, i.e., the two Meinongian objecta "pick out the same extension in some world" (Maida and Shapiro, 1982, p. 298).

Without stating them, I assume the existence of axioms branding Equiv as an equivalence "relation".

3.4 Time

Any account of situations will have to be based on some theory of time. The perspective that this work takes limits the options available for such a theory. Two points are, particularly, worth mentioning.

1. The mind of a cognitive agent does not contain an infinite (countable or uncountable) number of temporal individuals. At any point, there is only a finite number of times represented in Cassie's mind. Here, there is an underlying distinction between the contents of Cassie's mind and the formal language used to represent those contents. The language may indeed provide an infinite supply of terms denoting temporal individuals. However, this does not mean that all of these terms are represented in Cassie's mind, and are merely *discovered* as she conceives of a particular time. Rather, the mental event of conceiving of a time corresponds to the introduction of one of these terms into Cassie's consciousness. Hence, contrary to other work (van Benthem, 1983; Habel, 1994, for instance), I shall not assume that time is dense (or

continuous); the ontology sought is an epistemological one, and assuming density means that the agent's mind contains an infinite number of terms. Note that Cassie may have a *belief* that time is dense. This, however, does not necessitate that the model itself be dense. The main reason I am pointing this out is to dismiss models of time based on the rationals (or the reals) that might be anticipated by some readers.

2. Cassie may believe that time is linearly ordered. That is, $\mathcal{FOCS}$ may include an axiom (representing a proposition that Cassie believes) to the effect that any two times are either contemporaneous or one of them precedes the other. Nevertheless, the structure of time in Cassie's mind need not be linear. That is, meta-theoretically, not all pairs of time-denoting terms are related by the precedence relation. For instance, Cassie may have beliefs about two times, t_1 and t_2, without having any (need for a) belief about their order. This point is discussed in detail by (Kamp, 1979).

What kind of temporal individuals are to be admitted in the ontology? Points, intervals, or both? To start with, it could be argued that, from a strictly objectivist standpoint, and according to some theories of physics, the notion of a time point is nothing other than a short interval of time. As (van Benthem, 1983, p. 11) noted, one has to remember that "short" is a context-dependent modifier whose interpretation is highly-dependent on the notion of comparison (Klein, 1980). That is, one cannot once and for all identify those intervals that may be considered points; it all depends on the level of granularity at which one is *viewing* things, a purely cognitive notion. For instance, consider the following sentences from (Talmy, 2000).

(3.5) She climbed up the fire-ladder in 5 minutes.

(3.6) Moving along on the training course, she climbed the fire-ladder at exactly midday.

Here, (5) invokes a conceptualization of the ladder-climbing situation as a temporally-extended event. The same situation is conceptualized as a point-like event through the use of the "*at*" + *point-of-time* construction and the contextual specification in (6).

From the above discussion, one might conclude that we need to recognize both time points and time intervals in our representations. A more economical decision could be made, however; only intervals are to be represented. Strictly speaking, nothing much hangs on the fact that we *think* of time terms as denoting intervals; we could view them as mere *times* without any commitment to their being intervals or points. It is convenient though to consider them intervals for the simplicity of giving the semantics of relations among them. I maintain the view that the times of punctual events (typically viewed as points) are simply very short intervals of time. However, I shall not resort to representation of the metrical concept of length for this matter. Instead, I shall use the mereological notion of a subinterval to identify short intervals (or, if you will, points) of time. Informally, a time interval is treated by Cassie as a time point if she has no beliefs about its subintervals.

Now, let us move to a more formal characterization of the temporal component of $\mathcal{FOCS}$.

Sort. There is a sort, $\mathcal{T}$, where $\mathcal{D}_{\mathcal{T}}$ is a set of time intervals. I use t, t', and t_n ($n \in \mathbb{N}$) as meta-variables ranging over members of $\mathcal{T}$.

Constants. Tc, Tc', and Tc_n ($n \in \mathbb{N}$) are constant symbols in $\mathcal{T}$.

Variables. Tv, Tv', and Tv_n ($n \in \mathbb{N}$) are variables in $\mathcal{T}$. All variables are universally-quantified with widest scope unless otherwise noted.

Function Symbols. There is a set of function symbols denoting functions from $\mathcal{T} \times \mathcal{T}$ into $\mathcal{P}$. These correspond to temporal relations (see the discussion on "relations" in Section 3.3.1). The particular suite of relations employed here does not reflect any particular logical commitments; I only introduce those relations that shall be used below; more relations could be introduced as needed.

- $\prec : \mathcal{T} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![t \prec t']\!]$ is the proposition that $[\![t]\!]$ precedes $[\![t']\!]$.
- $\sqsubseteq : \mathcal{T} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![t \sqsubseteq t']\!]$ is the proposition that $[\![t]\!]$ is a subinterval of $[\![t']\!]$.
- $\sqsubset : \mathcal{T} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![t \sqsubset t']\!]$ is the proposition that $[\![t]\!]$ is a proper subinterval of $[\![t']\!]$.
- $\supset\subset : \mathcal{T} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![t \supset\subset t']\!]$ is the proposition that $[\![t]\!]$ *meets* (in the terminology of (Allen, 1983)), or *abuts* (in the linguistic terminology (De Swart, 1998, for instance)), $[\![t']\!]$.

Axioms. Axioms represent beliefs that Cassie has about properties of time and temporal relations. The following set of axioms characterize $\prec$ as a strict partial order (note the "relation"-talk).

- **AT1.** $Tv \prec Tv' \Rightarrow \neg[Tv' \prec Tv]$.

- **AT2.** $[Tv_1 \prec Tv_2 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_1 \prec Tv_3$.

- **AT3.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow \neg[Tv \prec Tv']$.

Similarly, the following axioms characterize $\sqsubseteq$ as a partial order.

- **AT4.** $[Tv \sqsubseteq Tv' \wedge Tv' \sqsubseteq Tv] \Rightarrow \mathsf{Equiv}(Tv, Tv')$.

- **AT5.** $[Tv_1 \sqsubseteq Tv_2 \wedge Tv_2 \sqsubseteq Tv_3] \Rightarrow Tv_1 \sqsubseteq Tv_3$.

- **AT6.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow Tv \sqsubseteq Tv'$.

The strict version, $\sqsubset$, of $\sqsubseteq$ is defined by the following axiom.

- **AT7** $Tv_1 \sqsubset Tv_2 \Leftrightarrow [Tv_1 \sqsubseteq Tv_2 \wedge \neg\mathsf{Equiv}(Tv_1, Tv_2)]$

Although I shall not explicitly state theorems to the effect, we can prove that $\sqsubset$ is a strict partial-order. I also assume the existence of axioms defining Equiv as an equivalence relation We may now prove the following simple but important results (see Appendix B for the proofs). **TT1** states that temporal part-hood and temporal precedence are contradictory. **TT2** states that if parts of an interval precede each other, then they are *proper* parts of that interval.

- **TT1.** $Tv_1 \prec Tv_2 \Rightarrow [\neg[Tv_2 \sqsubseteq Tv_1] \wedge \neg[Tv_1 \sqsubset Tv_2]]$

- **TT2.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow [Tv_2 \sqsubset Tv_1 \wedge Tv_3 \sqsubset Tv_1]$

The relation $\supset\subset$ is characterized by the following axiom.

- **AT8.** $[Tv_1 \supset\subset Tv_2] \Leftrightarrow [Tv_1 \prec Tv_2 \;\; \wedge \;\; \neg\exists Tv_3[Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2]]$.

Using **AT8**, **AT1**, and **AT3**, we can prove that $\supset\subset$ is anti-symmetric, intransitive, and irreflexive.[9]

- **TT3.** $Tv \supset\subset Tv' \Rightarrow \neg[Tv' \supset\subset Tv]$.

- **TT4.** $[Tv_1 \supset\subset Tv_2 \wedge Tv_2 \supset\subset Tv_3] \Rightarrow \neg[Tv_1 \supset\subset Tv_3]$.

[9]Formal proofs for a selected subset of the theorems presented in this chapter are to be found in Appendix B.

Temporal Relation	Allen-style Disjunction
$\supset\subset$	(m)
$\prec$	(b, m)
$\sqsubseteq$	(s, d, f, =)
$\sqsubset$	(s, d, f)

Table 3.1: Temporal relations and the corresponding disjunctions of Allen's (Allen, 1983) mutually exclusive relations.

- **TT5.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow \neg[Tv \supset\subset Tv']$.

The interaction between $\prec$ and $\sqsubseteq$ is illustrated by the following axioms. **AT9** and **AT10** state that inclusion preserves order,[10] and **AT11** asserts that all intervals in $\mathcal{D}_{\mathcal{T}}$ are convex.

- **AT9.** $[Tv_1 \prec Tv_2 \wedge Tv_3 \sqsubseteq Tv_1] \Rightarrow Tv_3 \prec Tv_2$.
- **AT10.** $[Tv_1 \prec Tv_2 \wedge Tv_3 \sqsubseteq Tv_2] \Rightarrow Tv_1 \prec Tv_3$.
- **AT11.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_4 \wedge Tv_4 \prec Tv_3] \Rightarrow Tv_4 \sqsubseteq Tv_1$.

The reader should note that, except for $\supset\subset$, none of the temporal relations presented correspond to any of the thirteen relations of (Allen, 1983).[11] They do, however, correspond to disjunctions thereof. Table 3.1 shows our relations and the corresponding Allen-style disjunctions.[12] (Freksa, 1992) argues that it is more cognitively plausible (and even more efficient) to represent some relations as primitives rather than Allen-style disjunctions, so long as they satisfy a certain property—forming a *conceptual neighborhood*. Two relations between pairs of intervals are conceptual neighbors if they may be *directly* (i.e., without *passing* through other relations) transformed into one another by applying a continuous (topological) deformation on the intervals (for example, by moving, stretching, or contracting the intervals). A conceptual neighborhood is a set of relations whose "elements are path-connected through 'conceptual neighbor' relations" (Freksa, 1992, p. 204).

(Freksa, 1992) argues that, if one is interested in representing knowledge about temporal relations, reasoning with conceptual neighborhoods is more plausible than with disjunctions. For, in this case, less knowledge corresponds to simpler representations, and more knowledge corresponds

[10] These are the "left monotonicity" and "right monotonicity" axioms of (van Benthem, 1983, p. 67), respectively.

[11] The thirteen relations are: b(efore), m(eets), o(verlaps), s(tarts), d(uring), f(inishes), =, and their inverses (= is it own inverse). The inverse of a relation "r" is denoted by "ri".

[12] Note that $\prec$ is simply the transitive closure of $\supset\subset$.

to more complex ones (in particular, conjunctions). Compare this to Allen's system, where less knowledge is represented by complex disjunctions. Since I *am* representing the knowledge of an agent, I opt for more loaded primitives corresponding to conceptual neighborhoods. I believe that incomplete (and inaccurate) knowledge is the norm, and thus a primitive relation should have a broader range of coverage, one that may be restricted as more knowledge is acquired.

My precedence relation, $\prec$, represents Cassie's knowledge of one event being entirely before another. Granularity aside, whether something happens in between or not is often not available and usually not important to know. Thus, $\prec$ does not distinguish between Allen's "m" and "b" (witness **AT8**). Similarly, $\sqsubseteq$ represents simple temporal parthood. We often do not care to distinguish among whether an interval is an initial, internal, or final subinterval of another.

3.5 States, Change, and Temporal Stability

3.5.1 States

The most basic type of situations in the ontology developed here is what I call *states*. The term "state" has been used by different authors, sometimes referring to different notions in the temporal domain. To make clear how I use the term, let us look at three of the common usages of "state".

(McDermott, 1982, p. 105), informally, defines a state to be "an instantaneous snapshot of the universe". This is similar to the "situations" of the situation calculus (McCarthy and Hayes, 1969), although McDermott does not spell it out. A state, in this sense, is defined, or is identifiable, by the sequence of events leading to it. In a linear model (not McDermott's branching structure for states), states are isomorphic to the times at which they are located.

A second use of "state" comes from the situation calculus.[13] (Thielscher, 1999) explains:

> The union of the relevant fluents that hold in a situation is called the *state* (of the world) in that situation. Recall that a situation is characterized by the sequence of actions that led to it. While the world possibly exhibits the very same state in different situations, the world is in a unique state in each situation.

According to this definition, the state (of the world) at some time is the set of facts, "relevant" facts,

[13]The term is also used to refer to the same notion under a different formal guise by (Shanahan, 1995).

that are true then. Note that, though apparently different, this notion of a "state" is actually close to that of McDermott. The only difference lies in the restriction of relevance. Had a state been defined as the set of *all* facts that are true in a situation, there would have been only one situation corresponding to each state. For, in each new situation, it is a fact that some new event has just taken place.

A third use of the term "state", the one that I adopt, comes from linguistics and the philosophy of language (Vendler, 1957; Mourelatos, 1978; Galton, 1984, to mention a few). (Galton, 1984, pp. 27–28) explains that "a state is homogeneous, dissective, and negatable, and obtains or fail to obtain from moment to moment." Although there is no general agreement as to the exact domain of coverage of this use of the term, almost everybody would agree that the following sentences are examples of states.

(3.7) Mary knows the answer.

(3.8) John is in New York.

(3.9) The litmus paper is red.

My use of the term *state* indeed covers the above sentences. However, some of its uses may seem new, or even strange, to some. To prepare the scene for these uses, we need to provide a more basic definition of what a state is.[14]

Definition 3.1 *A* **simple state** *is the existence (or absence) of some relation, R^n, among the entities of a tuple $\langle x_1, x_2, \ldots, x_n \rangle$.*

Each simple state is thus identifiable by:

1. the relation R^n,
2. the tuple $\langle x_1, x_2, \ldots, x_n \rangle$, and
3. whether the relation exists among the entities in the tuple.[15]

---

[14] Note that this and Definition 3.2 are not intended to be formal definitions in the strict sense; they are merely there to precisely state certain intuitions.

[15] This should remind the readers with similar ideas in situation semantics (Barwise and Perry, 1983). See below for a closer look on the similarity.

This, I believe, is as basic and simple as one can get. For nothing is more basic than the notion of an unrestricted relation existing among entities (note that this includes properties an entity may exhibit, when $n = 1$). In particular, we do not perceive objects simpliciter; we perceive them exhibiting certain properties or standing in relations with other objects.

Simple states, however, are not all there is; they may be combined, in the sense of (Bach, 1986; Link, 1998), to form *complex states*.

Definition 3.2 *A* **complex state** *is the mereological fusion of other complex and simple states.*

Thus, whenever each of a collection of states holds over some time, t, they give rise to a whole, a complex state, that holds over t. To formally establish the relation between states and time, we need to introduce the state-component of $\mathcal{FOCS}$.

Sort. There is a sort, $\mathcal{S}$, such that $\mathcal{D}_\mathcal{S}$ is a set of states. I use s, s', s_n ($n \in \mathbb{N}$) as meta-variables ranging over members of $\mathcal{S}$.

Constants. Sc, Sc', and Sc_n ($n \in \mathbb{N}$) are constant symbols in $\mathcal{S}$.

Variables. Sv, Sv', and Sv_n ($n \in \mathbb{N}$) are variables in $\mathcal{S}$.

Function Symbols. Only two functions need to be introduced at this point. More will be introduced below as the system evolves.

- $\mathsf{Holds}: \mathcal{S} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![\mathsf{Holds}(s,t)]\!]$ is the proposition that the state $[\![s]\!]$ holds throughout the interval $[\![t]\!]$.
- $\mathsf{MHolds}: \mathcal{S} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![\mathsf{MHolds}(s,t)]\!]$ is the proposition that the state $[\![s]\!]$ *maximally* holds throughout the interval $[\![t]\!]$.

States are ontologically-fundamental entities, ones to whose presence in the ontology I am committed, regardless of the particular domain $\mathcal{FOCS}$ is used to represent (much like propositions in the propositional calculus, for example). Therefore, Holds is a *core* function in $\mathcal{FOCS}$; it is more like a logical connective, a syncategorematic symbol, rather than one that depends on the domain represented by $\mathcal{FOCS}$. For any given domain, however, there will be a domain-dependent collection of function symbols in $\mathcal{S}$. For example, in a domain where we are interested in spatial relations, an expression such as $\mathsf{On}(x,y)$, might be used to denote the state of object $[\![x]\!]$ being on object $[\![y]\!]$. Here On is a function symbol in $\mathcal{S}$—one whose range is $\mathcal{S}$. The state $\mathsf{On}(x,y)$ is a simple state, where the entities $[\![x]\!]$ and $[\![y]\!]$ are related by the relation $[\![\mathsf{On}]\!]$ (note the relation-talk).

Axioms. All states are characterized by the two following axioms. In what follows, they will guide our intuitions to what might be labeled "a state".

- **AS1.** $[\forall Tv' \; [Tv' \sqsubset Tv \Rightarrow \mathsf{Holds}(Sv, Tv')]] \Rightarrow \mathsf{Holds}(Sv, Tv)$.

- **AS2.** $\mathsf{Holds}(Sv, Tv) \Rightarrow \; \forall Tv'[Tv' \sqsubset Tv \Rightarrow \mathsf{Holds}(Sv, Tv')]$.

Using Krifka's terms (Krifka, 1989), **AS1** says that states are *cumulative*, whereas **AS2** states that they are *divisive*. Cumulativity means that, if a state holds over all proper subintervals of some interval, t, then it also holds over t. Divisivity means that, if a state holds over an interval, then it holds over all of its subintervals. Together, the two axioms state that an interval over which a state holds is *homogeneous*; all of its parts look alike in some respect.

The following axiom represents a meaning postulate for MHolds.

- **AS3.** $\mathsf{MHolds}(Sv, Tv) \Rightarrow [\mathsf{Holds}(Sv, Tv) \wedge \neg\exists Tv'[\mathsf{Holds}(Sv, Tv') \wedge Tv \sqsubset Tv']]$

The reader should note that my notion of a state is similar, but not identical, to the *situation types* of situation semantics (Barwise and Perry, 1983). The term *infon* has been introduced later to refer to the same concept (Devlin, 1991). Essentially, a situation type, or an infon, is a set of pairs $\langle\langle r, x_1, \ldots, x_n\rangle, i\rangle$, where r is a relation, $x_1, \ldots, x_1$ are objects, and i is 0 or 1. Whether the second element of a pair is 0 or 1 depends on whether the relation holds among the objects in the first element. Situation types, in general, correspond to my complex states; those that are singleton sets correspond to simple states. For example,

$$\{\langle\langle \mathsf{Color}, \text{litmus-paper-1, red}\rangle, 1\rangle, \langle\langle \mathsf{Color}, \text{litmus-paper-2, red}\rangle, 0\rangle\}$$

is a situation type in which $[\![\text{litmus-paper-1}]\!]$ is red and $[\![\text{litmus-paper-2}]\!]$ is not. Although they do not say it explicitly, (Barwise and Perry, 1983, pp. 55–56) seem to imply that situation types are homogeneous (as per **AS1** and **AS2**).

3.5.2 Ways of State Formation

Where do states come from? That is, how does Cassie become aware of a state? There is more than one way. Here, I list five of these. The list is not meant to be exhaustive, but it suffices for the purpose of the present investigation.

Perception

This is one of the most basic means by which an agent may come to conceive of a state. By "perception", I mean any process by which raw sensory data is translated into beliefs about the external environment in the "mind" of a sensing agent. In the GLAIR architecture, perception is represented by beliefs added to the KL by PML processes. In particular, since I assume that Cassie believes whatever she perceives, such beliefs are added to β, Cassie's belief space. The following sentences represent examples of perception-invoked states, where reference to sensory data is relatively explicit.

(3.10) The air smells of jasmine.[16]

(3.11) The litmus paper is red.

Presenting well-chosen examples like (3.10) and (3.11) should not lure the reader into believing that, given a state, it is always possible to determine whether it could be invoked by perception. For example, assuming that (10) and (11) may be perception-generated presupposes a perceptual system capable of categorizing the color red and the odor of jasmine (this, again, assuming "red" and "jasmine" mean for the agent what they mean for us). Whether this is indeed the case depends on the type of agent we are considering. A typical robot, for example, does not have any capabilities for odor-detection, and, depending on its design, it may or may not be sensitive to colors. Note that this does not rule out the agent's conceiving of, and reasoning about, states involving notions of color and odor. Such states cannot be invoked by perception though.

Inference

Consider the following sentence.

[16]The sentence is due to (Mourelatos, 1978).

(3.12) The litmus paper is not red.

The sentence represents a simple state of the litmus paper *lacking* some property. Could such a state be perceived? Evidently not. One can only perceive the existence of relations among entities, not their absence. However, one may *infer* that a relation is absent. Thus, by perceiving a blue litmus paper, one infers that it is not red. Of course, *negative* states such as that in (3.12) are not the only examples of inference-invoked states. Inference underlies conceiving of any state that requires, in addition to perceiving the immediate (spatial and temporal) surroundings, general knowledge of the domain and its history.

Envelope Formation

Suppose we set a video camera shooting the scene of a table with a pen on it. If we watch the resulting motion picture one frame at a time, we should notice that, in all the frames, the pen appears to be on the table. That is, the state referred to by the sentence "*the pen is on the table*" is perceived to hold in every frame. Now suppose we do the same thing but this time shooting John while running. If we watch the motion picture one frame at a time, we shall notice that there is no *one* frame from which we may come to believe that the state represented by "*John is running*" holds. The state cannot be invoked by any single frame, but by a minimum number of them: a number that is large enough for us to notice that John moves his body parts in a certain way that is consistently repeated.

The video-camera analogy may be a reflection of some cognitive operation that takes place and directs the conceptualization of a certain kind of state: an operation that creates an envelope around a uniform pattern of changing states, with the resulting gestalt being itself conceived of as a state. Thus, even though John is not perceived to be running in any single frame, we may still believe that he is running *throughout* the interval captured by the motion picture. States that can be invoked by envelope formation but not by perception are exactly the class of situations referred to as *processes* (Mourelatos, 1978, for example) or "states of change" (Galton, 1984). Other examples of processes include the following.

(3.13) The litmus paper is turning red.

(3.14) John is writing his dissertation these days.

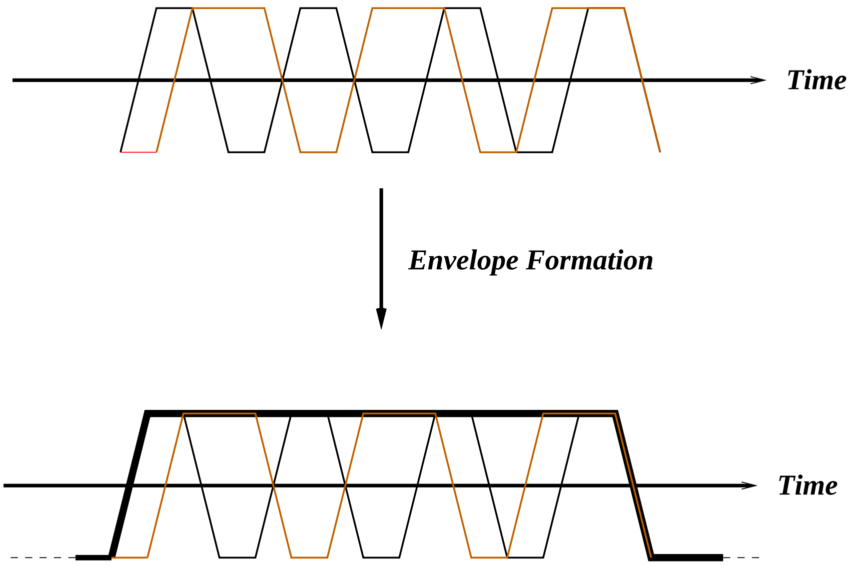

Figure 3.3: The process of envelope formation.

Note that, for envelope formation to work, there needs to be some kind of association between the length of time over which a process takes place and the rate of change of the states that it envelops. Uniform patterns of rapid state change require a relatively short period of time to form a process. On the other hand, patterns with sluggish changes require a long period of time. For instance, note the difference between *John is running* and (3.14).

First-Person Privileges

Among the states that Cassie may conceive of are states of her own self. English sentences representing such states are characterized by having as subject the first-person pronoun "I".

(3.15) I like John.

(3.16) I am standing still.

(3.17) I am running.

(3.18) I am running to the store.

Cassie's belief that she likes John is not the result of any particular perception event, nor of her being told so, but is a personal belief that only Cassie can assess. The same goes for feeling drowsy or being in pain. But examples like love and pain fall outside the scope of my concern, for they introduce issues that are much more ambitious than my attempt to account for rational action and sound, commonsense reasoning.[17] More down-to-earth examples such as (3.16)–(3.18) will be the focus of my discussion.

What exactly is involved in Cassie's belief that she is standing still? Assuming that she is in control of her own body (i.e., that she cannot be involuntarily moved), two things are involved: (i) the continuous proprioceptual feedback from her locomotive system (wheels, for instance) and (ii) her indisputable belief that she did not (intend and succeed to) start moving since the last time she (intended and successfully) stopped. Here, the information about whether she is standing still is all confined to Cassie's mind and body; there is no need to sense the external environment. The case of running (3.17) is quite similar but a little more involved. For running, feedback from the locomotive system is not sufficient; even though the wheels may be moving, Cassie should not believe that she is running if she has bumped into some obstacle. The difference here is that more proprioceptual information is combined to yield the feel for the state of running.

Despite their differences, the states represented by (3.16) and (3.17) are similar; they primarily involve proprioception, albeit with different degrees of information processing complexity at the PML. The example in (3.18) is more radically different from these two. As far as mere proprioception is concerned, there is no difference between (3.17) and (3.18); Cassie's body can only supply information about the "running" part of (3.18); it does not account for the belief that the running is to the store. This belief comes from Cassie's intentions, from a higher level of understanding what it is that she is doing. However, note that it is not just this *understanding* that makes the difference between (3.17) and (3.18); the two processes are, I believe, genuinely distinct. To run to the store, it is not sufficient for Cassie to start running and then believe that she is running to the store; she has to direct herself toward the store, maintain that direction while she is running, and make sure that she stops when, and only when, she is at the store (for more on this, see (Ismail and Shapiro, 1999)). That is, high-level intentions modulate the process of running giving rise to a substantially-different process. The importance of this point will become clear as I discuss the so-called "imperfective paradox" (Dowty, 1977).

[17] Although a low battery, with the battery conceived as part of Cassie's body, is analogous to pain.

It should now be clear that conceiving of a state due to first-person privileges is quite different from perceiving a state. However, a crucial difference has yet to be mentioned. For a first-person state, the agent *always* knows whether it holds or not. At any time, Cassie knows whether she likes John or not, whether she is standing still or running, and whether she is running to the store or to the park—no epistemological gaps, no "I don't know". On the other hand, Cassie may perceive a state holding but once her sensors are diverted to something else, it is possible for Cassie to have no beliefs whatsoever about whether the state continues to hold. For example, suppose Cassie is looking toward the walk-light, which happens to be green. Among others, there are two states that Cassie would believe to be holding.

(3.19) I am looking toward the walk-light.

(3.20) The walk-light is green.

The first state is invoked due to first-person privileges while the second is perceived. Now, if Cassie turns always from the walk-light she would believe that (3.20) no longer holds, but whether the walk-light is green or red is something that she cannot tell (at least after some time has elapsed).

Communication

Finally, Cassie may conceive of a state if she is told something about it. This, for example, is one of the main methods by which Cassie becomes aware of past states. By "telling", I mean another agent telling Cassie, with a human operator as the primary agent in mind. To be more precise, one should not single out communication as a separate way of state formation. After all, communication is a subtype of perception with inference. That is, by perceiving the communication event (which may actually involve speech perception), the gullible Cassie believes the contents of whatever she is being told. For the purpose of this study, however, communication shall be treated as an inspiration-like activity that interacts directly with the KL without passing through the PML.

3.5.3 Stability

States differ according to their degree of *temporal stability*. The notion of temporal stability goes back to (Givón, 1979), in an attempt to explain the lexical categories *noun, verb,* and *adjective*.

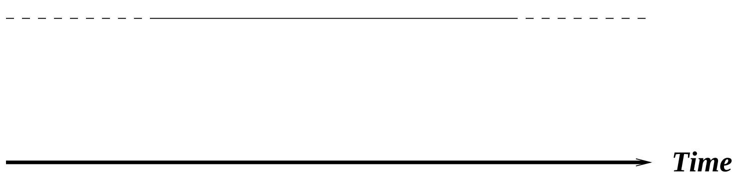

Figure 3.4: An eternally holding state.

> The *most time-stable* percepts, the ones that change slowly over time, the ones that are likely to be identical to themselves (in terms of properties), are lexicalized as nouns. The *least time-stable* percepts, events, and actions, which involve *rapid change* in the universe, are lexicalized as verbs, which by and large characterize *changes* from one steady state to another. (Givón, 1979, p. 321)

He later adds that adjectives lie somewhere between nouns and verbs on the scale of temporal stability. Such a scale is essentially a continuous rather than a discrete one. What is interesting here is that one may use that same notion (modulo some changes) to distinguish different types of states. Sentences representing those states would assume different positions on the scale of temporal stability. Three main types of states are readily available: *eternal, permanent*, and *temporary*. Formally, three sorts, ETERNAL, PERM, and TEMP, are subsorts of $\mathcal{S}$. Constants, variables, and meta-variables of these sorts shall be super-scripted by e, p, and t, respectively.

Eternal States

Eternal states do not start or cease; they either *always* hold or *never* hold. Sentences representing such states are the most temporally-stable; they (roughly) correspond to Quine's *eternal sentences* (Quine, 1960). Figure 3.4 depicts an eternal state believed by the agent to be holding. The following sentences are examples of eternal states.

(3.21) Whales are fish.

(3.22) God exists.

(3.23) The date of John's graduation is June 1st, 2001.

The first thing to notice is that a state's being eternal has nothing to do with its actually holding. Even though whales are not fish, (3.21) represents an eternal state. If someone believes that whales

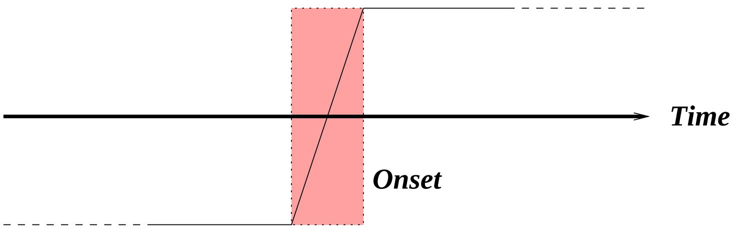

Figure 3.5: A permanent state.

are fish and later comes to believe that they are actually mammals, only one thing would have changed, namely the believer's state of mind. Whales did not become mammals; rather, the person would assume that they were and shall always be mammals. When it comes to beliefs regarding eternal states, whatever one believes is *now* the case is what they believe is always the case. This is captured by the following axiom.

- **AES1.** $\exists Tv[\mathsf{Holds}(S^e v, Tv)] \Rightarrow \forall Tv[\mathsf{Holds}(S^e v, Tv)]$.

That is, if Cassie believes that an eternal state holds over some time, then she may also believe that it holds over *all* times.

Permanent States

Unlike eternal states, permanent states may *start* to hold. Once a permanent state starts to hold, it never ceases (see Figure 3.5).[18] The *onset* referred to in the figure is the *event* of the permanent state starting to hold. Axiom **APS1** characterizes the temporal pattern governing permanent states; if a permanent state holds over some time, then it holds thenceforth.

- **APS1.** $\mathsf{Holds}(S^p v, Tv) \Rightarrow \forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(S^p v, Tv')]$

The following sentences represent permanent states.

(3.24) Fermat is dead.

(3.25) John has/had/will have turned 21.

(3.26) The Vietnam war is over.

[18]These are what (Galton, 1984) refers to as *irrevocable* states.

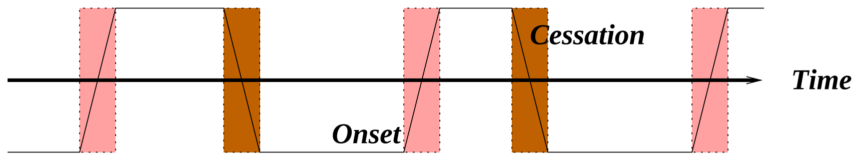

Figure 3.6: A temporary state.

There are two senses of "permanent" that should be distinguished. Let us refer to these as *naive permanence* and *logical permanence*. For example, the state reported by sentence (3.24) is permanent only in the naive sense. In reality, death might not be permanent (a common theme in many religions). Even more relevant than *reality*, many people do not believe that death is permanent. In what sense then does (3.24) refer to a permanent state? In the naive sense, in the sense that in our everyday planning and reasoning, and for all practical purposes (religion aside), we think of death as a permanent state (thus, nobody had seriously entertained the idea of waiting for Fermat to come back and show us the proof of his last theorem).

Sentence (3.25) reports what might be a logically-permanent state. Once John turns 21, the state of his *having turned 21* would always hold. Some have assumed that the perfect aspect/tense always signals a permanent state (Galton, 1984; ter Meulen, 1997). I shall not adhere to such a strong claim, though; the perfect has an evidently broader range of coverage (Comrie, 1976; Vlach, 1993; Michaelis, 1994). However, I will assume that the state of an event having occurred is permanent, corresponding to the tense-logical axiom that what is past will always be past (Prior, 1967). Such states are not always expressed using the perfect (as exemplified by (3.26)). Note that these states are *logically* permanent because of the (reasonable) assumption that time flows in one direction, that what is done cannot be undone.

Temporary States

As shown in Figure 3.6, a temporary state may repetitively start to hold and cease to hold. Unlike permanent states, temporary states do not just have an onset; they have both an onset and a cessation. These are events that mark the state starting to hold and ceasing to hold, respectively. For every time the state holds, there shall be a unique onset-cessation pair. Sentences (3.7)–(3.9) of Section 3.5.1 are typical examples of temporary states. In addition to such uncontroversial examples, I follow

(Galton, 1984) in admitting what are usually known as "processes" into the category of temporary states. As Galton puts it, these are "states of change".[19]

(3.27) I am crossing the street.

(3.28) The litmus paper is turning red.

Temporary states, being the least temporally-stable, do not require any stability-constraining axioms; they only need to satisfy **AS1** and **AS2**, which are characteristic of all states. Note, however, that both **AS1** and **AS2** are entailed by **AES1**. That is, eternal states are fully characterized by one axiom. Thus, using only **AES1**, we can prove, rather than posit, the following.

- **TES1.** $[\forall Tv' \; [Tv' \sqsubset Tv \Rightarrow \mathsf{Holds}(S^e v, Tv')]] \Rightarrow \mathsf{Holds}(S^e v, Tv)$.

- **TES2.** $\mathsf{Holds}(S^e v, Tv) \Rightarrow \; \forall Tv'[Tv' \sqsubset Tv \Rightarrow \mathsf{Holds}(S^e v, Tv')]$.

On the other hand, **APS1** entails neither **AS1** nor **AS2**; we need the three axioms to fully specify properties of permanent states. There is, therefore, some sense in which the nature of permanent states is more complex than that of eternal or temporary states. A possible explanation for this observation will be proposed below.

3.5.4 Complements and Change

As pointed out in Section 3.5.1, domain-dependent functions are used to form simple states. States formed thus, however, are only of a restricted type—the existence of some relation among a tuple of objects. Complex states and states expressing the absence of relations are to be formed using domain-independent syncategorematic operators. In this and the next section, I will introduce state operators that resemble standard truth-functional operators over propositions. In doing so, I am following (Allen, 1984). However, as shall be demonstrated in Chapter 6, these operators are needed for more than just increasing the expressivity of the language.

I will start with the traditionally simplest operator, yet the trickiest in our case: negation. I will use $\dot{\neg}$ to represent the state operator that corresponds to the propositional $\neg$.[20] The state $\dot{\neg}\, s$, called

[19] Adhering to this position does not mean that it is uncontroversial (Taylor, 1977; Gabbay and Moravcsik, 1979; Allen, 1984; Smith, 1999).

[20] I will drop the dot later, but this needs justification.

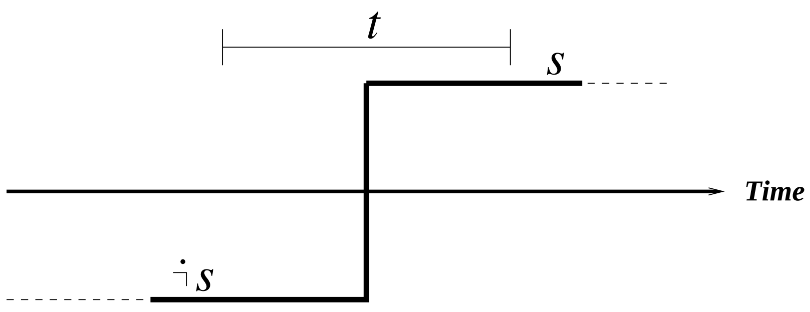

Figure 3.7: Neither s nor its complement hold over t.

the *complement* of s, is intuitively that *unique* state that holds over any interval *within* which s does not hold. It should be noted that it is not the case that, for any interval, either s or $\dot{\neg}\, s$ holds. Figure 3.7 shows such a situation, where neither state holds over t (the example was pointed out by other authors, for example (Herweg, 1991b; Koomen, 1991)). The following axiom comes from (Allen, 1984), stating a necessary and sufficient condition for $\dot{\neg}\, s$ to hold.

- **AS4.** $\mathsf{Holds}(\dot{\neg}\, Sv, Tv) \Leftrightarrow \forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \neg\mathsf{Holds}(Sv, Tv')]$

Given **AS4**, we can prove the following obvious, but important, result.

- **TS1.** $\mathsf{Holds}(\dot{\neg}\, Sv, Tv) \Rightarrow \neg\mathsf{Holds}(Sv, Tv)$

Using the above axiom and theorem, together with **AS2** (divisitivity), we can show that $\dot{\neg}\dot{\neg}\, s$ is equivalent to s.

- **TS2.** $\mathsf{Holds}(Sv, Tv) \Leftrightarrow \mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tv)$

Now, given **AS4** and **TS2**, we can show the following necessary and sufficient condition for a state's failing to hold.

- **TS3.** $\neg\mathsf{Holds}(Sv, Tv) \Leftrightarrow \exists Tv'[Tv' \sqsubseteq Tv \land \mathsf{Holds}(\dot{\neg}\, Sv, Tv')]$

Note what this is saying. It says that the only situation where a state s fails to hold over an interval is when there is a sub-interval over which $\dot{\neg}\, s$ holds. Similarly, the only situation where $\dot{\neg}\, s$ fails to hold over an interval is when there is a sub-interval over which s holds. Thus, *both* s and $\dot{\neg}\, s$ fail to hold over an interval t if and only if there are two sub-intervals, t' and t'', of t such that s holds over t' and $\dot{\neg}\, s$ holds over t''. This is exactly the situation depicted in Figure 3.7. Are there other situations where both s and $\dot{\neg}(s)$ fail to hold? Some hold that this is the case for *transitions* between the two states (for example, (Herweg, 1991b)).[21] Nevertheless, I choose to remain silent about this issue; the system will be developed in such way that we would not need to decide what happens during a transition. Note that this is consistent with our interpretation of the system as representing commonsense knowledge of an agent; agents do not know what happens during transitions, and the vast literature on this matter attests to that.

So far, I have not explicitly specified the domain and range of $\dot{\neg}$, implying that both are the entire set of states, $\mathcal{S}$. This is technically *correct*, but is also very weak. Such a characterization of $\dot{\neg}$ fails to capture strong intuitions that we have about complements. Assuming that *"whales are fish"* is eternal, what about *"whales are not fish"*? Intuitively, we have more to say about it than its merely being a state. In particular, *"whales are not fish"* should be an eternal state, for, otherwise, *"whales are fish"* may hold non-eternally. The same intuition applies to all eternal states, and similar claims could be made about other state sorts as well. Therefore, rather than interpreting $\dot{\neg}$ vaguely as a function over $\mathcal{S}$, I will overload it such that, for each sub-sort Σ of $\mathcal{S}$ (so far we only have three, but see below), $\dot{\neg}$ would have a signature with Σ as its domain and (generally) the union of a collection of subsorts of $\mathcal{S}$ as its range. Let us start with ETERNAL.

- $\dot{\neg}$: ETERNAL $\longrightarrow$ ETERNAL.

Given the above definition, the complement of an eternal state is restricted to be eternal. This is intuitive enough and should not pose problems to the soundness of our logic. In fact, we can prove the following result.

- **TES3.** $\exists Tv \mathsf{Holds}(\dot{\neg} Sv^e, Tv) \Rightarrow \forall Tv \mathsf{Holds}(\dot{\neg} Sv^e, Tv)$

That is, the complement of an eternal state satisfies **AES1**. Note, however, that this is *not* a proof that the complement of an eternal state is eternal, since **AES1** is not a sufficient condition for eternal states (see Section 3.7).

Continuing with our intuitions, the complement of a temporary state is, in general, also temporary—a state unrestricted by any stability axioms.

- $\dot{\neg}$: TEMP $\longrightarrow$ TEMP.

[21] See (Kamp, 1979) for a general discussion of this issue.

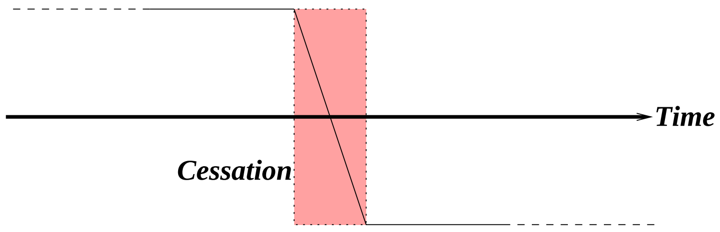

Figure 3.8: The complement of a permanent state.

	-onset	*+onset*
-cessation	ETERNAL	PERM
+cessation	?	TEMP

Figure 3.9: State types characterized by their onset-cessation possibilities.

For permanent states, things are not that straight-forward. Similar to **TES3**, we can prove the following result.

- **TPS1.** $\forall Tv[\mathsf{Holds}(\dot{\neg}\, Sv^p, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]]$

TPS1 states that if the complement of a permanent state holds over some time, then it holds over all preceding times. Both **TES3** and **TPS1** may be illustrated by simply *flipping* the timing diagrams of figures 3.4 and 3.5 (same for the assumption that $\dot{\neg}$ maps TEMP into TEMP). What sort of state is the complement of a permanent state? As per **TPS1**, such a state necessarily behaves as shown in Figure 3.8. It could be easily shown that the complement of a permanent state is neither eternal nor permanent. Is it temporary? If it is, it would be a very special type of temporary state; being extended indefinitely back in time makes it unintuitive to assume that it is temporary. One possibility is to assume a fourth sort of states, say CO-PERM, to which the complements of permanent states belong. Another insight may provide further merit for this assumption. Note that one may characterize the different types of states by the possibility of their having onsets and/or cessations—the patterns of change they are susceptible to. This is shown in the form of a matrix in Figure 3.9. Clearly, the empty slot could be readily filled by CO-PERM states. For completeness, the following axiom characterizes CO-PERM states (where the super-script *cp* indicates co-permanence).

- **AcPS1.** $\forall Tv[\mathsf{Holds}(Sv^{cp}v, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(Sv^{cp}v, Tv')]]$

Similar to **TPS1**, we can show that the complement of a CO-PERM state satisfies **APS1**.

- **TcPS1.** $\forall Tv[\mathsf{Holds}(\dot{\neg}\, S^{cp}, Tv) \Rightarrow \forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(\dot{\neg}\, S^{cp}, Tv')]]$

Armed with intuitions, **TPS1**, and **TcPS1**, we can now complete the characterization of $\dot{\neg}$.

- $\dot{\neg}$: PERM $\longrightarrow$ CO-PERM.
- $\dot{\neg}$: CO-PERM $\longrightarrow$ PERM.

Although, logically, CO-PERM is an independent distinctive sort, it seems that commonsense and, especially, language do not provide much support for this hypothesis. English, for example, does not provide a simple way for expressing CO-PERM states; outright negation of permanent states seems to be the only way of expressing them. To further illustrate this point, consider linguistic tests that could be used to distinguish sentences reporting different types of states. First, one may distinguish eternal states using the "Now" test.

(3.29) Now, Fermat is dead.

(3.30) Now, the Vietnam war is over.

(3.31) Now, the litmus paper is red.

(3.32) Now, Mary knows the answer.

(3.33) *Now, whales are fish.

(3.34) *Now, blood is red.

What "now" does, is that it (with a particular intonation) presupposes a period of time, preceding the time of the utterance, over which the reported state did not hold. This directly rules out eternal states.

Temporary states may be distinguished from eternal and permanent states by inserting "still" in a given sentence.

(3.35) The litmus paper is still red.

(3.36) Mary still knows the answer.

(3.37) ?Fermat is still dead.

(3.38) ?Whales are still fish.

The construct *still* directs attention to a possible period following the cessation of the state. The notion of the state ceasing to hold is implied and thus the awkwardness of (3.37) and (3.38).

In general, a state that may start to hold (+*onset*) would pass the "Now" test, and one that may cease to hold (+*cessation*) would pass the "still" test. A sentence denoting a CO-PERM state will have to fail the former and pass the latter. Evidently, it is very hard to come up with sentences that would exhibit such a behavior. One possible example is the following.[22]

(3.39) The temperature of the universe is above absolute zero.

(3.40) ?Now, the temperature of the universe is above absolute zero.

(3.41) The temperature of the universe is still above absolute zero.

It should be noted, however, that the co-permanence of (3.39) is not at all *logical* (cf. *logical permanence* in Section 3.23); it is contingent on the laws of physics. Similar to the situation with PERM states, it seems that the only example of a logically CO-PERM state is the state of some event being in the future (with PERM states, the event is in the past).

3.5.5 The Mereology of States

To form complex-state expressions, we need to introduce some operator to model the notion of mereological fusion of states (see Definition 3.2). For states, mereological fusion is essentially the counterpart of propositional conjunction; I will use $\overset{\bullet}{\wedge}$ to represent it.

- **AS5.** $\mathsf{Holds}(Sv_1 \overset{\bullet}{\wedge} Sv_2, Tv) \Leftrightarrow [\mathsf{Holds}(Sv_1, Tv) \wedge \mathsf{Holds}(Sv_2, Tv)]$

The above axiom states that, unlike with negation, Holds is transparent to $\overset{\bullet}{\wedge}$. Using this fact, we can prove that, like $\wedge$, $\overset{\bullet}{\wedge}$ is cumulative, associative, and idempotent.[23]

[22] The example was suggested by Stuart Shapiro in personal communication.

[23] Idempotence means that, $s \overset{\bullet}{\wedge} s$ holds if and only if s holds.

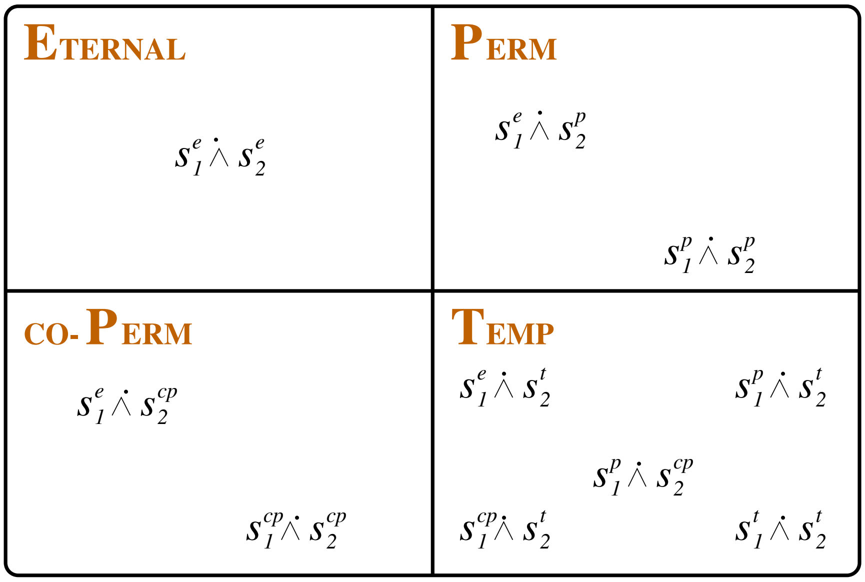

Figure 3.10: The ten signatures of $\overset{\bullet}{\wedge}$.

Following the same strategy adopted in the previous section, I will overload $\overset{\bullet}{\wedge}$ so that it has a distinct range for every possible domain. However, since the domain of $\overset{\bullet}{\wedge}$ is necessarily the cross-product of two sub-sorts of $\mathcal{S}$, we technically have sixteen possible signatures, not four as in the simple case of $\overset{\bullet}{\neg}$. But, given the cumulativity of $\overset{\bullet}{\wedge}$, we only need to consider ten signatures. Even then, it is a cumbersome task to list the ten signatures given that we can summarize them under the fold of a single intuition. Recall that $s_1 \overset{\bullet}{\wedge} s_2$ holds only when *both* s_1 and s_2 do (**AS5**). Therefore, should either of them cease to hold, their conjunction would also have to cease. Thus, if either s_1 or s_2 has the $+$*cessation* feature, so will $s_1 \overset{\bullet}{\wedge} s_2$. On the other hand, if both s_1 and s_2 have the $-$*cessation* feature, then so will their conjunction. Similarly, $s_1 \overset{\bullet}{\wedge} s_2$ have the $+$*onset* feature if and only if either s_1 or s_2 does. The ten different signatures are schematized in Figure 3.10.

We can readily prove the following theorems to support the distribution depicted in Figure 3.10. Note that there are no theorems corresponding to the TEMP corner, since no axioms characterize TEMP states.[24]

- **TEES.** $\exists Tv \mathsf{Holds}(Sv_1^e \overset{\bullet}{\wedge} Sv_2^e, Tv) \Rightarrow \forall Tv \mathsf{Holds}(Sv_1^e \overset{\bullet}{\wedge} Sv_2^e, Tv)$

[24]But maybe we can use counter examples to eliminate the possibility of a conjunction with a TEMP component being anything but TEMP.

- **TEPS.** $\forall Tv[\mathsf{Holds}(Sv_1^e \overset{\bullet}{\wedge} Sv_2^p, Tv) \Rightarrow$
$$\forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(Sv_1^e \overset{\bullet}{\wedge} Sv_2^p, Tv')]]$$
- **TPPS.** $\forall Tv[\mathsf{Holds}(Sv_1^p \overset{\bullet}{\wedge} Sv_2^p, Tv) \Rightarrow$
$$\forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(Sv_1^p \overset{\bullet}{\wedge} Sv_2^p, Tv')]]$$
- **TEcPS.** $\forall Tv[\mathsf{Holds}(Sv_1^e \overset{\bullet}{\wedge} Sv_2^{cp}, Tv) \Rightarrow$
$$\forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(Sv_1^e \overset{\bullet}{\wedge} Sv_2^{cp}, Tv')]]$$
- **TcPcPS.** $\forall Tv[\mathsf{Holds}(Sv_1^{cp} \overset{\bullet}{\wedge} Sv_2^{cp}, Tv) \Rightarrow$
$$\forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(Sv_1^{cp} \overset{\bullet}{\wedge} Sv_2^{cp}, Tv')]]$$

Using only $\overset{\bullet}{\neg}$ and $\overset{\bullet}{\wedge}$, we can define other operators similar to the propositional $\vee$, $\Rightarrow$, and $\Leftrightarrow$. Care should be taken though, since these operators employ $\overset{\bullet}{\neg}$ in their definition and, thus, cannot freely move across Holds. To illustrate, let us define $\overset{\bullet}{\vee}$ so that it satisfies the following axiom, akin to the corresponding DeMorgan's law.

- **AS6.** $\mathsf{Holds}(Sv_1 \overset{\bullet}{\vee} Sv_2, Tv) \Leftrightarrow \mathsf{Holds}(\overset{\bullet}{\neg}[\overset{\bullet}{\neg} Sv_1 \overset{\bullet}{\wedge} \overset{\bullet}{\neg} Sv_2], Tv)$

By defining $\overset{\bullet}{\vee}$ this way, we are forcing a certain interpretation of its relation to $\vee$. The following result comes from (Allen, 1984).

- **TS4.** $\mathsf{Holds}(Sv_1 \overset{\bullet}{\vee} Sv_2, Tv) \Leftrightarrow$
$$\forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]$$

Admittedly, as (Shoham, 1987) points out, the above theorem would come out in a much simpler form had I adopted a point-based ontology of time. Intuitively, $s_1 \overset{\bullet}{\vee} s_2$ may hold over an interval even if neither s_1 nor s_2 do, provided that at least one of them holds over any *atomic* subinterval (or point) thereof. The reason why there is a need for the doubly-nested use of the sub-interval relation is the two-level negation employed in the definition of $\overset{\bullet}{\vee}$. In particular, the following is not correct.

(3.42) $\mathsf{Holds}(Sv_1 \overset{\bullet}{\vee} Sv_2, Tv) \Leftrightarrow \forall Tv'[Tv' \sqsubseteq Tv \Rightarrow [\mathsf{Holds}(Sv_1, Tv') \vee \mathsf{Holds}(Sv_2, Tv')]]$

This is illustrated in Figure 3.11, where $s_1 \overset{\bullet}{\vee} s_2$ holds over t, but neither s_1 nor s_2 holds over t' (or over t for that matter).

Using the standard definitions of propositional connectives, we can likewise define $\overset{\bullet}{\Rightarrow}$ and $\overset{\bullet}{\Leftrightarrow}$. For completeness, we can introduce quantifying operators similar to $\forall$ and $\exists$.

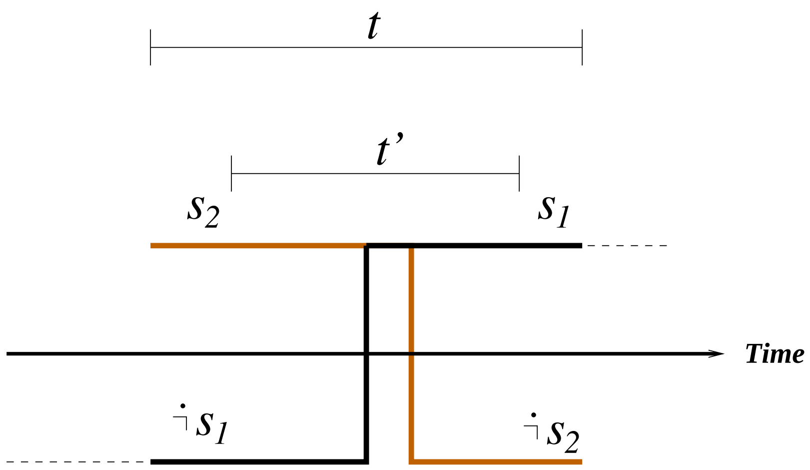

Figure 3.11: $s_1 \dot{\vee} s_2$ holds over t, but neither s_1 nor s_2 holds over t'.

- **AS7.** $\mathsf{Holds}(\dot{\forall} x[Sv], Tv) \Leftrightarrow \forall x[\mathsf{Holds}(Sv, Tv)]$
- **AS8.** $\mathsf{Holds}(\dot{\exists} x[Sv], Tv) \Leftrightarrow \forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge \exists x[\mathsf{Holds}(Sv, Tv'')]]]$

The above axioms are justified by the conjunctive and disjunctive interpretations of $\dot{\forall}$ and $\dot{\exists}$, respectively.

3.6 Propositions as States

In section 3.5.3, the English "whales are fish" was introduced as an example of a sentence representing an eternal state. To recapitulate, an eternal state is a state of affairs that either always holds or never holds (or is so conceived by an agent). "Whales are fish" represents a state of affairs that never holds, namely the state of whales' being fish. Formally, we might attempt to represent such a state by the following expression.

(3.43) Isa(whale, fish).

But now there is something curious about (3.43). Is there a reason, logical or ontological, why Isa should be a function to ETERNAL and not to $\mathcal{P}$? In other words, by looking at (3.43) with the intended meaning in mind, it is not clear why it denotes an eternal state rather than a proposition. The same applies to $\mathcal{FOCS}$ expressions such as $Tc_1 \sqsubseteq Tc_2$ or $Tc_1 \prec Tc_2$. Why do these forms represent propositions rather than eternal states? Does it matter? Are there differences between propositions and eternal states, or between propositions and states in general? A proposition is the kind of entity that may be believed (denied, asserted, wished, etc.) by an agent. When an agent believes a proposition, it holds a belief *about* something. Assuming that Isa is a function to $\mathcal{P}$, believing Isa(whale, fish), implies having a belief about whales and fish. But it also implies having a belief about a specific relation between them. Similarly, believing $Tc_1 \sqsubseteq Tc_2$ implies having the belief that the parthood relation exists between Tc_1 and Tc_2. The existence of a relation among entities is basically what a state is. Thus, in a general sense, an *atomic* proposition, one involving no connectives and/or quantifiers, describes some state holding.

A formal rendering of the last sentence may be reflected by adopting the principle that all atomic propositions are of the form "$\mathsf{Holds}(s,t)$". In a sense, that's the point. However, note that what Holds does is denote a relation between a state and a time. Now, the existence of such a relation is itself a state, an eternal state of some particular state holding over some particular time. What this amounts to is that, for every atomic proposition, there is a corresponding *eternal* state that it describes. Conversely, every *simple* eternal state has a unique corresponding atomic proposition. To see this, note that unlike temporary states, for instance, eternal states do not start, do not cease, and particularly, cannot be perceived. I can see John's crossing the street, but never his being mortal. Note that I may come to *believe* that John is mortal based on my perception of some temporary state holding. Because of their peculiar temporal properties (which amount to their being atemporal), eternal states can never be experienced; they may only be *believed* to be holding, *wished* to be holding, *denied* to be holding, etc. That is, an agent can only have what have been called *propositional* attitudes toward eternal states. But since this *is* the property defining propositions (see above), then the set of simple eternal states maps one-to-one to that of atomic propositions. Thus, the two sets stand in a one-to-one correspondence; everything that may be said of an atomic proposition can be said about a simple eternal state with a slight shift in perspective. In particular, Cassie's believing an atomic proposition is identical to her believing that some simple eternal state holds.

But the result can be extended to complex eternal states and non-atomic propositions. In particular, note that, given **AS5**, a proposition made up of the conjunction of atomic propositions, corresponds to a complex eternal state constructed using $\dot{\wedge}$. Negation, however, does not carry its traditional propositional semantics into the domain of states. That is, **AS4** does not represent a strong equivalence between $\neg$ and $\dot{\neg}$ in the same way that **AS5** does between $\wedge$ and $\dot{\wedge}$. Nevertheless, the equivalence can be established for the special case of ETERNAL states.

- **TES4.** $\mathsf{Holds}(\dot{\neg}\, S^e v, T v) \Leftrightarrow \neg\mathsf{Holds}(S^e v, T v)$

Thus, since the complement of an eternal state is eternal, every negated proposition has a corresponding eternal state (and vice versa). We can similarly extend the correspondence to propositions involving other connectives and quantifiers. For this reason, I shall henceforth drop the distinction between propositions and eternal states. Formally, I will dismiss the sort $\mathcal{P}$, using ETERNAL instead, with the belief space, β, being a subset of ETERNAL—the set of eternal states that Cassie believes to be holding. Informally, I will use the terms "proposition" and "eternal state" interchangably, preferring one over the other depending on the context. By this move, we get two things: (i) more ontological economy and (ii) a logic in which temporality is the default, eternal states emerging only as a special case of temporal situations. In addition, we can dismiss all the propositional connectives and quantifiers, replacing them with their state-based counterparts. I will do this by simply dropping the dot from the state operators, and use the familiar $\neg$, $\wedge$, $\vee$, $\Rightarrow$, $\Leftrightarrow$, $\forall$, and $\exists$ for the set of connectives and quantifiers forming state expressions.

One vexing notational issue remains to be settled. What does an assertion of the form $\mathsf{Holds}(s,t)$ convey? Two things. First, it tells us that the state s holds, and, second, that it holds over time t. Now, suppose that s is an eternal state. In this case, only the first piece of information is interesting, for, by **AES1**, if an eternal state ever holds, it always holds. Thus, the second argument of Holds in the case of an ETERNAL first argument is actually redundant; nothing is special about any particular time when it comes to eternal states. Thus, we might employ a unary version of Holds so that terms of the form $\mathsf{Holds}(s^e)$ are valid. Following the dismissal of $\mathcal{P}$, such terms are themselves ETERNAL, which means that $\mathcal{FOCS}$ can generate obnoxious-looking terms of the form $\mathsf{Holds}(\mathsf{Holds}(\ldots(\mathsf{Holds}(s^e)\ldots))$. Such terms are certainly syntactically valid, but they make little semantic sense. There is no *interesting* conceptual distinction between $\mathsf{Holds}(\mathsf{Holds}(s^e))$ and $\mathsf{Holds}(s^e)$.[25] In fact, there is no major difference between $\mathsf{Holds}(s^e)$ and s^e, especially at the level of commonsense reasoning in which I am interested. In particular, note that it does not make sense for Cassie to believe one and not the other; either both are in β, or neither is. In truth-functional terms,

[25] In this respect, the unary Holds resembles a truth connective in a deflationary theory of truth (see (Stoljar, 1997)) or the modal operators of the system *S5* (Hughes and Cresswell, 1996).

the unary version of Holds seems to be the complement of negation—the identity mapping. Hence, we might add the following axiom to $\mathcal{FOCS}$.

- $\mathsf{Equiv}(\mathsf{Holds}(S^e v), S^e v)$

However, I will resort to a more convenient measure; I will drop the unary Holds altogether. That is, as a matter of notation, I will use s^e to represent $\mathsf{Holds}(s^e)$. This will both simplify the notation, and give ETERNAL-terms a more traditional, proposition-like look.

3.7 The Problem with Sufficient Conditions

Now is a good time to explain what axioms like **APS1** and **AES1** represent. They represent *necessary* conditions to be satisfied by PERM and ETERNAL states, respectively. Note, however, that these axioms do not represent sufficient conditions; just because a state happens to satisfy **APS1** (**AES1**) does not mean that it is permanent (eternal). For example, suppose that a particular rock, r, is located at some particular location, l, somewhere in the desert. Further, suppose that this state of affairs is represented by the state-term $s = \mathsf{In}(r,l)$. What sort of state is s? Intuitively, it should be temporary, since it is possible for some agent to move r to some other location. But suppose that this never happens. That is, at all times r is located at l. In such a model, *s behaves* exactly as an eternal state and indeed satisfies **AES1**. Yet, that does not make it eternal, for, to be eternal, it must be *impossible* for s to cease to hold.

Similarly, just because some state satisfies **APS1** does not mean that it is necessarily permanent. First, one can construct a model in which the state $\mathsf{In}(r,l)$ from above satisfies **APS1**. Second, **APS1** is implied by **AES1**, so, for all we know, a state that satisfies **APS1** might actually be an eternal state (and not just happens to behave as one). For a state to be permanent, it should be *impossible* for it to cease (thus dismissing temporary states) and it should be *possible* that it once did not hold (thus dismissing eternal states).

It, therefore, seems that, in order to accommodate sufficient conditions, one needs to represent the notion of possibility, thereby stepping out of the coziness of ordinary first-order logic into the realm of modal logic. In modal logic, the necessary conditions **AES1**, **APS1**, and **AcPS1** may be replaced by the following necessary-and-sufficient conditions, respectively.

- **AES2.** $\Box[\exists Tv[\mathsf{Holds}(S^e v, Tv)] \Rightarrow \forall Tv[\mathsf{Holds}(S^e v, Tv)]]$.

- **APS2**. $\Box[\forall Tv[\mathsf{Holds}(S^p v, Tv) \Rightarrow \forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(S^p v, Tv')]]]$
  $\wedge\Diamond[\exists Tv, Tv'[Tv' \prec Tv \wedge \mathsf{Holds}(S^p v, Tv) \wedge \neg\mathsf{Holds}(S^p v, Tv')]]$.

- **AcPS2**. $\Box[\forall Tv[\mathsf{Holds}(S^{cp} v, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(S^{cp} v, Tv')]]]$
  $\wedge\Diamond[\exists Tv, Tv'[Tv \prec Tv' \wedge \neg\mathsf{Holds}(S^{cp} v, Tv) \wedge \mathsf{Holds}(S^{cp} v, Tv')]]$.

Using modal logic, we can, in fact, introduce a necessary-and-sufficient condition for temporary states.

- **ATS1**. $\Diamond[\exists Tv_1, Tv_2[Tv_1 \prec Tv_2 \wedge \neg\mathsf{Holds}(S^t v, Tv_1) \wedge \mathsf{Holds}(S^t v, Tv_2)]]$
  $\wedge\Diamond[\exists Tv_3, Tv_4[Tv_3 \prec Tv_4 \wedge \mathsf{Holds}(S^t v, Tv_3) \wedge \neg\mathsf{Holds}(S^t v, Tv_4)]]$.

Essentially, the first conjunct of **ATS1** asserts that temporary states may start to hold, the second asserts that they may cease. Note that this is exactly how TEMP states were characterized using the *onset* and *cessation* binary features in Figure 3.9. In fact, with simple logical manipulation, we can rewrite **AES2**, **APS2**, and **AcPS2** in such a way that the distribution of the *onset* and *cessation* features is readily readable from the axioms.

- **AES2**. $\neg\Diamond[\exists Tv_1, Tv_2[Tv_1 \prec Tv_2 \wedge \neg\mathsf{Holds}(S^e v, Tv_1) \wedge \mathsf{Holds}(S^e v, Tv_2)]]$
  $\wedge\neg\Diamond[\exists Tv_3, Tv_4[Tv_3 \prec Tv_4 \wedge \mathsf{Holds}(S^e v, Tv_3) \wedge \neg\mathsf{Holds}(S^e v, Tv_4)]]$.

- **APS2**. $\Diamond[\exists Tv_1, Tv_2[Tv_1 \prec Tv_2 \wedge \neg\mathsf{Holds}(S^p v, Tv_1) \wedge \mathsf{Holds}(S^p v, Tv_2)]]$
  $\wedge\neg\Diamond[\exists Tv_3, Tv_4[Tv_3 \prec Tv_4 \wedge \mathsf{Holds}(S^p v, Tv_3) \wedge \neg\mathsf{Holds}(S^p v, Tv_4)]]$.

- **AcPS2**.$\neg\Diamond[\exists Tv_1, Tv_2[Tv_1 \prec Tv_2 \wedge \neg\mathsf{Holds}(S^{cp} v, Tv_1) \wedge \mathsf{Holds}(S^{cp} v, Tv_2)]]$
  $\wedge\Diamond[\exists Tv_3, Tv_4[Tv_3 \prec Tv_4 \wedge \mathsf{Holds}(S^{cp} v, Tv_3) \wedge \neg\mathsf{Holds}(S^{cp} v, Tv_4)]]$.

Using the above axioms, rather than posit, we can now *prove* that the complements of ETERNAL, PERM, CO-PERM, and TEMP states are ETERNAL, CO-PERM, PERM, and TEMP, respectively. We can also prove similar results for complex (conjunctive) states. Although this is a desirable feature, we should keep in mind what it involves—modality. As should now be clear, using modal logic would certainly complicate the exposition, infesting $\mathcal{FOCS}$ with axioms and rules of inference for the modal operators $\Box$ and $\Diamond$. In addition, we would need to provide semantics for the modal operators that fit within our state-based, intensional framework.

One way to salvage the situation, and still remain within the confines of ordinary first-order logic, is to reify possible worlds, moving them from the semantic to the syntactic realm (in the spirit of (Moore, 1988)). For example, **ATS1** may be replaced by the following axiom, where $\mathsf{HoldsIn}$ is a function from eternal states and possible worlds to eternal states.

- **ATS1'**. $\exists Wv \mathsf{HoldsIn}(\exists Tv_1, Tv_2[Tv_1 \prec Tv_2 \wedge \neg \mathsf{Holds}(S^t v, Tv_1) \wedge \mathsf{Holds}(S^t v, Tv_2)], Wv) \wedge \exists Wv' \mathsf{HoldsIn}(\exists Tv_3, Tv_4[Tv_3 \prec Tv_4 \wedge \mathsf{Holds}(S^t v, Tv_3) \wedge \neg \mathsf{Holds}(S^t v, Tv_4)], Wv')$.

But such a move is not any better than using modal logic. For it involves introducing a new sort for possible worlds, figuring out the semantics of that sort, and, more importantly, introducing axioms for reasoning about the Kripkian accessibility relation (Hughes and Cresswell, 1996). This might all be important and useful, but it will certainly take us far beyond the limits of this research. I choose, then, to adhere to my previous position, adopting only the necessary conditions **AES1**, **APS1**, and **AcPS1**, and assigning states to sorts by fiat. This certainly misses the benefits gained by introducing sufficient conditions (being able to prove intuitions about the sorts of complement and conjunctive states), but is sufficient for the purpose of my investigation.

To support my claim, recall that $\mathcal{FOCS}$ axioms are taken to be beliefs of the *acting* cognitive agent Cassie. As such, an axiom is necessary in as much as it is useful for Cassie. Axioms **AES2**, **APS2**, **AcPS2**, and **ATS1** (or their variants using reified possible worlds) are useful for inferring the sorts of complement and conjunctive states. But such results are once-and-for-all results; they make claims about whole sorts of states, not particular instances. Using **AES2**, for example, once Cassie infers that the complement of an eternal state is eternal, she can repeatedly use *that* result to infer the eternity of complements of particiular eternal states; **AES2** is no longer needed in this regard. Thus, whether we allow Cassie to infer the sorts of complements, or build that into the system is a matter of demonstrating the ability to achieve logical elegance (which has been done above), but is of no practical relevance.

But sufficient conditions might be useful for more than just inferring the sorts of complements and conjunctions. In general, by inferring a sufficent condition, necessary conditions follow, and that is useful. However, in our case, there are two reasons why the argument does not hold. First, note that all our sufficient conditions are themselves the most useful necessary conditions. Therefore, by inferring a sufficient condition, Cassie would have inferred all the important results. Thus, introducing sufficient conditions does not get us anything for free. Second, our sufficient conditions involve the notion of possibility, and I believe that inferring such conditions would require analogical or inductive modes of reasoning that call for a full investigation of these issues. Since this is certainly beyond the scope of my interest, I shall step aside and leave the issue for future research.

3.8 Conclusions

In this chapter, we have developed a formal theory of states. We have identified what states are and how a cognitive agent may come to conceive of them. Four types of states have been distinguished based on the patterns of change they are susceptible to. Eternal states neither start, nor cease; permanent states may start, but cannot cease; co-permanent states cannot start, but may cease; and temporary states may repetitively start or cease. We have concluded that eternal states can fill the role of propositions and have extended the standard propositional connectives and quantifiers into the domain of states. In the next chapater we will be looking at how the theory of states developed here can form the basis of a theory of events.

Chapter 4

The Logic of Events

So far, I have considered only states—temporally-unbounded situations that homogeneously hold over time. With states alone, there is a lot of useful information that may be represented and a lot of interesting inferences that a cognitive agent can make. Nevertheless, more is needed. Consider, for example, sentence (3.2), *I crossed the street*. It represents a particular belief of a cognitive agent. As per Section 3.6, beliefs are all about eternal states holding. At least on one reading, the eternal state here is that of a particular street-crossing situation taking place over a particular time. Consider the following $\mathcal{FOCS}$ representation of (3.2).

(4.1) $\mathsf{Holds}(\mathsf{Cross}(\mathrm{I},\ \mathrm{STREET}), t)$.

Such a representation interprets the situation described by (3.2) as a state. This, however, is obviously wrong, for then (4.1) implies that the situation reported by (3.2) took place over all subintervals of $[\![t]\!]$. (4.1) is a possible representation for *I was crossing the street* not *I crossed the street*. As pointed out in Section 3.2, sentence (2) describes a situation as an *event*—a temporally-bounded situation wholly occurring at some time.

In this chapter, I shall introduce and axiomatize events. The sort $\mathcal{E}$ for events is partitioned into two subsorts: $\overset{\bullet}{\mathcal{E}}$ for *punctual events* and $\overset{\bullet\!-\!\bullet}{\mathcal{E}}$ for *durative events*. The distinction between these two shall be made clear below (it is not as simple as it may seem). For now, however, it would be helpful to think of a punctual event as the onset or the cessation of some state. A durative event, on the other hand, is made up of the punctual event of some state starting to hold, the state's holding for a while, and the punctual event of the state ceasing. Thus, punctual events represent the temporal boundaries of durative events. A member of $\mathcal{D}_{\mathcal{E}}$ is a topologically-closed situation: it is either durative, and therefore has its boundaries as parts, or punctual and, thus, a boundary, which is itself closed according to classical topology. Before formalizing these intuitions, however, let me first draw the distinction between types and tokens of events.

4.1 Types and Tokens of Events

The study of categories and category systems is among the classical favorites of cognitive science (Rosch, 1978; Lakoff, 1987, for instance). Category systems prove to be quite interesting and complicated. Here, however, I will only consider a simple, fairly unproblematic aspect of categorization, namely that our minds group entities into categories. Categories are conceptual structures that capture particular aspects of entities we encounter in our day-to-day experience. Among the categorizable entities are events. When we perceive an event, we seem to have no difficulty reporting it as one of crossing the street, of moving a box, or of buying a bottle of milk. Why and how this is done is an issue that falls outside the scope of my investigation. Nevertheless, the seemingly uncontroversial categorizability of events is something that the event-component of $\mathcal{FOCS}$ ought (and need) to account for.

Let $\mathcal{EC}$ be a $\mathcal{FOCS}$ sort for event categories.[1] $\mathcal{EC}$ is partitioned into two components, $\overset{\bullet}{\mathcal{EC}}$, for categories of punctual events, and $\overset{\bullet\!-\!\bullet}{\mathcal{EC}}$, for categories of durative events. The basic act of categorizing an event will be represented by the function symbol Cat.

- $\mathsf{Cat}\colon (\overset{\bullet}{\mathcal{E}} \times \overset{\bullet}{\mathcal{EC}}) \cup (\overset{\bullet\!-\!\bullet}{\mathcal{E}} \times \overset{\bullet\!-\!\bullet}{\mathcal{EC}}) \longrightarrow$ ETERNAL, where $[\![\mathsf{Cat}(e, ec)]\!]$ is the eternal state of event $[\![e]\!]$ belonging to event category $[\![ec]\!]$.[2]

To illustrate some of the important features of the categorizing function Cat, and for notational convenience, I will introduce a *meta-theoretical* function, $\mathcal{C}$, from events to sets of event categories (recall that β represents the set of propositions that Cassie believes).[3]

[1] These are similar to Galton's *event radicals* (Galton, 1984; Galton, 1987a). However, I opt for the cognitive, rather than the logical, terminology.

[2] Some readers might be sceptical about the inclusion of *both* event tokens and event categories as independent sorts in the ontology. I will not try to defend this position at this point; as the discussion proceeds it will become apparent why we need both sorts. For arguments supporting this choice, however, see (Link, 1987).

[3] Whether $\mathcal{C}$ should be defined in the semantic or syntactic domain (which I here choose) is an interesting issue. However, given the isomorphism between SNePS terms and their denotata, nothing major seems at stake by adopting either strategy.

- $\mathcal{C} : \mathcal{E} \longrightarrow 2^{\mathcal{EC}}$, where $\mathcal{C}(e) = \{ec : \beta \vdash \mathsf{Cat}(e, ec)\}$.

Three points to note. First, $\mathcal{C}$ maps an event to a *set* of categories. This is important, since one and the same event may be viewed from different perspectives. Examples may include an event of running to the store that is also an event of running, an event of shooting that is also an event of killing, etc. Of course, this invokes issues of event individuation (Davidson, 1970; Goldman, 1970; Pfeifer, 1989, for instance), which is beyond the scope of this work. However, allowing an event to be associated with more than one category seems to be the less controversial choice over associating it with a single category.

Second, the mapping is not *into*; two events may belong to exactly the same categories. Ultimately, however, any two events are distinct, particularly because of some categories being suitable for one and not the other. Such categories may be as exotic as the categories of events occurring over a particular time.[4] Nevertheless, because of my commitment to a common-sense epistemic ontology, I shall not thus constrain $\mathcal{C}$.

Third, the image of $\mathcal{C}$ does not cover $\mathcal{EC}$. This simply allows for categories that an agent may conceive of without ever conceiving of particular instances.[5]

One may similarly define a function that would map an event category to its instances.

- $\mathcal{I} : \mathcal{EC} \longrightarrow 2^{\mathcal{E}}$, where $\mathcal{I}(ec) = \{e : \beta \vdash \mathsf{Cat}(e, ec)\}$.

Unlike $\mathcal{C}$, the image of $\mathcal{I}$ covers $\mathcal{E}$, since every event has to be given at least one categorization to be conceived of in the first place. Sentences like *"Something happened at 3 p.m. yesterday"* might cause one to be suspicious about such an assumption. Nevertheless, one possible (though strange) category for whatever *"something"* refers to is the category of events that happened at the time indicated. Like $\mathcal{C}$, $\mathcal{I}$ is not *into*. If we think of $\mathcal{I}$ as mapping event intensions (categories) to event extensions (sets of events), then its not being one-to-one is reasonable, since two intensions may have the same extension (the President of the US visiting Greece in 1999 and Bill Clinton visiting Greece in 1999).

Unlike states, event tokens (elements of $\mathcal{E}$) do not *hold* homogeneously over time; rather, they occur as complete wholes. We, therefore, need some function, other than Holds, to relate events to

[4]Less exotic are Galton's *once-only* categories (Galton, 1984).

[5]This is stronger than just saying that $\mathcal{C}$ is not onto. Being onto is unnatural since it would imply that every possible collection of event categories share some event.

the times at which they occur.

- $\mathsf{Occurs}: \mathcal{E} \times \mathcal{T} \longrightarrow \textsc{Eternal}$, where $[\![\mathsf{Occurs}(e,t)]\!]$ is the eternal state of event $[\![e]\!]$ occurring over time $[\![t]\!]$.

The difference between Occurs and Holds is one of axiomatization. Event occurrence has a peculiar uniqueness property; an event token occurs only once.

- **AE1.** $[\mathsf{Occurs}(Ev, Tv_1) \wedge \mathsf{Occurs}(Ev, Tv_2)] \Rightarrow \mathsf{Equiv}(Tv_1, Tv_2)$

Given the above axiom, we can prove the following theorems stating that events are heterogeneous.

- **TE1.** $\mathsf{Occurs}(Ev, Tv) \Rightarrow \forall Tv'[Tv' \sqsubset Tv \Rightarrow \neg\mathsf{Occurs}(Ev, Tv')]$.

- **TE2.** $\mathsf{Occurs}(Ev, Tv) \Rightarrow \forall Tv'[Tv \sqsubset Tv' \Rightarrow \neg\mathsf{Occurs}(Ev, Tv')]$.

That is, if an event occurs over some interval, then it does not occur over any proper sub-interval, or super-interval, thereof. Some linguists (for example, (Partee, 1984; Herweg, 1991a; De Swart, 1998)) would reject **TE2**, asserting exactly its opposite: if an event occurs over some interval, then it occurs over all of its super-intervals. The point of such an assertion would be to account for the plausibility of sentences such as (4.2) where the trip to the store certainly did not take the whole day.

(4.2) John went to the store yesterday.

My inclusion of **AE1** (and, hence, **TE2**) explicates the particular way in which I am interpreting occurrence: an event occurs over an interval if and only if it fits exactly in that interval, not if it falls somewhere *within* it. This will become more precise as I discuss the internal part-structure of events. For now, however, let me introduce the following meta-theoretical *partial function* associating events with the times of their occurrence. This is the *temporal trace* function of (Link, 1987). Note that the function is partial since one can conceive of an event without it actually occurring.

- $\tau: \mathcal{E} \rightharpoonup \mathcal{T}$, where $\tau(e) = t$ if and only if $\beta \vdash \mathsf{Occurs}(e,t)$.

Both the temporal trace and the possible categorizations are properties that are shared by co-referential event terms. This is expressed by the following pair of axioms.

- **AE2.** $[\mathsf{Occurs}(Ev_1, Tv) \wedge \mathsf{Equiv}(Ev_1, Ev_2)] \Rightarrow \mathsf{Occurs}(Ev_2, Tv)$.

- **AE3.** $[\mathsf{Cat}(Ev_1, ECv) \wedge \mathsf{Equiv}(Ev_1, Ev_2)] \Rightarrow \mathsf{Cat}(Ev_2, ECv)$.

With the distinction between event types and tokens firmly in hand, we now turn to the topology of events.

4.2 Boundaries

Boundaries are punctual events that chop the temporal continuum into pieces, each characterizable by some state. In other words, boundaries are simply the onset and cessation events discussed in Section 3.5. Sentences denoting states are described by linguists as *unbounded*. The basic insight is that such sentences describe situations whose boundaries are out of the scope of attention.

> [A] speaker uses a[n unbounded] constituent to refer to an entity whose boundaries are outside the current field of view. This does not entail that the entity is absolutely unbounded in space or time; it is just that we can't see the boundaries from the present vantage point. (Jackendoff, 1991, p. 19)

I shall deviate from the standard linguistic terminology and use the standard topological *closed* and *open* for *bounded* and *unbounded*, respectively. The reason is that I would like to reserve the terms *bounded* and *unbounded* for those situations that Jackendoff refers to as "absolutely unbounded", namely eternal, permanent, and co-permanent states.

According to its type, a state s may be associated with event categories $\uparrow s$ and $\downarrow s$, representing categories of its onsets and cessations, respectively. More precisely:

- $\uparrow$: $\textsc{Perm} \cup \textsc{Temp} \xrightarrow{into} \dot{\mathcal{EC}}$, where $[\![\uparrow s]\!]$ is the category of onsets of state $[\![s]\!]$.
- $\downarrow$: $\textsc{co-Perm} \cup \textsc{Temp} \xrightarrow{into} \dot{\mathcal{EC}}$, where $[\![\downarrow s]\!]$ is the category of cessations of state $[\![s]\!]$.

The functions $\uparrow$ and $\downarrow$ are into; there is a unique event category for the onset/cessation of a state. I assume that the union of the images of $\uparrow$ and $\downarrow$ makes up most of $\dot{\mathcal{EC}}$. Except for what are called *mental* event categories (see Chapter 9), a punctual event category can only be a state transition. Note that this is a purely technical assumption. Traditionally, the term *punctual event* covers examples such as blinking, coughing, and flashing.[6] According to (Talmy, 2000), these events involve

[6] These are referred to as *semelfactives* by (Smith, 1997).

two transitions—one into a state and one out of it. The punctuality comes from the fact that, *linguistically*, whatever happens between the two transitions in not accessible (which probably stems from the typically brief duration of such events). I believe, however, that, *conceptually*, the state demarcated by the two transitions is accessible (the closed eyes, the contraction of the respiratory track muscles, and the flash light's being on), and I will, therefore, not include these events within $\dot{\mathcal{E}}$.

Some axioms are now necessary to illustrate the exact relation among s, $\uparrow s$, and $\downarrow s$. The first two come from standard topology: entities share their boundaries with their complements.

- **AOC1.** $\mathsf{Cat}(\dot{E}v, \uparrow Sv) \Leftrightarrow \mathsf{Cat}(\dot{E}v, \downarrow(\neg Sv))$.
- **AOC2.** $\mathsf{Cat}(\dot{E}v, \downarrow Sv) \Leftrightarrow \mathsf{Cat}(\dot{E}v, \uparrow(\neg Sv))$.

By chaining the above axioms, we can prove interesting theorems like the following (basically, the arrows flip with each application of $\neg$).

- **TOC1.** $\mathsf{Cat}(\dot{E}v, \uparrow Sv) \Leftrightarrow \mathsf{Cat}(\dot{E}v, \uparrow(\neg\neg Sv))$.
- **TOC2.** $\mathsf{Cat}(\dot{E}v, \downarrow Sv) \Leftrightarrow \mathsf{Cat}(\dot{E}v, \downarrow(\neg\neg Sv))$.

But **AOC1** and **AOC2** together with the two theorems do not tell us exactly how the occurrence of onsets and cessations relate to states. **AOC3** and **AOC4** axiomatize the *boundary-hood* of onsets, and **TOC3** and **TOC4** provide similar results for cessations. These could be seen as specifying necessary conditions for the occurrence of onsets and/or cessations, respectively. Intuitively, any transition is immediately preceded by a state and immediately followed by its complement (or vice versa).

- **AOC3.** $[\mathsf{Occurs}(\dot{E}v, Tv) \;\wedge\; \mathsf{Cat}(\dot{E}v, \uparrow Sv)] \;\Rightarrow$
  $$\exists Tv' [Tv \supset\subset Tv' \wedge \mathsf{Holds}(Sv, Tv')]$$
- **AOC4.** $[\mathsf{Occurs}(\dot{E}v, Tv) \;\wedge\; \mathsf{Cat}(\dot{E}v, \uparrow Sv)] \;\Rightarrow$
  $$\exists Tv' [Tv' \supset\subset Tv \wedge \mathsf{Holds}(\neg Sv, Tv')]$$
- **TOC3.** $[\mathsf{Occurs}(\dot{E}v, Tv) \;\wedge\; \mathsf{Cat}(\dot{E}v, \downarrow Sv)] \;\Rightarrow$
  $$\exists Tv' [Tv' \supset\subset Tv \wedge \mathsf{Holds}(Sv, Tv')]$$

- **TOC4.** $[\mathsf{Occurs}(\overset{\bullet}{E}v,Tv) \wedge \mathsf{Cat}(\overset{\bullet}{E}v,\downarrow Sv)] \Rightarrow$

$$\exists Tv[Tv \supset\subset Tv' \wedge \mathsf{Holds}(\neg Sv,Tv')]$$

The following axiom (**AOC5**) states a sufficient condition for the occurrence of onsets. Basically, Cassie may conclude that an onset of a state s has occurred if she is aware of a time over which s does not hold followed by another over which it does. Theorem **TOC5** states a similar result for cessations.

- **AOC5.** $[\neg\mathsf{Holds}(Sv,Tv_1) \wedge \mathsf{Holds}(Sv,Tv_2) \wedge Tv_1 \prec Tv_2] \Rightarrow$

$$\exists \overset{\bullet}{E}v,Tv_3[\mathsf{Cat}(\overset{\bullet}{E}v,\uparrow Sv) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v,Tv_3)$$
$$\wedge Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2].$$

- **TOC5.** $[\mathsf{Holds}(Sv,Tv_1) \wedge \neg\mathsf{Holds}(Sv,Tv_2) \wedge Tv_1 \prec Tv_2] \Rightarrow$

$$\exists \overset{\bullet}{E}v,Tv_3[\mathsf{Cat}(\overset{\bullet}{E}v,\downarrow Sv) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v,Tv_3)$$
$$\wedge Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2].$$

The signatures of $\uparrow$ and $\downarrow$ restrict the type of state to which they can apply. But even for those states, there are restrictions on the occurrence patterns of onsets, or cessations, that depend on the very nature of the state. The following pair of theorems state those restrictions dictated by the necessary conditions for PERM and CO-PERM states. Theorem **TOC6** states that once an onset of some PERM state occurs, no other onset can occur in the future. Theorem **TOC7** is the mirror image of **TOC6**.

- **TOC6.** $[\mathsf{Occurs}(\overset{\bullet}{E}v_1,Tv_1) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_1,\uparrow S^p v)] \Rightarrow$

$$\forall \overset{\bullet}{E}v_2,Tv_2[[Tv_1 \prec Tv_2 \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2,Tv_2)] \Rightarrow \neg\mathsf{Cat}(\overset{\bullet}{E}v_2,\uparrow S^p v)]$$

- **TOC7.** $[\mathsf{Occurs}(\overset{\bullet}{E}v_1,Tv_1) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_1,\downarrow S^{cp} v)] \Rightarrow$

$$\forall \overset{\bullet}{E}v_2,Tv_2[[Tv_1 \prec Tv_2 \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2,Tv_2)] \Rightarrow \neg\mathsf{Cat}(\overset{\bullet}{E}v_2,\downarrow S^{cp} v)]$$

For TEMP states, the pattern of occurrence of onsets and cessations exhibit a more complex temporal structure: any two onsets (cessations) are separated by at least on cessation (onset). Note that this may only apply to TEMP states since they are the only sort of state that can start, or cease, more than once.

- **TOC8.** $[\mathsf{Occurs}(\overset{\bullet}{E}v_1,Tv_1) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_1,\uparrow S^t v)$

$$\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_2, \uparrow S^t v) \wedge Tv_1 \prec Tv_2] \Rightarrow$$
$$\exists \overset{\bullet}{E}v_3, Tv_3[\mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \downarrow S^t v)$$
$$\wedge Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2]$$

- **TOC9.** $[\mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_1, \downarrow S^t v)$
$$\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_2, \downarrow S^t v) \wedge Tv_1 \prec Tv_2] \Rightarrow$$
$$\exists \overset{\bullet}{E}v_3, Tv_3[\mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow S^t v)$$
$$\wedge Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2]$$

These two theorems point to a particularly interesting meta-theoretical structure associated with each TEMP state.[7]

- $\mathcal{OC}$:TEMP$\longrightarrow 2^{\overset{\bullet}{\mathcal{E}} \times \overset{\bullet}{\mathcal{E}}}$, where
$$\mathcal{OC}(s^t) = \{\langle \overset{\bullet}{e}_1, \overset{\bullet}{e}_2 \rangle : \overset{\bullet}{e}_1 \in I(\uparrow s^t) \wedge \overset{\bullet}{e}_2 \in I(\downarrow s^t) \wedge \tau(\overset{\bullet}{e}_1) \prec \tau(\overset{\bullet}{e}_2)$$
$$\wedge \neg \exists \overset{\bullet}{e}_3 [\overset{\bullet}{e}_3 \in I(\downarrow s^t) \wedge \tau(\overset{\bullet}{e}_1) \prec \tau(\overset{\bullet}{e}_3) \wedge \tau(\overset{\bullet}{e}_3) \prec \tau(\overset{\bullet}{e}_2)]\}$$

The set $\mathcal{OC}(s^t)$ is said to be the set of *onset-cessation* pairs of s^t. Two events form an onset-cessation pair of some TEMP state if the state does not cease at any time between them. The following obvious observation sheds more light on the intuitive significance of onset-cessation pairs.

Observation 4.1 *For every $s^t \in$ TEMP, if $\langle \overset{\bullet}{e}_1, \overset{\bullet}{e}_2 \rangle \in \mathcal{OC}(s^t)$, then there is no $\overset{\bullet}{e}_3 \in I(\uparrow s^t)$ such that $\tau(\overset{\bullet}{e}_1) \prec \tau(\overset{\bullet}{e}_3)$ and $\tau(\overset{\bullet}{e}_3) \prec \tau(\overset{\bullet}{e}_2)$.*

Proof. Let $\langle \overset{\bullet}{e}_1, \overset{\bullet}{e}_2 \rangle \in \mathcal{OC}(s^t)$, for some $s^t \in$ TEMP. Assume that there is some $\overset{\bullet}{e}_3 \in I(\uparrow s^t)$ such that $\tau(\overset{\bullet}{e}_1) \prec \tau(\overset{\bullet}{e}_3)$ and $\tau(\overset{\bullet}{e}_3) \prec \tau(\overset{\bullet}{e}_2)$. Since $\tau(\overset{\bullet}{e}_1) \prec \tau(\overset{\bullet}{e}_3)$, then, by **TOC8**, there is some $\overset{\bullet}{e}_4 \in I(\downarrow s^t)$ such that $\tau(\overset{\bullet}{e}_1) \prec \tau(\overset{\bullet}{e}_4)$ and $\tau(\overset{\bullet}{e}_4) \prec \tau(\overset{\bullet}{e}_3)$. But since $\tau(\overset{\bullet}{e}_3) \prec \tau(\overset{\bullet}{e}_2)$, then by the transitivity of $\prec$ (**AT2**), $\tau(\overset{\bullet}{e}_4) \prec \tau(\overset{\bullet}{e}_2)$. Therefore $\tau(\overset{\bullet}{e}_1) \prec \tau(\overset{\bullet}{e}_4)$ and $\tau(\overset{\bullet}{e}_4) \prec \tau(\overset{\bullet}{e}_2)$, which, by the definition of $\mathcal{OC}$, implies that $\langle \overset{\bullet}{e}_1, \overset{\bullet}{e}_2 \rangle \notin \mathcal{OC}(s^t)$. Since this is contradictory, then no such $\overset{\bullet}{e}_3$ exists. $\square$

Thus, an onset-cessation pair for some state, s^t, bounds a *maximal* interval over which the state persists. Do such pairs correspond to anything in our ontology? As shall be argued below, each pair picks out a particular durative event. But, first, a digression.

[7] Here I'm allowing myself to be a little sloppy, using "$\wedge$" as meta-logical conjunction.

4.3 Simultaneous Events and the "Once-Only" Effect

The reader should notice that theorems **TOC6** and **TOC7** look exactly the same except for the different event categories employed. If we replace $\uparrow S^p v$ in **TOC6**, or $\downarrow S^{cp} v$ in **TOC7**, by some generic category, *ec*, we get the following schema.

$$(4.3)\quad [\mathsf{Occurs}(Ev_1, Tv_1) \land \mathsf{Cat}(Ev_1, ec)] \Rightarrow \\ \forall Ev_2, Tv_2[[Tv_1 \prec Tv_2 \land \mathsf{Occurs}(Ev_2, Tv_2)] \rightarrow \neg\mathsf{Cat}(Ev_2, ec)]$$

This schema may be interpreted as defining a particular property that may be exhibited by some event categories, namely that the occurrence of an instance of the category cannot be followed by another. Intuitively, this means that there can only be *one* occurrence of categories satisfying (4.3) (which makes perfect sense in the case of $\uparrow s$ and $\downarrow s$, for a PERM or a CO-PERM s). (Galton, 1984) calls such event categories *once-only*, and, in his system, he proves that the category of onsets of PERM states is once-only. Nevertheless, a careful look at (4.3) (or **TOC6** or **TOC7**) reveals that one cannot draw the conclusion that there is only one instance for categories satisfying it. The reason is that nothing in the system of axioms presented so far dismisses the possibility of having distinct events of the same category occurring *simultaneously*.[8] For example, there is no inconsistency in having all of the following in Cassie's belief space in addition to **TOC6**:

- $\mathsf{Occurs}(\dot{E}c_1, Tc)$.
- $\mathsf{Occurs}(\dot{E}c_2, Tc)$.
- $\mathsf{Cat}(\dot{E}c_1, \uparrow S^p c)$.
- $\mathsf{Cat}(\dot{E}c_2, \uparrow S^p c)$.

Thus, **TOC6** allows the existence of more than one onset for a PERM state, provided that they all occur simultaneously. How does Galton get around this difficulty? Actually, (Galton, 1984) is silent about this issue, and his axiomatic system admits multiple simultaneous occurrences of once-only events. The problem does not explicitly manifest itself since Galton's ontology does not include event tokens. In (Galton, 1987a), however, the model-theoretic semantics reveals how this is dealt

[8]Note that this also poses problems for the intuitive intepretation of the definition of onset-cessation pairs of TEMP states.

with. (Galton, 1987a) identifies the denotation of an event category with its set of tokens. Since his ontology does not include event tokens as primitives, he identifies these (roughly) with the intervals over which they occur. Thus, the tacit assumption is that there can only be one instance of any event category at any given time. The problem with $\mathcal{FOCS}$, however, is that event tokens are first-class entities in the ontology, independent of their time of occurrence. Nevertheless, one may achieve the same effect as Galton's by positing the following axiom:

(4.4) $[\mathsf{Cat}(Ev_1, ECv) \wedge \mathsf{Cat}(Ev_2, ECv) \wedge \mathsf{Occurs}(Ev_1, Tv) \wedge \mathsf{Occurs}(Ev_2, Tv)] \Rightarrow \mathsf{Equiv}(Ev_1, Ev_2)$

Given such an axiom, we can get the once-only effect for categories of onsets of PERM states and cessations of CO-PERM states. Unfortunately, however, (4.4) is not always valid as it stands. Consider the (de dicto reading of the) following English sentence and a possible representation of it in $\mathcal{FOCS}$:

(4.5) A man is crossing the street.

(4.6) $\exists x[\mathsf{Man}(x) \wedge \mathsf{Cross}(x, \mathrm{STREET})]$[9]

Let s refer to the temporary state represented by (4.6). According to (4.4), there can be only one onset of this state at any given time. But suppose that both John and Bill start crossing the street exactly at the same time (for example, when the walk-light turns on). Since both of these events (John's and Bill's starting to cross the street) are in the category $\uparrow s$, axiom (4.4) would render them one and the same, which, on any account, is just wrong.

For a PERM state example, consider (4.7), where both Bill and John turn 21 at exactly the same time (if this seems unlikely, think of the legal sense of turning 21.)

(4.7) A man has turned 21.

The problem is that, for two events e_1 and e_2, (4.4) only requires that $\mathcal{C}(e_1) \cap \mathcal{C}(e_2)$ be non-empty, which is too weak for co-referentiality.

What we need is to restrict (4.4) so that it only applies if $\mathcal{C}(e_1) = \mathcal{C}(e_2)$—if all categories of one event are categories of the other.

[9]"STREET" is a constant denoting a particular street.

$$\text{(4.8)}\quad [\mathsf{Occurs}(Ev_1, Tv) \wedge \mathsf{Occurs}(Ev_2, Tv) \\ \wedge \forall ECv[\mathsf{Cat}(Ev_1, ECv) \Leftrightarrow \mathsf{Cat}(Ev_2, ECv)]] \Rightarrow \mathsf{Equiv}(Ev_1, Ev_2)$$

Although (4.8) is reasonable enough for the purpose of the current discussion, whether it is always valid is something that has to be investigated. But even assuming that it is, an interesting question is whether its antecedent could ever be satisfied. That is, is it possible for Cassie to have two distinct mental representations of events and yet have no way of categorizing them differently? To appreciate where the difficulty lies, consider event categories such as "*events that John told me about*" or "*events that I became aware of at time t*". The very fact that there are *two* mental representations means that Cassie became aware of each under different circumstances, and, considering categories like the above two, there would have to be at least one category to which one event belongs and the other does not. As the reader should notice, event categories that seem to raise such problems are those that introduce some sort of an opaque context into the scene. Accordingly, (4.8) has to be further restricted so that it takes into account such opaque categories. However, for the purpose of the current discussion, I will settle for (4.8) as is.

Putting the above difficulties aside, (4.8) provides a sufficient condition for the co-referentiality of simultaneous events. Nevertheless, since simultaneity is by itself not sufficient, then we cannot get the once-only effect for event categories satisfying (4.3). The strongest we can get is that there is a single *time* over which possibly multiple occurrences of these categories take place.[10]

The above notwithstanding, it is still possible that, although not all categories satisfying (4.3) are strictly once-only in the intuitive sense, a subset of them are. In particular, consider categories, *ec*, satisfying the following minimality condition:

(4.9) For all $ec' \in \mathcal{EC}$ if $I(ec) \cap I(ec')$ is not empty, then $I(ec) \cap I(ec') = I(ec)$

If we envision event categories as forming a lattice ordered by the subset relation over sets of their instances, then categories satisfying (4.9) are the minimal elements of the lattice. Intuitively, these are categories that characterize their instances in the most specific way possible. Thus, the characterization of strict-once-only event categories would have to involve both (4.3) and the following $\mathcal{FOCS}$ schema.

[10]Note that (Galton, 1984; Galton, 1987a) face the same difficulty. In fact, Galton's system cannot distinguish these multiple occurences.

(4.10) $\mathsf{Cat}(Ev, ec) \Rightarrow \forall ECv' [\mathsf{Cat}(Ev, ECv') \Rightarrow \forall Ev' [\mathsf{Cat}(Ev', ec) \Rightarrow \mathsf{Cat}(Ev', ECv')]]$

Note that the above implies that $\mathcal{C}(e) = \mathcal{C}(e')$ for all pairs of instances of ec. Therefore, such pairs, if simultaneous (which is dictated by (4.3)), would have to be co-referential as per (4.8), and hence the strict once-only effect.

4.4 Temporal Closures

A durative event corresponds to some temporary state starting to hold, holding for a while, and then ceasing to hold. I will call this complex a *temporal closure* for the state. The terminology, of course, alludes to the mereotopological intuition underlying the notion of a durative event—a stretch of time over which a state persists, together with the transitions that demarcate it. Before formalizing these intuitions in $\mathcal{FOCS}$, a discussion of the meta-theoretical structure of durative events will be enlightening. First, let me define the following mapping.

- $\mathcal{CL} : \bigcup_{s^t} \mathcal{OC}(s^t) \longrightarrow \overset{\bullet\!-\!\bullet}{\mathcal{E}}$.

$\mathcal{CL}$ (for *closure*) maps an onset-cessation pair of some state s to a durative event of s maximally holding for a while.[11] What are reasonable properties of $\mathcal{CL}$? First, is it a function? Interestingly, it is not. The reason is that elements of $\overset{\bullet\!-\!\bullet}{\mathcal{E}}$ denote mental conceptualizations of events, rather than *actual* events. Thus, there may be distinct terms in $\overset{\bullet\!-\!\bullet}{\mathcal{E}}$ corresponding to multiple conceptualizations of the same actual event, which is captured by the function Equiv (see Section 3.3.4).[12] However, mereotopological properties are attributed to actual events, not conceptualizations. Therefore, $\mathcal{CL}$ maps an onset-cessation pair to multiple, but co-referential, elements of $\overset{\bullet\!-\!\bullet}{\mathcal{E}}$. More precisely:

- If $\langle \overset{\bullet}{e}_1, \overset{\bullet}{e}_2, \overset{\bullet\!-\!\bullet}{e_3} \rangle$ and $\langle \overset{\bullet}{e}_1, \overset{\bullet}{e}_2, \overset{\bullet\!-\!\bullet}{e_4} \rangle$ are in $\mathcal{CL}$, then $\beta \vdash \mathsf{Equiv}(\overset{\bullet\!-\!\bullet}{e_3}, \overset{\bullet\!-\!\bullet}{e_4})$

Whether $\mathcal{CL}$ is onto is not immediately obvious either. The whole idea behind $\mathcal{CL}$ is to provide a unified account of all durative events in terms of state transitions and, thus, ultimately in terms of

[11] Some authors (Herweg, 1991b, for instance) refer to intervals over which a state s maximally holds as *phases* of s.

[12] Even though multiple conceptualizations of the same actual event may be captured by multiple categorizations, there are reasons why we might still need distinct terms to represent them. For example, Cassie comes to believe that John attended a linguistics talk and, in a different occasion, she learns that Mary too attended a linguistics talk. Later, Cassie may realize that both attended the same talk, and that can be modeled by believing that the mental entities corresponding to the talk attended by John and that attended by Mary are co-referential (Maida and Shapiro, 1982).

states. A non-onto $\mathcal{CL}$ would mean that there are durative events that do not correspond to onset-cessation pairs. Ontologically, this seems to contradict my informal characterization of durative events. Nevertheless, epistemologically, Cassie may conceive of a durative event without explicitly conceiving of its boundaries. For example, a reasonable $\mathcal{FOCS}$ representation of (3.2) (*I crossed the street*) need not mention the boundaries demarcating the particular event being reported. Similarly, Cassie may conceive of an onset-cessation pair without conceiving of the durative event it defines (and, thus, $\mathcal{CL}$ would only be a partial function). To allow these possibilities, $\mathcal{CL}$ should not be required to be onto or total.

To formalize the above intuitions in $\mathcal{FOCS}$, some definitions will be introduced to render the notation more convenient. Note that the following are not official $\mathcal{FOCS}$ expressions, but macros to be expanded wherever they are used (expansion might also involve variable renaming).

- $\mathsf{NoOcc}(ec, t_1, t_2) =_{\text{def}} \neg\exists Ev, Tv[\mathsf{Cat}(Ev, ec) \land t_1 \prec Tv \land Tv \prec t_2 \land \mathsf{Occurs}(Ev, Tv)]$

- $\mathsf{OCPair}(\overset{\bullet}{e}_1, \overset{\bullet}{e}_2, s^t) =_{\text{def}}$
$$\begin{aligned}&[\mathsf{Cat}(\overset{\bullet}{e}_1, \uparrow s^t) \land \mathsf{Cat}(\overset{\bullet}{e}_2, \downarrow s^t)\\ &\land \forall Tv_1, Tv_2[[\mathsf{Occurs}(\overset{\bullet}{e}_1, Tv_1) \land \mathsf{Occurs}(\overset{\bullet}{e}_2, Tv_2)] \Rightarrow\\ &\qquad\qquad [Tv_1 \prec Tv_2 \land \mathsf{NoOcc}(\downarrow s^t, Tv_1, Tv_2)]]\end{aligned}$$

- $\mathsf{Covers}(t_1, t_2, t_3) =_{\text{def}}$
$$\begin{aligned}&t_2 \sqsubseteq t_1 \land t_3 \sqsubseteq t_1\\ &\land \forall Tv[Tv \supset\subset t_1 \Leftrightarrow Tv \supset\subset t_2]\\ &\land \forall Tv'[t_1 \supset\subset Tv' \Leftrightarrow t_3 \supset\subset Tv']\end{aligned}$$

The first definition specifies conditions for there to be no occurrences of a given event category between two given times. The second identifies two punctual events as forming an onset-cessation pair for a given state (note that the definition does not require the two events to actually occur). The third defines what it means for an interval to cover two other intervals. In the language of (Allen, 1983), t_1 covers t_2 and t_3 if it is *started* by t_2 and *finished* by t_3 (a similar definition appears in (Koomen, 1991)).

Corresponding to the meta-logical $\mathcal{CL}$, the following $\mathcal{FOCS}$ function forms $\overset{\bullet\!-\!\bullet}{\mathcal{E}}$ terms out of $\overset{\bullet}{\mathcal{E}}$ pairs.

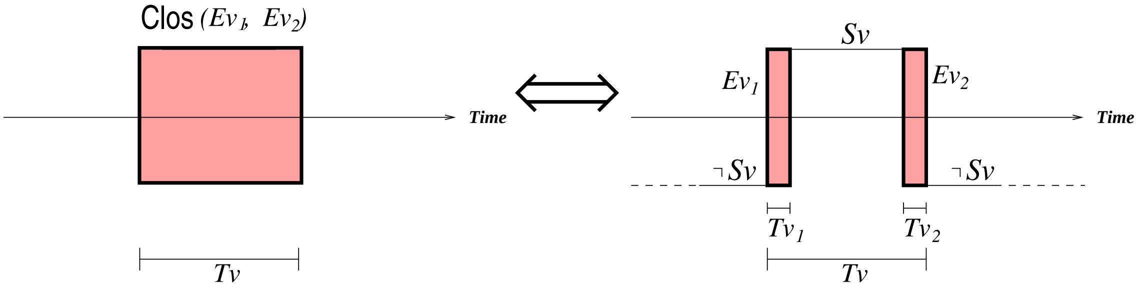

Figure 4.1: The relation between durative events and onset-cessation pairs.

- $\mathsf{Clos}\colon \dot{\mathcal{E}} \times \dot{\mathcal{E}} \longrightarrow \overset{\bullet-\bullet}{\mathcal{E}}$, where $[\![\mathsf{Clos}(\dot{e}_1, \dot{e}_2)]\!]$ is the durative event temporally-bounded by $[\![\dot{e}_1]\!]$ and $[\![\dot{e}_2]\!]$.

Since onset-cessation pairs do not correspond to a $\mathcal{FOCS}$ sort, we need to axiomatize Clos so that it only applies to onset-cessation pairs. In addition, we need to explicate the temporal-bounding relation that holds between $\mathsf{Clos}(\dot{e}_1, \dot{e}_2)$ on one hand and e_1 and e_2 on the other. The structure embodied in the following axiom is depicted in Figure 4.1.

- **AE4.** $\mathsf{Occurs}(\mathsf{Clos}(\dot{Ev}_1, \dot{Ev}_2), Tv) \Leftrightarrow$

$$\exists S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\dot{Ev}_1, \dot{Ev}_2, S^t v) \wedge \mathsf{Occurs}(\dot{Ev}_1, Tv_1) \wedge \mathsf{Occurs}(\dot{Ev}_2, Tv_2) \wedge \mathsf{Covers}(Tv, Tv_1, Tv_2)]$$

All that the above axiom mandates is that *some* durative events correspond to the temporal closures of TEMP states. But we need a stronger statement, requiring this correspondence to cover *all* durative events.

- **AE5.** $\forall \overset{\bullet-\bullet}{Ev}\ [\exists\, \dot{Ev}_1, \dot{Ev}_2\ [\mathsf{Equiv}(\overset{\bullet-\bullet}{Ev}, \mathsf{Clos}(\dot{Ev}_1, \dot{Ev}_2))]]$

Given **AE2**, **AE4**, and **AE5** we can prove that occurrences of durative events, in general, are associated with onset-cessation pairs satisfying the temporal pattern depicted in the right-hand part of Figure 4.1.

- **TE3.** $\exists \overset{\bullet-\bullet}{Ev}\ [\mathsf{Occurs}(\overset{\bullet-\bullet}{Ev}, Tv)] \Leftrightarrow$

$$\exists\, \dot{Ev}_1, \dot{Ev}_2, S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\dot{Ev}_1, \dot{Ev}_2, S^t v) \wedge \mathsf{Occurs}(\dot{Ev}_1, Tv_1) \wedge \mathsf{Occurs}(\dot{Ev}_2, Tv_2)$$

$$\wedge \mathsf{Covers}(Tv, Tv_1, Tv_2)]$$

Note what we have done so far. We have characterized durative event tokens in terms of state transitions. More precisely, a durative event comes out as a complex closed situation made up of a boundary, an onset-cessation pair, and an interior homogeneously *filled* with a state. Thus, communication aside, the occurrence of a durative event can only be *inferred*, using **TE3**. The only thing that Cassie needs in order to infer occurrences of events, in general, is to monitor various states. If Cassie determines that $\neg s$ holds and then determines that s holds, she may infer the occurrence of an onset of s (**AOC5**). If later she determines that $\neg s$ holds again, she may infer that a cessation of s has occurred (**TOC5**). Using **TE3**, she can then infer the occurrence of a durative event. Thus, events are purely conceptual entities. They are so in the sense that an agent, Cassie for example, can never experience an event, only conceives of it as it experiences states holding (see the discussion in Section 3.2). But conceiving of an event essentially involves categorizing it, and it should be intuitive enough that there must be some relation between the event cetegories associated with some durative event and a state that fills its interior. Hence, to this point we should now turn.

4.5 Packaging and Grinding

Given **TE3**, what do we know about a given durative event? If nothing else, we know that it is an event of some state holding for a while. For example, consider the following English sentences.

(4.11) John slept for two hours.

(4.12) For a while, John was jogging.

(4.13) John was in New York from 1970 to 1975.

All of these sentences refer to some state holding for a specified or unspecified, but bounded, period of time. On the account presented here, each sentence reports the occurrence of an event (also see (Depraetere, 1995)): an event of John's sleeping, his jogging, and his being in New York, respectively. The descriptions that I just gave to the three events are possible because of knowledge of the states that fill their interiors. Following (Galton, 1984) (and (Herweg, 1991b), who actually follows Galton), I will introduce a function, PO,[13] that acts like a universal packager for states (see Section 3.2).

- PO:$\textsc{Temp} \longrightarrow \overset{\bullet-\bullet}{\mathcal{EC}}$, where $[\![\mathsf{PO}(s^t)]\!]$ is the event category of the state $[\![s^t]\!]$ maximally holding for a while.

The following axiom captures the intended semantics of PO.

- **AE6.** $\mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, S^t v) \Leftrightarrow \mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \mathsf{PO}(S^t v))$

Following many authors (Dowty, 1977; Galton, 1984; Kamp and Reyle, 1993, for instance), I introduce a function, Prog to act as a universal grinder for event categories.

- Prog: $\overset{\bullet-\bullet}{\mathcal{EC}} \longrightarrow \textsc{Temp}$, where $[\![\mathsf{Prog}(\overset{\bullet-\bullet}{ec})]\!]$ is the state that holds whenever an event of category $[\![\overset{\bullet-\bullet}{ec}]\!]$ is in progress.

Axiom **AE7** explicates the semantics of Prog.

- **AE7.** $$[\mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, S^t v) \wedge \mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{EC}v)] \Rightarrow$$
$$\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$$
$$\wedge Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow$$
$$\mathsf{Holds}(\mathsf{Prog}(\overset{\bullet-\bullet}{EC}v), Tv_3)]$$

There are three things to note.

1. It might seem that a simpler version of the above axiom would be sufficient. In particular, consider the following:

   (4.14) $[\mathsf{Occurs}(\overset{\bullet-\bullet}{E}v, Tv) \wedge \mathsf{Cat}(\overset{\bullet-\bullet}{E}v, \overset{\bullet-\bullet}{EC}v)] \Rightarrow \forall Tv'[Tv' \sqsubseteq TV \Rightarrow \mathsf{Holds}(\mathsf{Prog}(\overset{\bullet-\bullet}{EC}v), Tv')]$

   Although this seems to capture the *intuition*, it is actually incorrect. The problem is that (4.14) requires the progressive state to be holding over *all* subintervals of the interval over which the event occurs. But remember that the interval over which a durative event occurs is made up of three parts: two intervals associated with an onset-cessation pair and the interval, delimited by these two, over which a state persists. It is not clear whether an event is in progress at the

[13](Galton, 1984) makes this particular choice of name for the operator to allude to the Russian prefix *po-* which has the same aspectual effect as the above-introduced PO.

onset-cessation intervals. For this reason, we need to allude to the internal structure of events, and to only require the progressive state to hold over their interior.

2. **AE7** only requires the progressive state to hold over the *interior* of an event; it is silent, not only about the boundaries as discussed above, but also about the *exterior* of the event. In particular, the axiom does not require the onset-cessation pair mentioned therein to coincide with (or be identical to) an onset-cessation pair for the progressive state. The reason is that the domain of Prog is the set of event *categories*, not that of events. Two distinct events, $\overset{\bullet-\bullet}{e_1}$ and $\overset{\bullet-\bullet}{e_2}$, belonging to the same category, $\overset{\bullet-\bullet}{ec}$, may overlap in time giving rise to a continuous stretch of time (the one Cover-ing the intervals over which the two events occur) over which $\mathsf{Prog}(\overset{\bullet-\bullet}{ec})$ holds and that is not within the interior of either event. For example, suppose that *Lecture* is a durative event category (with the obvious semantics). Further, suppose that there are two particular lectures, L_1 and L_2, such that L_1 starts at 9 a.m. and ends at 11 a.m., and that L_2 starts at 10 a.m. and ends at 12 p.m. The state $\mathsf{Prog}(\mathit{Lecture})$ holds over the entire period starting at 9 a.m. and ending at 12 p.m. even though neither lecture extends over this period.

3. As it stands, the axiom represents a *sufficient* condition for progressive states to hold. It does not exhaust all the cases in which such states may hold. In particular, note that the axiom states that whenever an event of category $\overset{\bullet-\bullet}{ec}$ *occurs*, the state $\mathsf{Prog}(\overset{\bullet-\bullet}{ec})$ holds over its interior. However, a well-studied feature of the English progressive (and the imperfective aspect in general) is that sentences in the progressive often do not entail their non-progressive counter-parts (Dowty, 1977; Galton, 1984; Parsons, 1989; Landman, 1992, for example). In particular, the entailment fails when the event category involved is *telic*, or an *accomplishment* in the terminology of (Vendler, 1957). For example, (3.1) does not entail (3.2) (repeated here for convenience)

   (3.1) I am crossing the street.

   (3.2) I crossed the street.

   If we represent the state reported in (3.1) by $\mathsf{Prog}(\mathsf{Cross}(\mathrm{I},\ \mathrm{STREET}))$ then even though Cassie may believe that such a state holds, it is possible that no event of category $\mathsf{Cross}(\mathrm{I},\ \mathrm{STREET})$

ever occurs (i.e., Cassie may never manage to successfully cross the street). Most scholars adopt a possible worlds/modal approach in order to give necessary conditions for the obtaining of progressive states (Dowty, 1977; Galton, 1984; Landman, 1992, for instance). My approach, which will be presented in Section 4.6, is closer to the spirit of (Parsons, 1989), relying on the intensional semantics of SNePS (and, hence, $\mathcal{FOCS}$).

Composing Prog and PO gives rise to what may *roughly* be called a *fixed-point* property.

- **TE4.** $\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, S^t v) \Rightarrow$

$$\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\dot{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\dot{E}v_2, Tv_2) \\ \wedge Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow \\ \mathsf{Holds}(\mathsf{Prog}(\mathsf{PO}(S^t v)), Tv_3)]$$

Again, there are three points to note.

1. We cannot prove (or even require) the following seemingly simpler version of **TE4**:

   (4.15) $\mathsf{Holds}(S^t v, Tv) \Rightarrow \mathsf{Holds}(\mathsf{Prog}(\mathsf{PO}(S^t v)), Tv)$

   The reason is that a state s^t, though temporary, may hold indefinitely (see Section 3.7) and, thus, does not give rise to the onset-cessation pair required by PO.

2. Interestingly, composing PO and Prog does not give rise to any fixed-point properties like **TE4** for the composition of Prog and PO. First, the following is not correct (in general) as pointed out above; although Cassie can spend a while *crossing* the street, she may never actually *cross* the street.

   (4.16) $\mathsf{Cat}(Ev, \mathsf{PO}(\mathsf{Prog}(ECv))) \Rightarrow \mathsf{Cat}(Ev, ECv)$

   Second, the other direction of (4.16) is not correct either.

   (4.17) $\mathsf{Cat}(Ev, ECv) \Rightarrow \mathsf{Cat}(Ev, \mathsf{PO}(\mathsf{Prog}(ECv)))$

   The reason, again, is the possibility of overlapping events of the same category. To use the same example cited above, although L_1 is of category *Lecture*, it is not of category $\mathsf{PO}(\mathsf{Prog}(\mathit{Lecture}))$ since it does not occur over an interval over which $\mathsf{Prog}(\mathit{Lecture})$ *maximally* holds.

4.6 Telicity

One dimension along which linguists characterize events is that of telicity (Garey, 1957; Comrie, 1976; Declerck, 1979; Dahl, 1981; Krifka, 1998, for instance). An event is telic if it is described (or conceived of) as having a built-in *definite* ending state, otherwise it is *atelic*. For example, the following pair of sentences describe certain situations (possibly the same) as telic and atelic, respectively.

(4.18) John ran to the store.

(4.19) John ran.

The notion of telicity is an important one in linguistics and is certainly worth a precise characterization. As it turns out, however, telicity is also crucial for the specification of correct execution of sequences of acts (see Chapter 9). For this reason, the analysis that follows will be making distinctions that linguistics is not traditionally interested in but that are important for action execution. On the other hand, there will be no discussion of how telic features of an event category are derived from its compositional makeup—an issue that, for obvious reasons, linguists are concerned with (Declerck, 1979; Verkuyl, 1989; Jackendoff, 1996; Krifka, 1998, for instance)

4.6.1 Completion

Central to the notion of telicity is that of completion. Let me be precise about what I mean by that. Suppose that we instruct Cassie to run to the store and that she actually starts running. When would Cassie's act be considered *complete*? Intuitively, when she is at the store; she may then stop running and would have correctly performed an act of running to the store. Had the instruction been to just run, Cassie's act would be complete whenever she stops running, with no restrictions whatsoever. The only difference between the two instructions is one of telicity: "run to the store" is telic and "run" is atelic. The same applies to general events not just acts; the telic features of an event category determine *completion conditions* for any of its instances.

But now, someone might object, if being *complete* is a possible property of an event, wouldn't this mean that there could be *incomplete* events? It does, and that should not be problematic. Given the intensional semantics of SNePS, incomplete entities do not pose any ontological problems. In fact, representing such entities is one of the reasons behind the choice of intensional semantics (Shapiro and Rapaport, 1987). Recall that $\mathcal{FOCS}$ terms denote objects of *thought*, not entities out there in the world. Thus, whether there are incomplete events *out there* is irrelevant. What is relevant is whether incomplete events are possible objects of thought. (Parsons, 1989) faces a similar situation since his analysis of the progressive commits him to the existence of incomplete (or "unfinished") objects.[14] In defence of incomplete objects, Parsons argues:

> In Northern California there is a state park – Jack London State Park. One can go there and see the house that Jack London was building when he died. At least this is what the tourists say. It isn't much of a house – only a foundation and parts of small walls. But native speakers of English call it a house. What evidence could there be that they are wrong? (Parsons, 1989, p. 225)

What I would like to take out of this is that incomplete objects are, at least, objects of thought. The same argument could be made for events. To take the above example, suppose that Cassie is instructed to run to the store, starts runnnig, but is interrupted by some unfortunate mishap (a dead battery, a broken wheel, etc.). It is arguable that before heading to the store, Cassie conceives of the act that she is about to perform as one of running to (and reaching) the store. Neverthless, the act is never completed as such; an event of Cassie's running to the store never occurs.[15] What occurs is another event, perhaps one of her running, or running toward the store.

Although there could be incomplete events, once complete, an event is always complete; you cannot go back in time and *un*-complete it.[16] Thus, the completion of a *particular* event is permanent. In fact, it is the prototypical example of logical permanence (see Section 3.23).

- $\mathsf{Complete}: \mathcal{E} \longrightarrow \text{PERM}$, where $[\![\mathsf{Complete}(e)]\!]$ is the state of event $[\![e]\!]$'s being complete.

Completion of an event is closely tied to its occurrence; if an event is complete, then it must have already occurred.

- **AE8.** $\mathsf{Holds}(\mathsf{Complete}(Ev), Tv) \Rightarrow \exists Tv'[Tv' \prec Tv \wedge \mathsf{Occurs}(Ev, Tv')]$

[14] And in the case of (Parsons, 1989), these objects exist out there in the world.
[15] In this respect, also see (Hirst, 1991).
[16] But you can *change your mind* about whether it is complete.

It is tempting to think that axiom **AE8** should be a bi-conditional. That is, to assert that if an event has occurred, then it must be complete. Admittedly, this sounds intuitive and reasonable enough that one has to justify why it is not the case. Consider sentence (4.18). In order to run to the store, John starts some process (namely, running) and stops when, and only when, he is at the store. In this case, the completion of the act is simultaneous with John's stopping to run. There are cases, however, where the simultaneity of the cessation of the process and the onset of the act completion fails to hold. Consider the following examples.

(4.20) John pushed the rock down the hill into the river.

(4.21) John slid the pen across the table to Mary.

Note that, in the case of (4.20), whatever John is *doing* ceases once he pushes the rock. Nevertheless, a period of time elapses before the rock is actually in the river. During this period, John is not, and typically cannot, do anything to help achieve that state. Note that it is not appropriate for John to utter (4.22) after having pushed the rock.

(4.22) I am pushing the rock down the hill into the river.

More importantly, John cannot truthfully utter (4.23).

(4.23) I have pushed the rock down the hill into the river.

In examples like (4.18), the achievement of the state signalling the completion of the act (being at the store) is totally dependent on the agent's (John's) behavior, and the cessation of the process of running takes place when and only when the act is complete. In examples like (4.20) and (4.21), on the other hand, achieving that state is only partially dependent on the agent's behavior. The agent merely initiates a sequence of events that (may) result in achieving the intended goal and hence completing the act. These two cases correspond to sentences that (Talmy, 2000) describes as involving *extended causation* and *onset causation*, respectively. Telic events may therefore be categorized into two types: *telic with extended causation* (denoted $\overset{\rightarrow\bullet}{\mathsf{telic}}$), where the cessation of the process is almost simultaneous with the completion of the event; and *telic with onset causation* (denoted $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$), where a temporal gap separates the cessation of the process from the completion of the act.[17] All events that are *not* $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ satisfy the converse of **AE8**.

(4.24) $[\mathsf{Occurs}(Ev,Tv) \wedge Tv \prec Tv'] \Rightarrow \mathsf{Holds}(\mathsf{Complete}(Ev),Tv')$

I will later come back to the issue of $\overset{\rightarrow\bullet}{\mathsf{telic}}$ and $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ events. But now that the notion of completion has been introduced, a more tight characterization of what it means and how exactly it is related to telicity is due.

4.6.2 The Unfolding of Events

We often conceive of events before they occur. Every year, as I plan to fly to Cairo, I keep on thinking of the details of the trip: how I am going to get to the airport, when I should be there, how I will spend my time in transit, which books I should take with me for entertainment, and so on. I almost have a complete conceptualization of the event, awaiting its occurrence to fill in the fine details. As the time comes, I find myself getting *into* the event, watching it as it gradually unfolds in time. In the meantime, I witness various states: driving, flying, sleeping, standing in the check-in line, etc. These states hold only during portions of the trip, not throughout its entire extent. What then is a state that characterizes the interior of the whole event? One possible candidate is the progressive state associated with the *travelling-to-Cairo* event category. Although this state holds throughout the period of the trip, it does not exactly capture what I have in mind during any *particular* trip. In particular, note that this state holds throughout *any* trip to Cairo, although each trip precipitates a different experience. I suggest that, just as there is a specific (progressive) state associated with each event category, there is a specific state associated with each *event*—the state of *that* event being in progress. Thus, for the particular trip I made to Cairo in December 2000, there is a unique state that I got into as I started my trip, a state that held throughout the whole event and that never held before, nor will it hold after, the event.

But is there a way to independently characterize such a state? The formalization to be developed below will *not* provide such a characterization; I will simply assume, by fiat, the existence of a unique state for every particular event. There is a reason for this, however. Unlike perceptually-grounded states such as the light's being on or the air's smelling of jasmine, the state of a particular

[17] A famous example of $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ acts is that of (Thomson, 1971), where a person is shot at t_1 but dies at some later time, t_2. The problem that Thomson addresses is whether the shooting event and the killing event are identical and how that may be, given that one of them is *complete* before the other. More on this below.

event in progress is primarily conceptual. Note, in particular, that such a state cannot be *recognized* as a result of raw sensations; one has to first conceive of an event, and then decide that it is in progress. Neverthless, there is one way to describe these states in more basic terms. In Section 4.1, I pointed out that two events are distinct in as much as there are categories to which one belongs and the other does not. One might formalize this by representing a particular event as some sort of a collection of all the event categories it belongs to (cf. the set $\mathcal{C}$ in Section 4.1). The problem, of course, is that conscious knowledge of an agent need not be fine-grained and extensive enough to render any two events distinguishable based solely on their categories. For this reason, we introduced event tokens as an independent sort of the logic.[18] Similarly, we can think of the state of a particuilar event in progress as the mereological fusion of the progressive states associated with *all* event categories to which the event belongs. Since not all of these categories are available at the conscious level, I posit the existence of the state as a primitive of the theory. If you do not find the above convincing, then I would have to say that the existence of such a thing as the state of a particular event being in progress is justified because it is *needed* for further development of the system.

Formally, there is a (one-to-one) mapping from durative events to their *interiors*.

- Int: $\overset{\bullet-\bullet}{\mathcal{E}} \xrightarrow{into} \textsc{Temp}$, where $[\![\mathsf{Int}(\overset{\bullet-\bullet}{e})]\!]$ is the state of event $[\![\overset{\bullet-\bullet}{e}]\!]$ being in progress.

The following axiom establishes a basic relation between an event $\overset{\bullet-\bullet}{e}$ and the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$.

- **AE9.** $\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev})))$.

Given **AE9** (together with **AE3**, **AE4**, **AE5**, and **AE6**), we can prove that whenver a durative event, $\overset{\bullet-\bullet}{e}$, occurs, the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ fills its interior. Note that the opposite is not true; the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ may hold for a while without $\overset{\bullet-\bullet}{e}$ ever occurring. This is similar to the case with Prog states. For example, let e be the event of my travelling to Cairo in December 2000. Suppose that, on the day of the flight, I leave Buffalo and drive to Toronto to catch the plane. During this period, the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ holds. However, I reach Toronto's airport only to find that all flights have been canceled due to a severe blizzard. This results in a total change in my plans, forcing me to drive back to Buffalo through the snow. Thus, even though the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ held for a while, the event e never

[18]Note that it is also this need to represent only partial information that makes a situation in the situation calculus (McCarthy and Hayes, 1969) different from the set of fluents that hold in it.

occurred; if it had, I would have been in Cairo. What *did* occur is a *part* of $\overset{\bullet-\bullet}{e}$, an event of category $\mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{e}))$, one that, had it unfolded in a certain way, would have been $\overset{\bullet-\bullet}{e}$. Note that this is exactly similar to the unfinished house of (Parsons, 1989).

Another important property of states formed by Int is that, if they hold, they must hold for only a bounded period of time. This is a property inherited from the *closed-ness* of events.

- **AE10.** $\forall \overset{\bullet-\bullet}{Ev}\ [\exists\ \overset{\bullet}{E}v_1, Tv_1[\mathsf{Cat}(\overset{\bullet}{E}v_1, \uparrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1)] \Leftrightarrow$
  $\exists\ \overset{\bullet}{E}v_2, Tv_2[\mathsf{Cat}(\overset{\bullet}{E}v_2, \downarrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)]]$

In the informal discussion, it was noted that Int states only hold once. That is, $\mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{e}))$, for any $\overset{\bullet-\bullet}{e} \in \overset{\bullet-\bullet}{\mathcal{E}}$, is a once-only event category à la (Galton, 1984).[19]

- **AE11.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev_1}, \mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev}))) \wedge \mathsf{Cat}(\overset{\bullet-\bullet}{Ev_2}, \mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev})))$
  $\wedge \mathsf{Occurs}(\overset{\bullet-\bullet}{Ev_1}, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet-\bullet}{Ev_2}, Tv_2)] \Rightarrow$
  $\mathsf{Equiv}(\overset{\bullet-\bullet}{Ev_1}, \overset{\bullet-\bullet}{Ev_2})$

Over any interval, if the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ holds, then it must be the case that all the progressive states associated with event categories of $\overset{\bullet-\bullet}{e}$ hold as well.

- **AE12.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \mathsf{Holds}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev}), Tv)] \Rightarrow \mathsf{Holds}(\mathsf{Prog}(\overset{\bullet-\bullet}{ECv}), Tv)$

The following axiom expresses an important relation between an event $\overset{\bullet-\bullet}{e}$ and temporal closures of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$.

- **AE13.** $[\mathsf{Occurs}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), Tv) \wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev}))] \Rightarrow$
  $[\mathsf{Equiv}(\overset{\bullet-\bullet}{Ev}, \mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2)) \Leftrightarrow \exists Tv'[\mathsf{Holds}(\mathsf{Complete}(\overset{\bullet-\bullet}{Ev}), Tv')]]$

Note what this saying. If you know that an event $\overset{\bullet-\bullet}{e}$ is complete, then the *unique* occurrence (see **AE11**) of the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ is itself *the* occurrence of $\overset{\bullet-\bullet}{e}$. In addition, if you know that the unique occurrence of the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ is the occurrence of $\overset{\bullet-\bullet}{e}$, then $\overset{\bullet-\bullet}{e}$ must be complete. The second part basically asserts that completion is transparent to equivalence. The first part is the crucial one. It defines the condition under which a durative event is to be considered complete, namely if it is identical to the unique closure of its progressive state. Now the crucial point is how this identity is to be established. Consider the following *potential* axiom.

[19]Note that this is a strong statement about the "once-only" effect. In particular, not only are we requiring events of category $\mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{e}))$ to be simultaneous, but actually identical (see Section 4.3).

$$(4.25)\,[\mathsf{OCPair}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, \mathsf{Int}(\overset{\bullet\!-\!\bullet}{Ev}))$$
$$\wedge\ \forall \overset{\bullet\!-\!\bullet}{ECv}\,[\mathsf{Cat}(\overset{\bullet\!-\!\bullet}{Ev}, \overset{\bullet\!-\!\bullet}{ECv}) \Rightarrow \mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2), \overset{\bullet\!-\!\bullet}{ECv})]] \Rightarrow \mathsf{Equiv}(\overset{\bullet\!-\!\bullet}{Ev}, \mathsf{Clos}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2))$$

This says that an event $\overset{\bullet\!-\!\bullet}{e}$ is identical to the closure of $\mathsf{Int}(\overset{\bullet\!-\!\bullet}{e})$ if that closure falls under all categories that $\overset{\bullet\!-\!\bullet}{e}$ falls under. To illustrate, consider the following example.[20] John decides to play the piano and, thereby, to wake up Brown. Formally, let $\overset{\bullet\!-\!\bullet}{e_1}$ be the particular act that John sets out to perform. This act is restricted to belong to two categories: $\overset{\bullet\!-\!\bullet}{ec_1}$, for playing the piano, and $\overset{\bullet\!-\!\bullet}{ec_2}$, for waking up Brown.[21] Once John starts the performance of $\overset{\bullet\!-\!\bullet}{e_1}$ (by starting to play the piano), the state $\mathsf{Int}(\overset{\bullet\!-\!\bullet}{e_1})$ starts to hold (at least as far as John in concerned). Eventually, the state $\mathsf{Int}(\overset{\bullet\!-\!\bullet}{e_1})$ ceases (probably by John's ceasing to play the piano) giving rise to a temporal closure, $\overset{\bullet\!-\!\bullet}{e_2}$, of category $\mathsf{PO}(\mathsf{Int}(\overset{\bullet\!-\!\bullet}{e_1}))$. How can we tell if $\overset{\bullet\!-\!\bullet}{e_2}$ is identical to $\overset{\bullet\!-\!\bullet}{e_1}$, the act that John has set out to perform? Basically, by checking if it satisfies all the restrictions imposed on $\overset{\bullet\!-\!\bullet}{e_1}$. For this example, it is unproblematic that $\overset{\bullet\!-\!\bullet}{e_2}$ is of category $\overset{\bullet\!-\!\bullet}{ec_1}$. Nevertheless, whether $\overset{\bullet\!-\!\bullet}{e_2}$ is of category $\overset{\bullet\!-\!\bullet}{ec_2}$ depends on whether John actually manages to wake up Brown. If he does, then $\overset{\bullet\!-\!\bullet}{e_2}$ is identical to $\overset{\bullet\!-\!\bullet}{e_1}$—an event of $\overset{\bullet\!-\!\bullet}{e_1}$ in progress that satisfies all the restrictions imposed on $\overset{\bullet\!-\!\bullet}{e_1}$. On the other hand, if Brown does not wake up, then $\overset{\bullet\!-\!\bullet}{e_2}$ is *not* the same as $\overset{\bullet\!-\!\bullet}{e_1}$ even though it is an event of $\overset{\bullet\!-\!\bullet}{e_1}$ in progress.

Although it captures these intuitions, (4.25) is, unfortunately, too strong. In particular, in order to use (4.25) to prove that an event $\overset{\bullet\!-\!\bullet}{e}$ is identical to the closure of $\mathsf{Int}(\overset{\bullet\!-\!\bullet}{e})$ we have to prove that every *possible* category of $\overset{\bullet\!-\!\bullet}{e}$ is *necessarily* a category of that closure. This, in general, is not feasible since most categories are assigned to events by fiat; there is no general characterization of all the possible categories of a given event. However, as mentioned above, this is too strong for what is actually required. The only categories that the closure of $\mathsf{Int}(\overset{\bullet\!-\!\bullet}{e})$ needs to satisfy are those that $\overset{\bullet\!-\!\bullet}{e}$ is restricted to belong to just before $\mathsf{Int}(\overset{\bullet\!-\!\bullet}{e})$ starts to hold. These are not all the possible categories of $\overset{\bullet\!-\!\bullet}{e}$, only the ones $\overset{\bullet\!-\!\bullet}{e}$ happens to be envisioned as an instance of. Thus, instead of the elegant general axiom in (4.25), for any given event, there will be a more *blunt* axiom specifying when that event is identical to the closure of its Int state. Such axioms will follow the following general schema.

[20] The example is a slight variant of one due to (Goldman, 1970, pp. 2–3) who introduces it in a discussion of event indviduation.

[21] Here I am adopting a Davidsonian view (Davidson, 1963) by assuming that there is a *single* act under multiple categories.

$$(4.26)\,[\mathsf{OCPair}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, \mathsf{Int}(\overset{\bullet-\bullet}{e}))$$
$$\wedge \bigwedge_{\overset{\bullet-\bullet}{ec} \in \mathcal{C}(\overset{\bullet-\bullet}{e})} \mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2), \overset{\bullet-\bullet}{ec})] \Rightarrow \mathsf{Equiv}(\overset{\bullet-\bullet}{e}, \mathsf{Clos}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2))]$$

Note that, since, at any time, $\mathcal{C}(\overset{\bullet-\bullet}{e})$ is the set of all event categories that, Cassie believes, $\overset{\bullet-\bullet}{e}$ is an instance of, (4.26) captures the intuition. For a given event, $\overset{\bullet-\bullet}{e}$, Cassie comes to believe an appropriate instance of (4.26) once she determines that the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ holds (see Chapter 5). The next step now is to introduce the axioms necessary to establish which categories the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ belongs to.

4.6.3 An Ontology of Event Categories

The difference between telic event categories (whether $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ or $\overset{\rightarrow\bullet}{\mathsf{telic}}$) and atelic ones hinges on the existence of certain temporal constraints on when an instance of the category is considered complete. For telic categories, there is a state that any instance cannot be considered complete before, nor could it extend after, it starts. For atelic categories, no state so constrains the completion of its instances. On a more abstract level, let *ec* be an event category, *e* an arbitrary member of *ec*, and *s* a TEMP or PERM state (i.e., one that has the +*onset* feature). The telic features of *ec* are determined by the kinds of constraints on the temporal position of the onset of $\mathsf{Complete}(e)$ with respect to the *first* onset of *s* following the start of *e*. For example, suppose that *ec* is the category of events of running to the store and that *s* is the state of being at the store. In this case, any instance *e* of *ec* starts to be complete as *s* starts to hold (for the first time, following the start of *e*). Note that this constraint does not hold for other choices of *s*—being at the park for example. Thus, an event category is telic if there is at least one state *s* whose onset signals the completion of any of its instants. On the other hand, if we take *ec* to be the category of running, then there is no particular state *s* that signals the completion of all of its instants. This is the sign of atelicity.

Formally, we define the following macro for telic event categories. An event category *ec* is telic if there is a state *s* such that $\mathsf{telic}(ec, s)$.[22]

- $\mathsf{telic}(\overset{\bullet-\bullet}{ec}, s) =_{\text{def}}$
  $\forall \overset{\bullet-\bullet}{Ev}, \overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, \overset{\bullet}{Ev}_3, Tv_1, Tv_2, Tv_3[[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ec})$

[22]Note that I'm assuming the notion of telicity applies only to durative events.

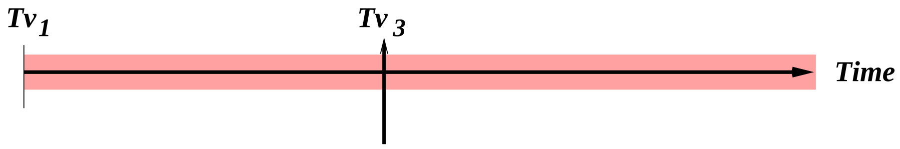

Figure 4.2: The structure of telic events; the event starts to be complete at the time the state starts to hold.

$$
\begin{aligned}
&\wedge \mathsf{Cat}(\dot{E}v_1, \uparrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\dot{E}v_1, Tv_1)\\
&\wedge \mathsf{Cat}(\dot{E}v_2, \uparrow \mathsf{Complete}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\dot{E}v_2, Tv_2)\\
&\wedge \mathsf{Cat}(\dot{E}v_3, \uparrow s) \wedge \mathsf{Occurs}(\dot{E}v_3, Tv_3)\\
&\wedge \mathsf{NoOcc}(\uparrow s, Tv_1, Tv_3)] \Rightarrow\\
&\qquad\qquad\qquad \mathsf{Equiv}(Tv_2, Tv_3)]
\end{aligned}
$$

Figure 4.2 depicts the constraints on the temporal structure of a general telic event category, *ec*, as dictated by the above definition. The shaded part of the time line represents times at which Ev does *not* start to be complete, i.e., those that cannot be the location of Tv_2. The vertical arrow marks the only time at which Tv_3 may be located: the same time at which s starts to hold.

It is enlightening to view what is going on here in the following way. The linearity of time imposes the following general constraint on Tv_2 and Tv_3 from the above definition.

(4.27) $Tv_2 \prec Tv_3 \bigvee Tv_3 \prec Tv_2 \bigvee \mathsf{Equiv}(Tv_2, Tv_3)$.[23]

What the condition for telicty does is restrict (4.27) so that only the last disjunct is possible. Thus, the consequent in the above definition is exactly (4.27) with the first two disjuncts removed. Evidently, one can get four constraints corresponding to the presence, or absence, of either of the first two constraints in (4.27). If both are absent, we get the telicity constraint. If both are present, we have the full linearity constraint. If one is present and the other is absent, we get constraints that require a strict precedence relation between Tv_2, the time of the completion of the event, and Tv_3 the time of the onset of the state. Thus, together with telic event categories, we get three more types: *left-atelic* (denoted $\overleftarrow{\mathsf{atelic}}$), *right-atelic* (denoted $\overrightarrow{\mathsf{atelic}}$), and *left-right-atelic* (denoted $\overleftrightarrow{\mathsf{atelic}}$).

[23] Strictly speaking, equality of Tv_2 and Tv_3 should be replaced by their overlapping. Neverthless, I will allow myself to be a little sloppy here. First, both Tv_2 and Tv_3 are the times of punctual events which means that they should be short enough so that overlapping is not significantly different from equality. (In Chapter 7, the instantaneous nature of punctual events will be precisely chracterized). Second, whether they are identical or merely overlapping does not have any effect on what follows.

- $\overleftarrow{\mathsf{atelic}}(\overset{\bullet-\bullet}{ec}, s) =_{\text{def}}$

$$\forall \overset{\bullet-\bullet}{Ev}, \overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \overset{\bullet}{E}v_3, Tv_1, Tv_2, Tv_3[[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ec})$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_1, \uparrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1)$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_2, \uparrow \mathsf{Complete}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow s) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3)$$
$$\wedge \mathsf{NoOcc}(\uparrow s, Tv_1, Tv_3)] \Rightarrow$$
$$[Tv_2 \prec Tv_3 \ \vee \ \mathsf{Equiv}(Tv_2, Tv_3)]]$$

- $\overrightarrow{\mathsf{atelic}}(\overset{\bullet-\bullet}{ec}, s) =_{\text{def}}$

$$\forall \overset{\bullet-\bullet}{Ev}, \overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \overset{\bullet}{E}v_3, Tv_1, Tv_2, Tv_3[[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ec})$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_1, \uparrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1)$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_2, \uparrow \mathsf{Complete}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow s) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3)$$
$$\wedge \mathsf{NoOcc}(\uparrow s, Tv_1, Tv_3)] \Rightarrow$$
$$[Tv_3 \prec Tv_2 \ \vee \ \mathsf{Equiv}(Tv_2, Tv_3)]]$$

- $\overleftrightarrow{\mathsf{atelic}}(\overset{\bullet-\bullet}{ec}) =_{\text{def}}$

$$\forall Sv, \overset{\bullet-\bullet}{Ev}, \overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \overset{\bullet}{E}v_3, Tv_1, Tv_2, Tv_3[[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ec})$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_1, \uparrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1)$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_2, \uparrow \mathsf{Complete}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$$
$$\wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow Sv) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3)$$
$$\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_3)] \Rightarrow$$
$$[Tv_2 \prec Tv_3 \ \vee \ Tv_3 \prec Tv_2 \ \vee \ \mathsf{Equiv}(Tv_2, Tv_3)]]$$

Intuitively, an event category is $\overleftarrow{\mathsf{atelic}}$ if there is a particular state such that any instance of the category starts to be complete *before or at* the time at which the state starts to hold (see Figure 4.3). Similarly, an event category is $\overrightarrow{\mathsf{atelic}}$ if there is a particular state such that any instance of the category starts to be complete *after or at* the time at which the state starts to hold (see Figure 4.4). Finally, an event category is $\overleftrightarrow{\mathsf{atelic}}$ if no state onset restricts the completion of its instances (see Figure 4.5).

The four types of event categories may be characterized as points in a 2-dimensional feature space. An event category has the $+\mathbf{L}$ (for *left*) feature (or is $+\mathbf{L}$) if $Tv_2 \prec Tv_3$ appears as a disjunct

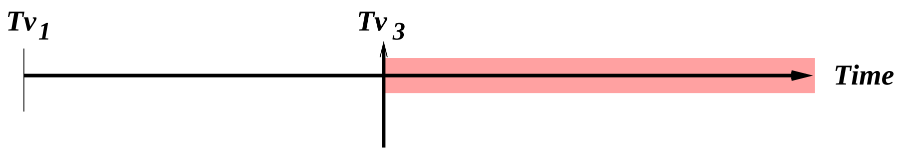

Figure 4.3: The structure of $\overleftarrow{\mathsf{atelic}}$ events; the event starts to be complete before, or at, the time the state starts to hold.

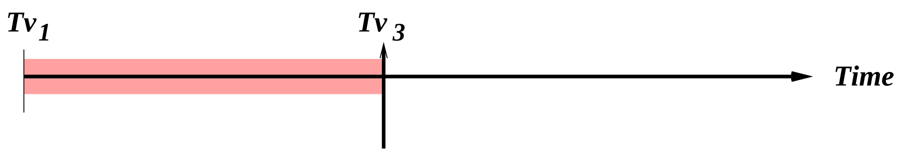

Figure 4.4: The structure of $\overrightarrow{\mathsf{atelic}}$ events; the event starts to be complete after, or at, the time the state starts to hold.

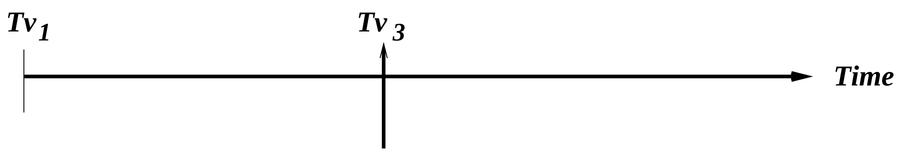

Figure 4.5: The structure of $\overleftrightarrow{\mathsf{atelic}}$ events; the time of the completion of the event is not constrained.

	$-\mathbf{R}$	$+\mathbf{R}$
$-\mathbf{L}$	telic *Run to the store*	$\overrightarrow{\text{atelic}}$ (right-atelic) *Run past the store*
$+\mathbf{L}$	$\overleftarrow{\text{atelic}}$ (left-atelic) *Run toward the store*	$\overleftrightarrow{\text{atelic}}$ (left-right-atelic) *Run*

Figure 4.6: The RL-matrix

in the consequent of the definition that applies to it. Similarly, an event category has the $+\mathbf{R}$ (for *right*) feature (or is $+\mathbf{R}$) if $Tv_3 \prec Tv_2$ appears as a disjunct in the cosequent of the definition that applies to it. This is illustrated in Figure 4.6 with examples of the four types.

Instructing Cassie to *Run* is instructing her to perform a $\overleftrightarrow{\text{atelic}}$ act; there is no state at/before/after which Cassie must stop running. Classical examples of atelicity are mostly of $\overleftrightarrow{\text{atelic}}$ events. *Run toward the store* represents a $\overleftarrow{\text{atelic}}$ act category. Cassie may stop running at any time before reaching the store. However, once at the store, she must stop running since continuing to run would be away from, not toward, the store. The class of $\overleftarrow{\text{atelic}}$ events also explains certain cases that (Dahl, 1981) discusses. For example, consider the following sentences.[24]

(4.28) John tried to build a house.

(4.29) The submarine moved toward the north pole.

According to (Dahl, 1981, p. 86), the existence of some state beyond which the process cannot continue rules out the possibility of the above sentences being atelic. Accordingly, Dahl treats them as telic. Such a move proves to be problematic as Dahl himself notices. However, given the proposed analysis, the presence of such a state beyond which the process cannot continue only means that the sentences are $-\mathbf{R}$. In that case, they could be either telic or $\overleftarrow{\text{atelic}}$, and according to our analysis they indeed are $\overleftarrow{\text{atelic}}$ (since they are $+\mathbf{L}$). This resolves the problems discussed by Dahl and at the same time supports the intuition that sentences like (4.28) and (4.29) are different from the more traditional atelic examples (i.e., those that are $\overleftrightarrow{\text{atelic}}$ according to our analysis).

[24]These are the perfective versions of (18), and (22) from (Dahl, 1981), respectively.

Examples of $\overrightarrow{\mathsf{atelic}}$ acts are those that essentially lack the **L** feature. Such acts have to reach some state but then may go on indefinitely (cf. *run past the store*). For example,

(4.30) John ran no less that 2 miles.

(4.31) John drank no less than three cups of coffee.

(4.32) They lifted at least four tables.[25]

Other examples are those analyzed by (Declerck, 1979) as sentences that "can be used to describe situations that are unnecessarily protracted beyond the potential terminal point" (Declerck, 1979, pp. 783–784).

(4.33) John painted the door.

(4.34) John sharpened the saw.

(4.35) John washed the sheet.[26]

A more elaborate discussion of the linguistic ramifications of the proposed analysis is beyond the scope of this paper; future work shall address these issues in more detail.

I will conclude this section with a couple of definitions for $\overset{\rightarrow\bullet}{\mathsf{telic}}$ and $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ event categories.

- $\overset{\rightarrow\bullet}{\mathsf{telic}}(\overset{\bullet-\bullet}{ec}, s) =_{\text{def}}$

$$
\begin{aligned}
&\mathsf{telic}(\overset{\bullet-\bullet}{ec}, s)\\
&\wedge \forall \overset{\bullet-\bullet}{Ev}, \overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, T_1, Tv_2[[\mathsf{cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ec})\\
&\qquad \wedge \mathsf{Cat}(\overset{\bullet}{E}v_1, \downarrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1)\\
&\qquad \wedge \mathsf{Cat}(\overset{\bullet}{E}v_2, \uparrow \mathsf{complete}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)] \Rightarrow\\
&\qquad\qquad \mathsf{Equiv}(Tv_1, Tv_2)]
\end{aligned}
$$

- $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}(\overset{\bullet-\bullet}{ec}, s) =_{\text{def}}$

$$\mathsf{telic}(\overset{\bullet-\bullet}{ec}, s)$$

[25] Due to (Verkuyl, 1989, p. 83).

[26] These are (respectively) sentences (15), (91a), and (91b) in (Declerck, 1979). Also see the examples in footnote 33 therein.

$$\wedge\forall \overset{\bullet-\bullet}{Ev}, \overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, T_1, Tv_2[[\mathsf{cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ec})$$

$$\wedge \mathsf{Cat}(\overset{\bullet}{Ev}_1, \downarrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_1, Tv_1)$$

$$\wedge \mathsf{Cat}(\overset{\bullet}{Ev}_2, \uparrow \mathsf{complete}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_2, Tv_2)] \Rightarrow$$

$$[\mathsf{Equiv}(Tv_1, Tv_2) \vee Tv_1 \prec Tv_2]]$$

Note that the classification induced by the **R** and **L** features is based on the temporal relation between the onset of the state *s* and that of the completion of the event. On the other hand, the internal distinction, within the class of telic categories, between $\overset{\rightarrow\bullet}{\mathsf{telic}}$ and $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ is based on the temporal relation between the end of the event and the onset of its completion. Note that, according to the above definition of $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$, an event categorized as such may *occur* but does not become *complete* until a period of time elapses.[27]

4.6.4 Recognizing an Event

In Section 4.6.2, the schema in (4.26) was introduced to indicate the condition under which, for some event *e*, the closure of $\mathsf{Int}(e)$ is identical to it. To recapitulate, the condition was that such a closure falls under all categories that *e* is believed to be an instance of. What now needs to be done is to state conditions under which the closure of $\mathsf{Int}(e)$ falls under a category of *e*. This primarily depends on what kind of category it is, according to the ontology developed in the previous section. Given the partitioning of telic into $\overset{\rightarrow\bullet}{\mathsf{telic}}$ and $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$, we have five types of categories and, hence, five axioms.[28]

- **AE14.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overset{\rightarrow\bullet}{\mathsf{telic}}\,(\overset{\bullet-\bullet}{ECv}, Sv)$

  $\wedge \mathsf{OCPair}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Cat}(\overset{\bullet}{Ev}_3, \uparrow Sv)$

  $\wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_2, Tv_2) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_3, Tv_2)$

  $\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_2)] \Rightarrow$

  $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2), \overset{\bullet-\bullet}{ECv})$

- **AE15.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overset{\rightarrow\ldots\bullet}{\mathsf{telic}}\,(\overset{\bullet-\bullet}{ECv}, Sv)$

[27] Interestingly, alluding to this distinction between completion and occurrence is how (Thomson, 1971) manages to explain the killing-shooting problem (see fn. 17).

[28] I am *not* assuming that $\overset{\bullet-\bullet}{\mathcal{EC}}$ is partitioned by these five categories. In chapter 9 a sixth category of structured events (there, I will concentrate on acts), notably sequences, is presented. Completion conditions for these events will be given within an operational semantics framework.

$$\wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow Sv)$$

$$\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3)$$

$$\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_3) \wedge [Tv_2 \prec Tv_3 \vee \mathsf{Equiv}(Tv_2, Tv_3)]] \Rightarrow$$

$$\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{ECv})$$

- **AE16.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overleftarrow{\mathsf{atelic}}\,(\overset{\bullet-\bullet}{ECv}, Sv)$

  $\wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev}))$

  $\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$

  $\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_2)] \Rightarrow$

  $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{ECv})$

- **AE17.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overrightarrow{\mathsf{atelic}}\,(\overset{\bullet-\bullet}{ECv}, Sv)$

  $\wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow Sv)$

  $\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3)$

  $\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_3) \wedge [Tv_3 \prec Tv_2 \vee \mathsf{Equiv}(Tv_2, Tv_3)]] \Rightarrow$

  $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{ECv})$

- **AE18.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overleftrightarrow{\mathsf{atelic}}\,(\overset{\bullet-\bullet}{ECv}) \wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev}))] \Rightarrow$

  $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{ECv})$

AE14 states that, if $\overset{\bullet-\bullet}{ec}$ is a $\overset{\rightarrow\bullet}{\mathsf{telic}}$ event category of $\overset{\bullet-\bullet}{e}$, then it is also an event category of the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ provided that $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ ceases at the onset of the state, s, marking the end of instances of $\overset{\bullet-\bullet}{ec}$. Similarly, **AE15** states that, if $\overset{\bullet-\bullet}{ec}$ is $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$, then the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ is an instance of it if $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ ceases at, or before, the time s starts.

AE16 states that, if $\overset{\bullet-\bullet}{ec}$ is $\overleftarrow{\mathsf{atelic}}$, then s may not start within the interior of $\overset{\bullet-\bullet}{e}$. That is, it may not start before $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ ceases. Now, one might think that this condition looks suspiciously similar to that for $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ event categories: s does not start before $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ ceases. The difference, however, is that, whereas **AE15** requires s to actually start, **AE16** does not. Thus, John cannot claim that he has pushed the rock down the hill into the river unless the rock actually gets into the river. On the other hand, John may claim that he has run toward the store although he never actually reaches the store.

AE17 should be obvious enough; if $\overset{\bullet-\bullet}{ec}$ is $\overrightarrow{\mathsf{atelic}}$, then $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ is an instance of it if it ceases only at, or after, the onset of s. **AE18** is interestingly simple. It states that the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$

is, unconditionally, an instnce of any $\overleftrightarrow{\mathsf{atelic}}$ category of $\overset{\bullet-\bullet}{e}$. Note that **AE18** does not even require that the closure actually occurs. For the other types of categories, the occurrence of the closure is required since its membership in any such category hinges on *how* it actually unfolds.

Given the above axioms, it is now possible for Cassie to *infer* when an event (most importantly, an act of hers) is complete. Suppose that Cassie conceives of some event, $\overset{\bullet-\bullet}{e}$, such that $\mathcal{C}(\overset{\bullet-\bullet}{e})$ is non-empty (i.e., $\overset{\bullet-\bullet}{e}$ is restricted to unfold in a certain way). Further, suppose that Cassie determines that $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ holds.[29] Now, if $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ ceases in such a way that its closure falls under all the categories in $\mathcal{C}(\overset{\bullet-\bullet}{e})$ (as per **AE14**–**AE18**), then, by the appropriate instance of (4.26), Cassie infers that this closure is identical to $\overset{\bullet-\bullet}{e}$. By **AE13**, Cassie may infer that $\overset{\bullet-\bullet}{e}$ is complete.

The above notwithstanding, there are two problems with the above axioms—one pertaining to the applicability of the axioms and one that has to do with their correctness. The first problem revolves around the NoOcc clauses that appear in the antecedents of **AE14** through **AE17**. These clauses are there for a good reason; for an event $\overset{\bullet-\bullet}{e}$, they temporally locate the *first* onset of the relevant state s following the onset of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$. To appreciate their role, consider what happens if the NoOcc clause is removed from **AE16** (the most striking case).

$$
\begin{aligned}
(4.36)\ & [\mathsf{Cat}(\overset{\bullet-\bullet}{Ev},\overset{\bullet-\bullet}{ECv}) \wedge \overleftarrow{\mathsf{atelic}}\ (\overset{\bullet-\bullet}{ECv},Sv) \\
& \wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1,\overset{\bullet}{E}v_2,\mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \\
& \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1,Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2,Tv_2)] \Rightarrow \\
& \qquad\qquad\qquad\qquad \mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1,\overset{\bullet}{E}v_2),\overset{\bullet-\bullet}{ECv})
\end{aligned}
$$

Essentially, (4.36) renders recognizing an event as $\overleftarrow{\mathsf{atelic}}$ almost identical to recognzing it as $\overleftrightarrow{\mathsf{atelic}}$. The NoOcc clause in **AE16** is the crucial component that distinguishes the $\overleftarrow{\mathsf{atelic}}$ and the $\overleftrightarrow{\mathsf{atelic}}$ cases; it ensures that the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ must cease before any onset of s following its own onset. Dropping that condition imposes no restrictions whatsoever on when the event reaches an end. For example, according to (4.36), it is fine for Cassie to recognize a running of hers as a running toward the store, even if that running continues beyond the store.

So the NoOcc clauses are dramatically important; why is that a problem? It is a problem because, in general, it is often impossible for Cassie to *infer* that an event does *not* occur within a

[29] This is easy if $\overset{\bullet-\bullet}{e}$ is an act to be performed by Cassie (see Chapter 5). For other types of events, see the discussion below.

period of time. A weaker requiremenet is needed. In particular, instead of requiring that Cassie knows that an event (namely an onset of the appropriate *s*) does not occur, we should only require Cassie *not to know* of any such occurrence. This seems both sufficient and realistic. However, it should be clear what it entails: non-monotonicity. For Cassie to reason based on lack of knowledge, she must be endowed with non-monotonic (auto-epistemic, for those who mind the distinction) reasoning capabilities.

Now, the axioms may be fixed by putting the NoOcc clauses within the scope of the non-monotonic modal operator **M** of (McDermott and Doyle, 1980) (which is the dual of the **L** operator of (Moore, 1985) and (Konolige, 1994)). A clause of the form $\mathbf{M}\mathsf{NoOcc}(ec, t_1, t_2)$ would mean that $\mathsf{NoOcc}(ec, t_1, t_2)$ is consistent with what Cassie believes. In other words, it means that $\beta \nvdash \neg\mathsf{NoOcc}(ec, t_1, t_2)$, which is what we need. Although this would work, incorporating non-monotonicity within the SNePS system is still under investigation (but see (Cravo and Martins, 1993)) and I am not willing to introduce it in the axioms without fully investigating all its side-effects. However, as will be shown in Chapter 9, there are certain heuristics that may be adopted in order to overcome the difficulties with using the above axioms.

Let us now turn to the second problem. The careful reader would notice that axioms **AE14**–**AE17** are actually not correct. Take **AE14** for example. The only restrictions on the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ to be of a $\overset{\rightarrow\bullet}{\mathsf{telic}}$ category of $\overset{\bullet-\bullet}{e}$ are purely temporal, mere constraints on the temporal positioning of some events. As is widely known, this is not correct. In particular, there has to be some sort of a *causal* relationship between the closure of $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ and the onset of the state *s* signalling its completion. To take the example of (Goldman, 1970) from Section 4.6.2, John might start playing the piano and Brown might eventually wake up. However, for John to claim that *he* has woken up Brown, his playing the piano should be the cause of Brown's waking up. This, of course, need not be the case, for one can come up with numerous scenarios in which, although John plays the piano, Brown wakes up as a result of something totally different. In fact, the semantics of accomplishment sentences (corresponding to telic event categories) traditionally have causality built in (Dowty, 1977; Moens and Steedman, 1988; Steedman, 2000, for example). In AI, causality has typically been represented as a primitive notion (McDermott, 1982; Allen, 1984; Lifschitz, 1987; Stein and Morgenstern, 1994; Lin, 1995; Steedman, 2000).[30] After discussing problems with

[30]But see (Shoham, 1986), for instance.

attempting an explanatory, non-primitive approach to causality, (McDermott, 1982) concludes:

> I assume that there is no way to get around this problem, and that there is no way to infer causality merely from correlation. So we will not try to define causality in terms of something more basic. Instead, we will assume whatever causal assertions we need, and infer events from them. (McDermott, 1982, p. 117)

This is fine as far as it goes, but, in our case, it does not go very far. First, most of the authors cited above define causality over the domain of event categories, primarily to state the effects of actions. For example, the following axiom appears in (Lifschitz, 1987):

causes(shoot, loaded, false)

The intended interpretation is that shooting a gun causes it to become unloaded. But this is a general action-effect axiom; Lifschtiz's *causes* does not state that a particular event actually causes anything.

(Allen, 1984), on the other hand, defines causality over particular occurrences of event categories (though he does not use these terms since the details of his system are different from ours). But this still does not help. Granted, we can introduce causality as a primitive—for example, a function, Cause, from $\mathcal{E} \times \mathcal{E}$ to ETERNAL. We can then add $\mathsf{Cause}(\mathsf{Clos}(\dot{E}v_1, \dot{E}v_2), \dot{E}v_3)$ (or something similar) as a conjunct to the antecedents of the above axioms. Although the axioms would then *look* correct, they would not be helpful for all practical purposes. As (McDermott, 1982) points out (see above), expressions such as $\mathsf{Cause}(e_1, e_2)$ are mainly used to infer the occurrence of events: given that e_1 has occurred and that e_1 causes e_2, we can infer that e_2 has also occurred (Axiom (O.4) of (Allen, 1984)). Or, within an abductive framework, given that e_2 occurs, one can defeasibly infer the occurrence of e_1 (to explain e_2). In our case, however, the situation is very different. Here, Cassie is given that both e_1 and e_2 have occurred and she is supposed to figure out whether e_1 has caused e_2. Evidently, an approach where causality is a primitve notion does not help; we need to have a theory of how agents can infer causality given a set of domain axioms and the information available in the pattern of states and their transitions that the agent has access to. Even non-primitive accounts of causality (Lewis, 1973; Shoham, 1986, for example) do not provide such a theory. I believe that this a major epistemological problem, and I do not think that any solutions could be provided here without taking us far beyond the scope of this work. Admittedly, this is a *hole* in the theory, but at least we have identified where it is (axioms **AE14–AE17**).[31]

The situation, however, is not too bad. Leaving causality out of axioms **AE14–AE17** amounts to making the following default assumption:

1. if some event $\overset{\bullet-\bullet}{e}$ is believed to be in progress, and

2. if $\overset{\bullet-\bullet}{e}$ is expected to have a certain effect s, and

3. if s starts to hold, then,

   implicitly, $\overset{\bullet-\bullet}{e}$ has caused the onset of s.[32]

I believe that the above assumption is reasonable for a naive theory of events. Of course, it does not always work. For, if one knows of an event e' that is the actual cause of s's starting to hold, then the conclusion would not be valid. Note, however, that this knowledge requires the reasoning agent to have some notion of causality. If the agent has no such notion—if it does not know anything about causality, only temporal correlations—then nothing can invalidate the (implicit) conclusion that $\overset{\bullet-\bullet}{e}$ has caused the onset of s. Of course, this does not mean that the conclusion is correct; it only means that it is *consistent* with the agent's view of the world. Given that the theory is supposed to account for Cassie's beliefs about the world, rather than the world itself, the above should be fine for the scope of this work.[33]

4.7 Concluding Remarks

The main result of this chapter is a fairly detailed account of how Cassie may construct the occurrence and completion of events by the mere monitoring of states—the only thing available to an agent reasoning and acting in time. To put it this way, however, is a little misleading, for there is another factor that is crucial to the account developed here. Not only can Cassie monitor *current* states, but she may also conceive of, and reason about, events before they actually occur. This is embodied in the function Int and everything that is built on it.

[31] Actually, it could be argued that the hole lies in disregarding causality in the definitions of the different types of telicity/atelicity.

[32] "Implicitly", because there is no actual account of causality in the theory.

[33] If, for example, the theory is extended to account for Cassie's reasoning about the actions of other agents in the domain, such that responsibility for an action becomes crucial, then an account of causality would be inevitable.

Implicit in the above-presented analysis is an account of the progressive aspect and Dowty's *imperfective paradox* (Dowty, 1977). First, it should be noted that the account is very similar in spirit to that of (Parsons, 1989). Parsons's account may be stated thus: a non-progressive (perfective) sentence asserts that an event of a particular category *culminates* (*completes* in our terminology), while the progressive form of the same sentence asserts that the event only *holds*. One difference between the current proposal and that of (Parsons, 1989) is that, in our case, events do not hold; a unique associated state does. This is, evidently, a superficial difference, but it fits better within the entire state-based framework.[34] The main difference, however, between the current approach and that of Parsons is that the current approach *explains* the relation between the occurrence of an event $\overset{\bullet-\bullet}{e}$ and the holding of the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ (**AE13–AE18** together with (4.25)) whereas (Parsons, 1989) does not provide any formal account of the relation between holding and culmination.

Dowty's imperfective paradox is basically concerned with explaining data like the following. In particular, the semantics of the progressive should account for why (4.37) entails (4.38), whereas (4.39) does not entail (4.40).

(4.37) Mary was pushing a cart.

(4.38) Mary pushed a cart.

(4.39) Mary was crossing the street.

(4.40) Mary crossed the street.

I believe that the current proposal predicts these entailments. Consider the following $\mathcal{FOCS}$ representation of the (tenseless) logical form of (4.37) and (4.38), respectively.

(4.41) $\exists\, \overset{\bullet-\bullet}{Ev}, Tv, Xv[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \mathsf{Push}(\mathsf{Mary}, Xv)) \wedge \mathsf{Cart}(Xv) \wedge \mathsf{Holds}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev}), Tv)]$

(4.42) $\exists\, \overset{\bullet-\bullet}{Ev}, Tv, Xv[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \mathsf{Push}(\mathsf{Mary}, Xv)) \wedge \mathsf{Cart}(Xv) \wedge \mathsf{Occurs}(\overset{\bullet-\bullet}{Ev}, Tv)]$

[34]In fact, (Parsons, 1989, note 16) briefly discusses this same approach of positing a unique progressive state corresponding to an event stating that "it may be equally as good, or even better, than [his own]" (Parsons, 1989, p. 239).

Now, given that pushing a cart is $\overset{\longleftrightarrow}{\mathsf{atelic}}$, then, by **AE18**, the closure of the state reported by (4.37) ($\mathsf{Int}(\overset{\bullet-\bullet}{Ev})$ in (4.41)) is an *event* of pushing a cart. Note that this event *has* to occur as per **AE10**. The occurrence of such an event is exactly what (4.38) reports.

Similarly, the follwing $\mathcal{FOCS}$ expressions represent the logical forms of (4.39) and (4.40), rep-spectively.

(4.43) $\exists \overset{\bullet-\bullet}{Ev}, Tv[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \mathsf{Cross}(\mathrm{Mary}, \mathrm{STREET})) \wedge \mathsf{Holds}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev}), Tv)]$

(4.44) $\exists \overset{\bullet-\bullet}{Ev}, Tv[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \mathsf{Cross}(\mathrm{Mary}, \mathrm{STREET})) \wedge \mathsf{Occurs}(\overset{\bullet-\bullet}{Ev}, Tv)]$

In this case, crossing the street is $\overset{\rightarrow\bullet}{\mathsf{telic}}$, and one cannot infer that the closure of the state reported by (4.39) is an event of Mary's actually crossing the street given (4.43) alone. In particular, one would also need something to the effect that Mary got to the other side of the street at the end of the state reported by (4.39) (so that **AE14** would apply).

The main difference between the current proposal (which, as mentioned above, is similar to (Parsons, 1989)) and the more common proposals of, for example, (Dowty, 1977; Landman, 1992) may be stated this way. In traditional analyses, a progressive sentence like (4.39), roughly, reports the occurrence of an event e which is a part (or a *stage* à la (Landman, 1992)) of an event, e', of Mary's successfully crossing the street that occurs in all possible worlds where things go as expected.[35] The current analysis does not allude to possible worlds at all; it trades possible worlds for the existence of incomplete events. This is justified since the ontology underlying $\mathcal{FOCS}$ is an epistemological one, and incomplete events are, arguably, possible objects of thought (see Section 4.6.1). A sentence in the progressive reports the holding of a state of the form $\mathsf{Int}(\overset{\bullet-\bullet}{e})$, where $\overset{\bullet-\bullet}{e}$ is an event, one that is not known to be complete, conceived of under certain categories. So what exactly is the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$? According to the above-presented system, it is a state which, if it comes to a closure in a certain way as dictated by the categorical restrictions on $\overset{\bullet-\bullet}{e}$, would give rise to an actual occurrence of $\overset{\bullet-\bullet}{e}$. The only catch is that the formal theory does not provide any criteria, independent of the actual occurrence of $\overset{\bullet-\bullet}{e}$, based on which it could be determined that $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ holds. The reply to this is two-fold. First, even if this is true, I am not aware of any other approach to the semantics of the progressive that does not embody a similar mystery. For example, as mentioned

[35] These are *inertia worlds* for (Dowty, 1977) and *reasonable options* for (Landman, 1992).

above, (Dowty, 1977) assumes the existence of some event that culminates in some possible world to account for progressive sentences. However, nowhere does Dowty specify what it is about the event that is actually going on in the real world that allows speakers to envision such a possible completion.[36] Second, if $\overset{\bullet-\bullet}{e}$ is one of Cassie's actions, then determining when $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ holds is easy; it is taken care of at the meta-theoertical level by an explicit assertion that is made once Cassie starts performing $\overset{\bullet-\bullet}{e}$ (see Chapter 5). If $\overset{\bullet-\bullet}{e}$ is some other external event, then, unless there are explicit domain rules specifying when $\overset{\bullet-\bullet}{e}$ is to be considered in progress, there is no way for Cassie to recognize that. For the purpose of this work, however, I am only interested in progressive states of Cassie's own actions and, as such, I can safely put the issue of external events aside.

[36]See the discussion of this point by Dowty himself (Dowty, 1977, fn. 8, p. 58).

Chapter 5

Temporal Progression

In this chapter, we turn, from the logical theory developed in the previous two chapters, to the *meta-theory*. Whereas the logical theory represents the contents of Cassie's mind (basically, $\mathcal{FOCS}$ expressions), the meta-theory does not. It represents interesting structures of those contents, relations among those structures, and a detailed account of how they evolve over time. Of particular importance, the meta-theory provides an account of temporal progression: how the passage of time is represented and how Cassie's beliefs are updated as time goes by. In what follows, I will be mainly concerned with temporary states. Therefore, in informal discussions, "state" would typically mean "temporary state". Formally, however, I shall clearly point out the state-sorts. It should be noted, though, that, since CO-PERM states are of no relevance to the system, and since ETERNAL states are more propositional in character (see Section 3.6), the formalism will only cover TEMP and PERM states.

5.1 The Cognitive Structure of Time

5.1.1 Temporal Frames

In Section 3.4, the basic logic of time was introduced. In this section, we investigate the meta-theoretical structures articulating the domain of time intervals. First, let us restrict the discussion to those intervals in Cassie's "consciousness".

Definition 5.1 *For any sort Σ, define $\Psi(\Sigma)$ as the largest subset of $\mathcal{D}_\Sigma$ (the set of denotata of terms of sort Σ) such that for every $[\![\tau]\!] \in \Psi(\Sigma)$, $[\![\tau]\!] \in \beta$ or there is a term τ' such that $[\![\tau']\!] \in \beta$ and τ is a* *subterm of* τ'.

Intuitively, $\Psi(\Sigma)$ represents those entities in $\mathcal{D}_\Sigma$ that Cassie conceives of.[1] It should be noted that $\Psi(\text{ETERNAL})$ is not identical to β, since Cassie may conceive of propositions she does not believe. For convenience, I shall henceforth drop the syntax-semantics distinction. Thus, I will use "τ" in place of "$[\![\tau]\!]$" and "Σ" in place of "$\mathcal{D}_\Sigma$". In addition, where $p \in$ ETERNAL, I will often use "p" where what is intended is "$[\![p]\!] \in \beta$". This should not be confusing given the one-to-one correspondence between terms and their denotations.

Definition 5.2 *An interval* $t \in \Psi(\mathcal{T})$ *is an* **atomic interval** *if there is no* $t' \in \Psi(\mathcal{T})$ *such that* $t' \sqsubset t$.[2]

Note that an interval being atomic is not an intrinsic property of the interval itself; it is totally dependent on Cassie's *state of mind*. Cassie's state of mind may be represented by the set of all terms in Cassie's mind:

$$\bigcup_{\Sigma} \Psi(\Sigma),$$

where Σ is a $\mathcal{FOCS}$ sort. This set changes with time, since Cassie may acquire more knowledge as time passes by.

Consider the following sentences.

(5.1) The vase fell on the floor.

(5.2) John tried to catch it.

Sentence (5.1) reports a punctual event. After hearing (5.1), Cassie would have certain beliefs about the time, t, of that event. None of these beliefs, however, are about any subinterval of t. In that *state of mind*, t is an atomic interval. Later, after hearing (5.2), this may change. Cassie would have a belief about a subinterval, t', of t, the interval over which John tried catching the falling vase. Now Cassie is in a state of mind in which t is not atomic.[3] Indeed, to be more precise, one

[1] See (Shapiro, 1991, p. 143) for a different presentation of the same notion.

[2] *Atomic intervals are, therefore, similar to the* moments *of (Allen and Hayes, 1985).*

[3] Recall a similar discussion in Section 3.4

should relativize all definitions to Cassie's state of mind. I choose, however, to be more liberal while stressing that all of the definitions should be interpreted in the context of a particular state of mind.

The relations $\prec$ and $\sqsubset$ articulate the set $\Psi(\mathcal{T})$ giving rise to what are called *temporal frames*.

Definition 5.3 *A set $\Phi \subseteq \Psi(\mathcal{T})$ is a* **temporal frame** *if there exists $t \in \Psi(\mathcal{T})$ such that, for every $t' \in \Phi$, $t \sqsubset t'$ (is in β) or $t' = t$.*[4]

The above definition may be interpreted as saying that a temporal frame is a set of intervals that share a common subinterval. It should be noted that, in the above definition, "$t \sqsubset t'$ (is in β) or $t' = t$" is not equivalent to "$t \sqsubseteq t'$ (is in β)", the latter is equivalent to "$t \sqsubset t'$ (is in β) or $\mathsf{Equiv}(t', t)$ (is in β)".

The above definition is by itself not very interesting since it covers a lot of trivial cases. For instance, any singleton subset of $\Psi(\mathcal{T})$ is a temporal frame. A more conservative notion is required.

Definition 5.4 *A temporal frame Φ is* **maximal** *if there is no temporal frame Φ' such that $\Phi \subset \Phi'$.*

It could be shown that, if Φ is a maximal temporal frame (henceforth, **MTF**), then $\langle \Phi, \lambda x \lambda y([\![x \sqsubset y]\!] \in \beta \vee x = y)\rangle$ is a poset with a smallest element (which makes it a *meet semilattice* (Link, 1998)).[5]

Observation 5.1 *If Φ is an MTF, then the poset $\langle \Phi, \lambda x \lambda y([\![x \sqsubset y]\!] \in \beta \vee x = y)\rangle$ has a smallest element.*

Proof. To show that $\langle \Phi, \lambda x \lambda y([\![x \sqsubset y]\!] \in \beta \vee x = y)\rangle$ has a smallest element, we need to prove that there is some $t \in \Psi(\mathcal{T})$ such that:

(i) for every $t' \in \Phi$, $\langle t, t'\rangle \in \lambda x \lambda y([\![x \sqsubset y]\!] \in \beta \vee x = y)$ and

(ii) $t \in \Phi$.

By Definition 5.3, there exists $t \in \Psi(\mathcal{T})$ such that for every $t' \in \Phi$ $t \sqsubset t' \in \beta$ or $t' = t$, which proves (i). Now, suppose that such an interval, t, is not in Φ. Then the set $\Phi \cup \{t\}$ is a super-set of Φ. Since, by designation, t is a sub-interval of every element of Φ, then, by Definition 5.3, $\Phi \cup \{t\}$ is itself a temporal frame. But, by Definition 5.4, this implies that Φ is not maximal, which leads to a contradiction. Therefore, $t \in \Phi$, which proves (ii). $\square$

[4] *Note that '=' is a meta-theoretical predicate representing term (or denotation) identity. This is not to be confused with the object language* Equiv *which corresponds to co-reference of denotations.*

[5] I am overloading '$\vee$', using it as meta-theoretical disjunction.

The smallest elements of MTFs have an interesting property: they are atomic.[6]

Proposition 5.1 *If Φ is an MTF, then a smallest element of Φ is an atomic interval.*

Proof. Let t be a smallest element of Φ. Assume that t is not atomic. By Definition 5.2, there is some $t' \in \Psi(\mathcal{T})$ such that $t' \sqsubset t$. Since, for all $t'' \in \Phi$, $t'' \neq t$ implies $t \sqsubset t''$, then, by the transitivity of $\sqsubset$, t' is a proper sub-interval of all elements of Φ. Thus, by Definition 5.3, $\Phi \cup \{t'\}$ is a temporal frame. If $t' \notin \Phi$, then $\Phi \subset \Phi \cup \{t'\}$ and, by Definition 5.4, Φ is not an MTF, leading to a contradiction. On the other hand, if $t' \in \Phi$, then t cannot be a smallest element of Φ, which also leads to a contradiction. Therefore, t must be atomic. $\square$

Not only is a smallest element of an MTF atomic, but it is the *only* atomic interval therein.[7]

Theorem 5.1 *For every MTF, Φ, there is one and only one atomic interval in Φ.*

Proof. Let Φ be an MTF. By Observation 5.1, Φ contains a smallest interval, t, which, by Proposition 5.1, is atomic. Now, we need to show that t, is the only atomic interval in Φ. Since t is a smallest element of $\langle \Phi, \lambda x \lambda y([\![x \sqsubset y]\!] \in \beta \vee x = y) \rangle$, then, for every $t' \in \Phi$, $t' \neq t$ implies that $t \sqsubset t' \in \beta$. Thus, by Definition 5.2, for every $t' \in \Phi$, $t' \neq t$ implies that t' is not atomic. Therefore, t is the only atomic interval in Φ. Since Φ is arbitrary, then the result applies to all MTFs. $\square$

The following corollary directly follows.

Corollary 5.1 *Every MTF has a unique smallest element.*

Proof. Follows directly from Observation 5.1, Proposition 5.1, and Theorem 5.1. $\square$

Given the above results, one can outline an algorithm for computing the MTFs of a set $\Psi(\mathcal{T})$.

1. For every atomic interval $t_i \in \Psi(\mathcal{T})$, let $\Phi_i \longleftarrow \{t_i\}$.

2. For every $t \in (\Psi(\mathcal{T}) - \bigcup_i \Phi_i)$

    3. For every Φ_i

        4. If $t_i \sqsubset t$, let $\Phi_i \longleftarrow \Phi_i \cup \{t\}$.

---

[6] In what follows, I will be talking about "a smallest element of Φ" (where Φ is an MTF) referring to a smallest element of the poset $\langle \Phi, \lambda x \lambda y([\![x \sqsubset y]\!] \in \beta \vee x = y) \rangle$.

[7] The notion of MTFs is related to that of *filters* (van Benthem, 1983, Ch. I.4). Informally, van Benthem's definition of filters maps as follows onto our system. A subset F of $\Psi(\mathcal{T})$ is a filter if (i) $t \in F$ implies that all super-intervals of t are in F and (ii) t and t' are in F implies that their maximal common sub-interval (if one exists) is also in F. Accordingly, an MTF is a filter, but not every filter is an MTF. In particular, a filter may include more than one atomic interval.

The resulting Φ_i's are the MTFs of $\Psi(\mathcal{T})$. Note that, for each Φ_i, t_i is its unique smallest element. The intuitive construction of MTFs represented by the above algorithm assumes the following result which, together with Theorem 5.1, point to a one-to-one correspondence between MTFs and atomic intervals.

Theorem 5.2 *For every atomic interval, t, there is one and only one MTF to which t belongs.*

Proof. Let t be an atomic interval. Consider the set consisting of *all* intervals, t', satisfying the property $t \sqsubset t' \in \beta$ or $t' = t$. By Definition 5.3, this set is a temporal frame, and is maximal since it includes all such t's. Now, we need to show that t belongs to a unique MTF. Suppose that Φ_1 and Φ_2 are distinct MTFs to which t belongs. Then there is an interval, t', that belongs to Φ_1 and not Φ_2. Since t is atomic, then, by Theorem 5.1, it is the smallest element of both Φ_1 and Φ_2 and, therefore, $t \sqsubset t'$. By Definition 5.3, the set $\Phi_2 \cup t'$ is a temporal frame which, since $t' \notin \Phi_2$, is a super-set of Φ_2. Thus, by Definition 5.4, Φ_2 is not maximal, leading to a contradiction. Therefore, t belongs to only one MTF. Since t is arbitrary, then the result applies to all atomic intervals. □

Given Theorems 5.1 and 5.2, the notation $\Phi(t)$ (where t is atomic) will be used to refer to the unique MTF whose smallest element is t. Going back to the above algorithm for the construction of MTFs, we can show that the union of the MTFs of a set $\Psi(\mathcal{T})$ is identical to $\Psi(\mathcal{T})$. More specifically, we can make the following observation.

Observation 5.2 *The collection of MTFs of* $\Psi(\mathcal{T})$ *constitutes a minimal cover of* $\Psi(\mathcal{T})$.

Proof. First, we show that the set of MTFs is a cover of $\Psi(\mathcal{T})$ and, then, we prove its minimality. Let C be the collection of MTFs of $\Psi(\mathcal{T})$. To show that C covers $\Psi(\mathcal{T})$, we only need to prove that, for every $t \in \Psi(\mathcal{T})$,

$$t \in \bigcup_{\Phi_i \in C} \Phi_i.$$

Consider an arbitrary $t \in \Psi(\mathcal{T})$. If t is atomic, then, by Theorem 5.2, $t \in \Phi(t)$. On the other hand, if t is not atomic, then there must be some atomic interval, t', such that $t' \sqsubset t$. But, by Definition 5.4, $t \in \Phi(t')$. Therefore, for every $t \in \Psi(\mathcal{T})$, there is some MTF, Φ, such that $t \in \Phi$. Thus, for every $t \in \Psi(\mathcal{T})$,

$$t \in \bigcup_{\Phi_i \in C} \Phi_i,$$

which means that the collection of MTFs covers $\Psi(\mathcal{T})$. To prove that the covering is minimal, we need to show that for every $\Phi \in C$, there is some $t \in \Psi(\mathcal{T})$ such that

$$t \notin (\bigcup_{\Phi_i \in C} \Phi_i) \setminus \Phi.$$

Let Φ be an arbitrary member of C. By Theorem 5.1, there is a unique atomic interval, t, such that $t \in \Phi$. By Theorem 5.2, Φ is the only MTF to which t belongs. Therefore,

$$t \notin (\bigcup_{\Phi_i \in C} \Phi_i) \setminus \Phi.$$

Since Φ is an arbitrary member of C, then C minimally covers $\Psi(\mathcal{T})$. $\square$

It should be noted that the collection of MTFs does not partition $\Psi(\mathcal{T})$. The reason is that a non-atomic interval, t, may span more than one MTF, i.e., MTFs are not disjoint.

MTFs are, in general, similar to the situations of the situation calculus (McCarthy and Hayes, 1969); they represent snapshots, not of the universe, but of Cassie's conceptualization of it. A general, not essentially maximal, temporal frame corresponds to the *states of affairs* of situation semantics (Barwise and Perry, 1983). Note, though, that whereas the latter is objective, the former is mental. As fluents may hold in situations, states hold in MTFs.

Definition 5.5 *For every* $s \in \Psi(\textsc{Temp}) \cup \Psi(\textsc{Perm})$ *and temporal frame* Φ*,* s **holds in** Φ*, if there is some* $t \in \Phi$ *such that* $\mathsf{Holds}(s,t)$.

Note, in particular, that if s holds in an MTF, $\Phi(t)$, then (by **AS2**) Cassie may conclude that $\mathsf{Holds}(s,t)$. In addition, if there is some t' such that $\mathsf{Holds}(s,t')$ and there is some atomic interval, t, such that $t \sqsubset t'$, then s holds in $\Phi(t)$ since otherwise $\Phi(t)$ would not be maximal. Thus, an MTF with a smallest element, t, corresponds to the set of all states that Cassie believes, or may conclude they, hold over t. Figure 5.1 shows two MTFs with the states holding in them. The MTFs are represented by rectangles and their contents by meta-variables. The smallest element of an MTF is shown near the center of the bottom side of the rectangle. States are represented by meta-variables, with lines connecting states and time intervals standing for the Holds relation.

A precise characterization of the difference between two MTFs may be given an epistemic interpretation.

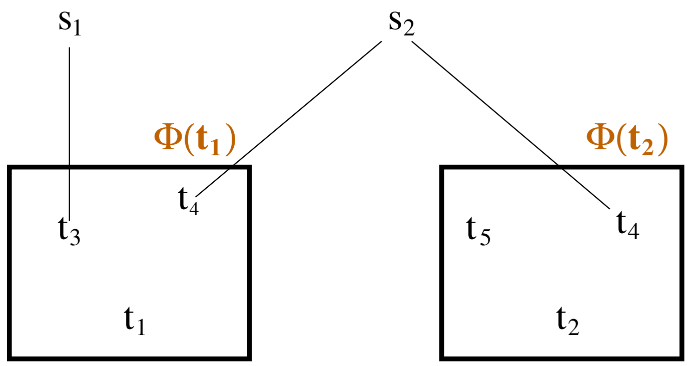

Figure 5.1: Two MTFs, $\Phi(t_1)$ and $\Phi(t_2)$. The state s_1 holds in $\Phi(t_1)$ and the state s_2 holds in both $\Phi(t_1)$ and $\Phi(t_2)$.

Definition 5.6 *The* **epistemic distance** *between two MTFs, Φ_1 and Φ_2, denoted $d_e(\Phi_1, \Phi_2)$, is the cardinality of their symmetric difference. That is, $d_e(\Phi_1, \Phi_2) = |\Phi_1 \triangle \Phi_2|$.*

For example, the epistemic distance between the two MTFs shown in Figure 5.1 is 4. Depending on their epistemic distance, two MTFs are more, or less, similar to each other as far as the states of affairs they correspond to are concerned. In fact, the set of MTFs together with d_e form a metric space.

Observation 5.3 *The function d_e defines a metric over the set of MTFs.*

Proof. For d_e to be a metric, it must satisfy the following (where $\Phi(t_1)$, $\Phi(t_2)$, and $\Phi(t_3)$ are MTFs):

(a) $d_e(\Phi(t_1), \Phi(t_2)) \geq 0$.

(b) $d_e(\Phi(t_1), \Phi(t_2)) = 0$ if and only if $\Phi(t_1) = \Phi(t_2)$.

(c) $d_e(\Phi(t_1), \Phi(t_2)) = d_e(\Phi(t_2), \Phi(t_1))$.

(d) $d_e(\Phi(t_1), \Phi(t_3)) \leq d_e(\Phi(t_1), \Phi(t_2)) + d_e(\Phi(t_2), \Phi(t_3))$.

The proof is straightforward and follows from the definition of d_e. In particular, the cardinality of the symmetric difference is a metric over any class of finite sets. (a), (b), and (c) are obvious, and (d) follows from the fact that $A \triangle B \subseteq (A \triangle C) \cup (C \triangle B)$, for any sets A, B, and C. $\square$

The above observation constrains epistemic distance to be non-negative. However, being a metric over *MTFs* imposes yet another constraint.

Theorem 5.3 *For any two distinct MTFs, $\Phi(t_1)$ and $\Phi(t_2)$, $d_e(\Phi(t_1),\Phi(t_2)) \geq 2$.*

Proof. Since $\Phi(t_1)$ and $\Phi(t_2)$ are distinct MTFs, then, by Thereom 5.2, t_1 and t_2 are distinct atomic intervals. By Theorem 5.1, an MTF may have one and only one atomic interval. It follows that, $\{t_1,t_2\} \subseteq \Phi(t_1)\triangle\Phi(t_2)$. Therefore, $d_e(\Phi(t_1),\Phi(t_2)) \geq 2$. $\square$

As shown above, the relation $\sqsubset$ provides the internal structure of MTFs. The relation $\prec$ provides the external structure.

Definition 5.7 *An MTF, $\Phi(t_1)$,* **precedes** *another MTF, $\Phi(t_2)$, (or $\Phi(t_2)$* **follows** *$\Phi(t_1)$) if and only if $t_1 \prec t_2$.*

Of course, since the precedence relation over MTFs is based on $\prec$, it is a strict partial order. Although there may be a situation where Cassie does not have any beliefs about the relative order of various MTFs, some structure may still be retrieved. In particular, MTFs form clusters corresponding to their intersections.

Definition 5.8 *For every $t \in \Psi(\mathcal{T})$, the* **span** *of t is the set $Span(t) = \{\Phi : \Phi$ is an MTF and $t \in \Phi\}$. For any $s \in \Psi(\text{TEMP}) \cup \Psi(\text{PERM})$ and any set, A, of MTFs, s* **spans** *A if there is some $t \in \Psi(\mathcal{T})$ such that $\mathsf{Holds}(s,t)$ and $A \subseteq Span(t)$.*

The MTFs in the span of an interval, t, correspond to different *pieces* of t. Since intervals are convex, those MTFs form clusters that, although not internally ordered, are certainly *closer* to each other than to MTFs not containing t. Such clusters of MTFs form *episodes* in Cassie's memory: a collection of related and temporally contiguous events (see (Rumelhart et al., 1972; Tulving, 1972)). Cassie may not know the exact order of intervals within an episode, but she may know that one episode is earlier or later than another if they correspond to the spans of some time intervals, t_1 and t_2, where $t_1 \prec t_2$ (or vice versa). Note that this knowledge is essentially based on **AT8** and **AT9**.

5.1.2 The Passage of Time

A particularly interesting subset of MTFs forms a linearly-ordered chain corresponding to the experienced progression of time. Cassie's sense of temporal progression is modeled by a deictic metavariable, NOW, that assumes values from amongst the members of $\mathcal{T}$ (Almeida and Shapiro, 1983; Almeida, 1995; Almeida, 1987; Shapiro, 1998; Ismail and Shapiro, 2000b). At any time,

Algorithm `move_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{{}^*\mathsf{NOW} \prec t\}$.
3. $\mathsf{NOW} \longleftarrow t$.

Figure 5.2: The algorithm `move_NOW`.

NOW *points* to a particular member of $\Psi(\mathcal{T})$. This represents Cassie's sense of the current time at the finest level of granularity (see Chapter 7).

The movement of time is represented by NOW's *moving* (changing its value) to a different term in $\Psi(\mathcal{T})$. Depending on what NOW exactly represents, there may, or may not, be restrictions on its movement. For example, if NOW represents a narrative *now-point* (Almeida and Shapiro, 1983; Almeida, 1987; Almeida, 1995; ter Meulen, 1997), then there may be no restrictions at all on the values it assumes; NOW may freely hop around in $\Psi(\mathcal{T})$. This is because narration may go back and forth in time and may be about temporally unrelated episodes. On the other hand, if NOW represents the *real* present for an acting agent, which is how I am using it, then there certainly are restrictions on its movement. First, whenever it moves, NOW moves to a new term. That is, a change in the value of NOW is always associated with a change in $\Psi(\mathcal{T})$, since at least the new present enters into Cassie's consciousness. Second, values of NOW form a chain of times linearly-ordered by $\prec$.[8] For any $\Psi(\mathcal{T})$, NOW is always pointing to the greatest element of the chain (the *newest* present). The movement of time is thus modeled by the algorithm `move_NOW` shown in Figure 5.2. *NOW denotes the term pointed to by NOW (i.e., "*" is a dereferencing operator). Note that although a change in NOW (step 3) is always associated with a change in Cassie's state of mind (step 2), the converse is not necessarily true. It all depends on what causes NOW to move. In the current status of the theory, NOW moves whenever Cassie becomes aware of a change in the environment. The "environment" here does not include Cassie's own state of mind. Thus, Cassie's noticing that the walk-light turns from red to green, her starting to move, or her sensing that her battery is low (for a battery-operated Cassie) results in NOW moving. However, mere inferences that do not involve any interaction with the environment (for example, inferring she can cross the

[8]Thus, the presented model is silent about the issue of *forgetting*.

Algorithm `initialize_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\mathsf{NOW} \longleftarrow t$.

Figure 5.3: The algorithm `initialize_NOW`.

street having perceived that the walk-light is green) and that change Cassie's state of mind do not move NOW. Thus, the actual movement of NOW (i.e., the implementation of `move_NOW`) is taken care of at the PML. Generally, however, NOW may move with every inference step, providing Cassie with a fine-grained sense of temporal progression.[9]

Algorithm `move_NOW` takes care of *changing* the value of NOW; we still need to account for *initializing* it. This is illustrated in Figure 5.3. Note that algorithm `initialize_NOW` is identical to algorithm `move_NOW` without the second step which requires a previous value of NOW. I make the following reasonable assumptions about the temporal career of the variable NOW.

1. Algorithms `move_NOW` and `initialize_NOW` are the only places where NOW is set. The first time NOW is set is by step 2 of `initialize_NOW`. Subsequent changes to NOW are the result of step 3 of `move_NOW`.

2. At any time there is at most one execution of algorithm `move_NOW` going on.

Together, the above two assumptions mean that the value of NOW changes sequentially with time. For ease of notation, I will use numerical subscripts to refer to the successive values of NOW. Thus, $*\mathsf{NOW}_i$ is the value of NOW at a time earlier than the time at which $*\mathsf{NOW}_j$ is the value of NOW if and only if i is less than j. Thus, $*\mathsf{NOW}_1$ is the first value of NOW, $*\mathsf{NOW}_2$ is the second, and so on. If i is the largest subscript, then both $*\mathsf{NOW}$ and $*\mathsf{NOW}_i$ refer to the latest value of NOW. Algorithm `move_NOW` guarantees that this real-time ordering of the values of NOW corresponds to a $\prec$-chain of those values.

Theorem 5.4 *For all $i \in \mathbb{N}$ ($i > 0$),* $\beta \vdash *\mathsf{NOW}_i \prec *\mathsf{NOW}_{i+1}$.

Proof. Let $i \in \mathbb{N}$ ($i > 0$). According to the above-stated assumptions, the value of NOW can change from $*\mathsf{NOW}_i$ to $*\mathsf{NOW}_{i+1}$ only by executing algorithm `move_NOW`. At the time of executing

[9]This is the idea behind *Active Logic* (Elgot-Drapkin and Perlis, 1990; Perlis et al., 1991). See Chapter 2 for an overview of that system.

the algorithm, $^*\mathsf{NOW}_i$ is the latest value of NOW, and is, therefore, identical to $^*\mathsf{NOW}$. Step 1 introduces a new interval t. By step 2, $^*\mathsf{NOW} \prec t \in \beta$. By step 3, t becomes the $i+1^{st}$ value of NOW, and is, thus, identical to $^*\mathsf{NOW}_{i+1}$. Therefore, $\beta \vdash {}^*\mathsf{NOW}_i \prec {}^*\mathsf{NOW}_{i+1}$. Since i is arbitrary, then the result applies to all $i \in \mathbb{N}$ $(i > 0)$. $\square$

5.1.3 Types of Intervals

The introduction of NOW induces a partitioning of the set of time intervals. In particular, we need to distinguish between those intervals that are introduced by step (1) of `move_NOW` (or `initialize-NOW`), those introduced by assertions about states holding, and those introduced by assertions about events occurring. There are three main types of intervals.

Definition 5.9 *For every* $t \in \Psi(\mathcal{T})$*:*

1. *t is an* **event interval** *if there is some* $e \in \Psi(\mathcal{E})$ *such that* $\beta \vdash \mathsf{Occurs}(e,t)$*. In that case, we say that t* **is an event interval associated with** *e.*

2. *t is a* **state interval** *if there is some* $s \in \Psi(\text{TEMP}) \cup \Psi(\text{PERM})$ *such that* $\beta \vdash \mathsf{MHolds}(s,t)$*. In that case, we say that t* **is a state interval associated with** *s.*

3. *t is a* **reference interval** *if it is not a state interval or an event interval.*

An event interval is the *unique* interval (modulo Equiv) over which an event occurs. Note that this uniqueness is enforced by **AE1**. A state interval is a maximal stretch of time over which Cassie believes that some state holds. State and event intervals are similar to what are referred to as *temporal traces* in the literature (Krifka, 1989; Link, 1987, for instance). For example, if the state associated with the state interval t is that of the walk-light being green, then, for Cassie to linguistically express t, she would use (variations of) the noun phrase "the time of the walk-light being green".[10] Similarly for event intervals. State and event intervals are, thus, ontologically-dependent on states and events, respectively (Chisholm, 1990; Pianesi and Varzi, 1996a).

A particularly interesting sub-type of event intervals are what will be referred to as *transition intervals*.

[10]Cassie's use of "the time" or "a time" in expressing t depends on whether there is more than one event of s holding.

Definition 5.10 *For every $t \in \Psi(\mathcal{T})$, t is a* **transition interval** *if there is some $\overset{\bullet}{e} \in \Psi(\overset{\bullet}{\mathcal{E}})$ such that* $\mathsf{Occurs}(\overset{\bullet}{e}, t)$.

Referring to the intervals defined thus as *transition* intervals underlines the assumption that I making about punctual events: they are primarily state transitions (see Section 4.2). The punctuality of transitions is captured by the following axiom.

Axiom 5.1 *A transition interval is always atomic.*

Given Theorem 5.2, we have the following piece of notation.

Definition 5.11 *An MTF $\Phi(t)$ is a* **transition-MTF** *if t is a transition interval.*

The role of *state* intervals in the theory is very crucial; they are used to model the persistence of states as NOW moves. To assert that a state, s, holds in the present, two propositions are involved:

1. $\mathsf{MHolds}(s,t)$, where t is a newly-introduced interval, and
2. $*\mathsf{NOW} \sqsubset t$.

As NOW moves, s's persistence is modeled by including each new value of NOW as a sub-interval of t.[11] Of course, this requires that t be associated *only* with the state s. This is captured by the following axiom.

Axiom 5.2 *For every every $s \in \Psi(\text{TEMP}) \cup \Psi(\text{PERM})$ and $i \in \mathbb{N}$ ($i > 0$), if $\beta \vdash \mathsf{Holds}(s, *\mathsf{NOW}_i)$, then there exists a unique state interval, $t \in \Psi(\mathcal{T})$, such that*

1. *$\beta \vdash \mathsf{MHolds}(s,t)$ and $\beta \vdash *\mathsf{NOW}_i \sqsubset t$, and*
2. *for every $s' \in \Psi(\text{TEMP})$, if $\beta \vdash \mathsf{MHolds}(s',t)$, then $s' = s$.*

The uniqueness of the state interval required by the axiom may actually be proved. However, to do that, we will need to introduce more notions that would complicate the exposition and that are only needed for the proof. Therefore, I only sketch an informal proof here. Suppose that there are two distinct intervals, t_1 and t_2, such that they both satisfy (1) in Axiom 5.2. Thus, t_1 and t_2 overlap,

[11] Note that we only need to worry about the persistence of TEMP states; the persistence of PERM states is inevitable, given **APS1**.

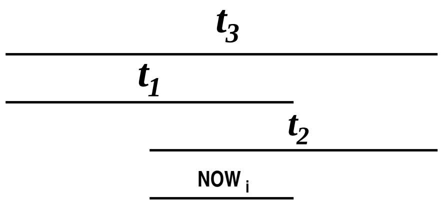

Figure 5.4: An impossible situation: s maximally holds over the overlapping intervals t_1 and t_2.

their common sub-interval being $*\mathsf{NOW}_i$. Without loss of generality, assume that t_1 starts before t_2. The situation is shown in Figure 5.4. Now, consider the interval t_3, the sum of t_1 and t_2. By the cumulativity of states, s holds over t_3. But since both t_1 and t_2 are proper sub-intervals of t_3, then axiom **AS3** is violated. That is, s does not hold maximally over t_1 or t_2.

Reference intervals, on the other hand, are not associated with particular states; they designate temporal perspectives from which Cassie *views* a situation. In particular, reference intervals are used to represent different granularities of the present (more on this in Chapter 7). Among the collection of reference intervals, an important sub-collection is the collection of NOW-**intervals**, that is made up of those intervals that were once the value of NOW. NOW-intervals are intervals representing the present at the finest level of granularity.

Axiom 5.3 *For all* NOW*-intervals, t, the following holds:*

1. *t is a reference interval.*

2. *There is no reference interval, $t' \in \Psi(\mathcal{T})$, such that $t' \sqsubset t$.*

Requiring NOW-intervals to be reference intervals excludes them form being maximal intervals over which a state is asserted to hold. That is, a proposition of the form $\mathsf{MHolds}(s, *\mathsf{NOW}_i)$, though syntactically and semantically valid, is pragmatically not possible. This reflects our discussion above that a super-interval of *NOW is introduced whenever a state is asserted to be holding in the present.

The reference interval pointed to by NOW (i.e., *NOW) is expressible by the English "now". A reference interval may also be given a *value*, "3:45 p.m." for instance. Otherwise, Cassie cannot linguistically-express reference intervals; they only determine the tense and aspect of sentences produced by Cassie (Almeida, 1987; Almeida, 1995). Reference intervals are similar (but not identical) to the reference times of (Reichenbach, 1947), and are not to be confused with the reference intervals of (Allen, 1983).

5.1.4 The Chain of Experience

The following axiom states a principle that follows from the informal assumption that NOW moves whenever there is a change.

Axiom 5.4 (The First Principle of Change) *NOW *is always atomic.*

Informally, suppose that *NOW is not atomic; then there is some t such that $t \sqsubset$*NOW. By Axiom 5.3, t cannot be a reference interval. Therefore, t is a situation interval. By Definition 5.9, there is a state s that maximally holds over t. But if s maximally holds *within* *NOW, then Cassie became aware of changes (s starting and/or ceasing to hold) without NOW moving. Since NOW moves whenever Cassie is aware of a change, then *NOW must be atomic. Note that, *theoretically*, a general NOW-interval need not be atomic. For example, Cassie may be told that some state held within $*\mathrm{NOW}_i$, where $*\mathrm{NOW}_i \prec *\mathrm{NOW}$. Thus, NOW-intervals are atomic so long as they are present (i.e., pointed to by NOW); as they become past, they may no longer be atomic. Nevertheless, note that except for *NOW, which is expressible by "now", one cannot refer to reference intervals in natural language. That is, once they become past, we cannot tell Cassie anything about them. Thus, I will assume as a working hypothesis that *all* NOW-intervals are atomic.[12] This way, we can talk about NOW-**MTF**s. An MTF, $\Phi(*\mathrm{NOW}_i)$, is said to be the i^{th} NOW-MTF. $\Phi(*\mathrm{NOW})$ will be referred to as the **current** MTF. Corresponding to the chain of NOW-intervals, NOW-MTFs form a chain ordered by the "precedes" relation, with the current MTF being the last element in the chain.

The chain of NOW-MTFs is not only distinguished for being linearly ordered; there is a genuine difference between NOW-MTFs and other MTFs in Cassie's mind. NOW-MTFs comprise direct, first-person experiences by Cassie. At least some of the states that hold in a NOW-MTF have been directly perceived by Cassie. Cassie's knowledge of states holding in other MTFs is either the result of inference or direct assertion, but never bodily feedback which may only take place in the present, within some NOW-MTF. As shall be shown later, it is because of this distinction that Cassie may

[12]Even if a NOW-interval is given a value, such as "3:45", we assume that these values are, linguistically, moments of time within which it is not reasonable to assert that something happened.

have a *feel* of the duration of a NOW-MTF (or the smallest element thereof), but only *knowledge* of the duration of a non-NOW-MTF.

NOW moves when, and only when, Cassie becomes aware of a change. There are two comments to make about this assertion. First, not any change moves NOW. For example, if Cassie infers that the walk-light changed from red to green yesterday, NOW shouldn't move. NOW moves when, and only when, some state holding in the current MTF ceases to hold, or some state not holding in the current MTF starts to hold. Note that, in such cases, NOW *must* move to reflect the state's holding, or not holding, being past. Second, the "when" part is sanctioned by the First Principle of Change (Axiom 5.4). The "only when" part is validated by another principle regarding the epistemic distances among NOW-MTFs. First, a definition.

Definition 5.12 *Two* NOW-*MTFs,* $\Phi(*\mathsf{NOW}_i)$ *and* $\Phi(*\mathsf{NOW}_j)$*, are* **epistemically equivalent** *if* $d_e(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_j)) \leq 2$.

What does the epistemic equivalence of $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ imply? It implies that, if distinct, $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ differ only in their smallest elements, their location in time if you will.

Proposition 5.2 *If* $\Phi(*\mathsf{NOW}_i)$ *and* $\Phi(*\mathsf{NOW}_j)$ *are distinct, epistemically-equivalent* NOW-*MTFs, then* $\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j) = \{*\mathsf{NOW}_i, *\mathsf{NOW}_j\}$.

Proof. Since $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ are distinct, then, by Theorem 5.2,

$$\{*\mathsf{NOW}_i, *\mathsf{NOW}_j\} \subseteq \Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j).$$

But since $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ are epistemically-equivalent, then, by Definition 5.6,

$$|\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j)| \leq 2.$$

Thus, by Theorem 5.3,

$$|\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j)| = 2.$$

Therefore, $\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j) = \{*\mathsf{NOW}_i, *\mathsf{NOW}_j\}$. $\square$

Other than their smallest elements, two epistemically-equivalent NOW-MTFs share all their intervals. Most importantly, they share all their state intervals, and thus the states that hold in $\Phi(*\mathsf{NOW}_i)$ are exactly those that hold in $\Phi(*\mathsf{NOW}_j)$. There are two points to note.

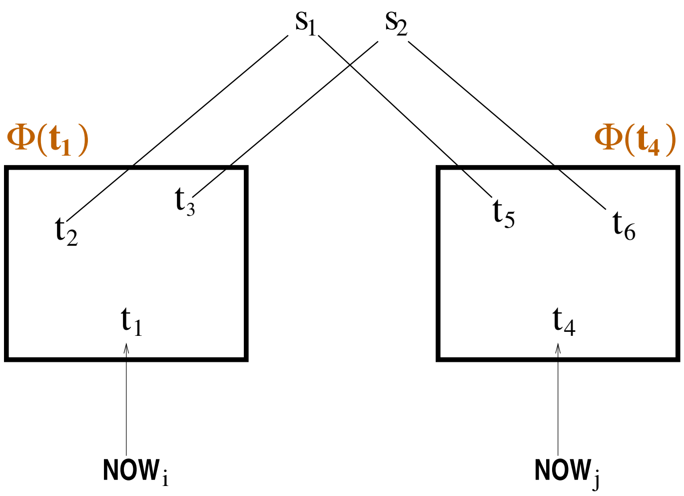

Figure 5.5: Two NOW-MTFs sharing the same set of states but not epistemically equivalent.

1. Epistemic equivalence (formally, $\lambda x \lambda y(d_e(x,y) \leq 2)$) is an equivalence relation. Symmetry stems from the commutativity of d_e (property (c) of d_e in the proof of Observation 5.3). Transitivity is based on noting that $\Phi(*\mathsf{NOW}_i) \setminus \{*\mathsf{NOW}_i\} = \Phi(*\mathsf{NOW}_j) \setminus \{*\mathsf{NOW}_j\}$, which follows from Proposition 5.2. Reflexivity follows from symmetry and transitivity, or could be independently established using property (b) of d_e.

2. As pointed out above, if $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ are epistemically equivalent, then the same collection of states hold in both. The converse is not true, however. That is, two NOW-MTFs may correspond to the same collection of states, yet fail to be epistemically equivalent. The reason is that the above definition is based on MTFs, mere collections of intervals, not states corresponding to MTFs. Thus, what matters is not whether the two MTFs share the same states, but rather, whether they share the same *events* of the same states holding. Figure 5.5 depicts two NOW-MTFs with the same states holding in them, yet with an epistemic distance of 6.

Having mentioned events, we need to consider how an event is distributed over NOW-MTFs. The following proposition establishes a result that, reflecting the supervenience of durative events (see Chapters 3 and 4), allows us to dismiss event intervals as having anything to do with differences among NOW-MTFs; the responsibility for these differences lies squarely on states.

Proposition 5.3 *For every non-transition event interval, t, and $i \in \mathbb{N}$ ($i > 0$), if $t \in \Phi(*\mathsf{NOW}_i)$, then there is some state interval, t', such that $t' \in \Phi(*\mathsf{NOW}_i)$.*

Proof. Since t is a non-transition event interval, then there is some $\overset{\bullet\!-\!\bullet}{e} \in \Psi(\overset{\bullet\!-\!\bullet}{\mathcal{E}})$ such that $\mathsf{Occurs}(\overset{\bullet\!-\!\bullet}{e},t)$. By **TE3** from Section 4.4, there are transition intervals, t_1 and t_2, such that $\beta \vdash \mathsf{Covers}(t,t_1,t_2)$. Again, by **TE3**, there is some $s \in \Psi(\textsc{Temp})$ such that t_1 and t_2 are the times of occurrence of an onset-cessation pair of s. Therefore, there is some state interval, t_3, associated with s such that $t_3 \sqsubseteq t$. By Definition 5.4 (MTFs), it follows that, for every $i \in \mathbb{N}$ ($i > 0$), if $t \in \Phi(*\mathsf{NOW}_i)$, then $t_3 \in \Phi(*\mathsf{NOW}_i)$. □

Given the above result, the intuition is that, if two NOW-MTFs are not epistemically-equivalent, then there must be some state that holds in one and not in the other. To ensure that NOW moves only when there is some state change, I adopt the following principle.

Axiom 5.5 (The Second Principle of Change) *For every $i \in \mathbb{N}$ ($i > 0$), $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_{i+1})$ are not epistemically equivalent.*

Thus, the structure of Cassie's memory is *concise* (Williams, 1986): Cassie's mind is not populated with chains of NOW-MTFs that only differ in their smallest elements. If a state interval is in $\Phi(*\mathsf{NOW}_i) \setminus \Phi(*\mathsf{NOW}_{i+1})$, then a state holding in $\Phi(*\mathsf{NOW}_i)$ has ceased to hold; if it is in $\Phi(*\mathsf{NOW}_{i+1}) \setminus \Phi(*\mathsf{NOW}_i)$, then a state not holding in $\Phi(*\mathsf{NOW}_i)$ has started to hold.

Theoretically, two consecutive NOW-MTFs may be disjoint. Nevertheless, except for the perceptually-crudest of agents, this seems very unlikely. First, this is certainly not the case for humans; our deep sense of the continuity (or density) of time would probably vanish if it weren't for the strong overlap between consecutive NOW-MTFs. If many changes appear to us to happen simultaneously, we would probably not be able to make any sense of the world. Second, as pointed out by many authors (McDermott, 1982; Shoham and McDermott, 1988; Morgenstern, 1996; Shanahan, 1997, to mention a few), most states are persistent, and change is generally an exception. Thus, typically, a state would span more than one MTF. However, this does not rule out cases in which two consecutive NOW-MTFs may have no states in common. For example, consider an incarnation of Cassie where her only task is to monitor the state of some gauges. The readings of gauges are prone to very rapid change, and Cassie only *samples* those readings, say, every five seconds. Naturally, the reading of any gauge would change from one sample to the next. In such a situation, every sample corresponds to a different NOW-MTF, and one may expect consecutive samples to be totally different from one another.

The above principle notwithstanding, it should be noted that the epistemic distance between two NOW-MTFs, $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ does not necessarily increase with their *temporal distance*, $|i-j|$. Figure 5.6 shows an example where $\Phi(t_4)$ separates the epistemically equivalent NOW-MTFs $\Phi(t_1)$ and $\Phi(t_6)$. Intuitively, NOW moves from t_1 to t_4 as a result of s_3 starting to hold. It then moves to t_6 when s_3 ceases to hold. Thus, except for their location in time, the states of affairs corresponding to $\Phi(t_1)$ and $\Phi(t_6)$ are identical. Figure 5.7 shows a situation where the two epistemically equivalent, $\Phi(t_1)$ and $\Phi(t_6)$, are separated by $\Phi(t_4)$ which is missing a state interval, t_2 that belongs to both. Note, however, that such a situation is impossible, for the convexity of intervals necessitates that t_2 be in $\Phi(t_4)$ (see **AT10**). Indeed, one may prove the following result.

Proposition 5.4 *If* $\Phi(*\mathsf{NOW}_i)$ *and* $\Phi(*\mathsf{NOW}_{i+n})$ $(n > 1)$ *are epistemically-equivalent* NOW*-MTFs, then for every m,* $0 < m < n$*,* $\Phi(*\mathsf{NOW}_i) \setminus \{*\mathsf{NOW}_i\} \subset \Phi(*\mathsf{NOW}_{i+m})$*.*

Proof. Since $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_{i+n})$ are epistemically-equivalent, then, by Proposition 5.2, for every t, $t \in \Phi(*\mathsf{NOW}_i) \setminus \{*\mathsf{NOW}_i\}$ implies $t \in \Phi(*\mathsf{NOW}_{i+n})$. Now suppose that $t \in \Phi(*\mathsf{NOW}_i) \setminus \{*\mathsf{NOW}_i\}$ and that it is not the case that $t \in \Phi(*\mathsf{NOW}_{i+m})$, for some m $(0 < m < n)$. Therefore, $*\mathsf{NOW}_i \sqsubset t$ and $*\mathsf{NOW}_{i+n} \sqsubset t$, but $*\mathsf{NOW}_{i+m} \not\sqsubset t$. But, by Theorem 5.4 and the transitivity of $\prec$ (**AT2**), $*\mathsf{NOW}_i \prec *\mathsf{NOW}_{i+m}$ and $*\mathsf{NOW}_{i+m} \prec *\mathsf{NOW}_{i+n}$. By the convexity of t (**AT10**) and **TT2**, $*\mathsf{NOW}_{i+m} \sqsubset t$, which leads to a contradiction. Therefore, $t \in \Phi(*\mathsf{NOW}_{i+m})$, for all m $(0 < m < n)$. Since t is arbitrary, then $\Phi(*\mathsf{NOW}_i) \setminus \{*\mathsf{NOW}_i\} \subset \Phi(*\mathsf{NOW}_{i+m})$, for all m $(0 < m < n)$. □

5.2 The Dynamics of Time

In the previous section, we investigated the meta-theoretical structure of time and outlined general principles that govern the movement of NOW. In this section, we look in more detail at how Cassie's beliefs evolve over time and how this interacts with the dynamics of NOW. Before doing that, however, I need to digress.

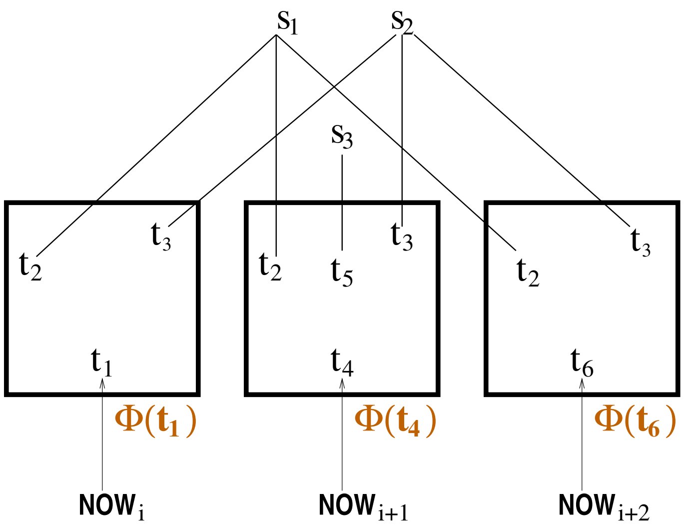

Figure 5.6: MTF $\Phi(t_4)$ separates the epistemically equivalent $\Phi(t_1)$ and $\Phi(t_6)$.

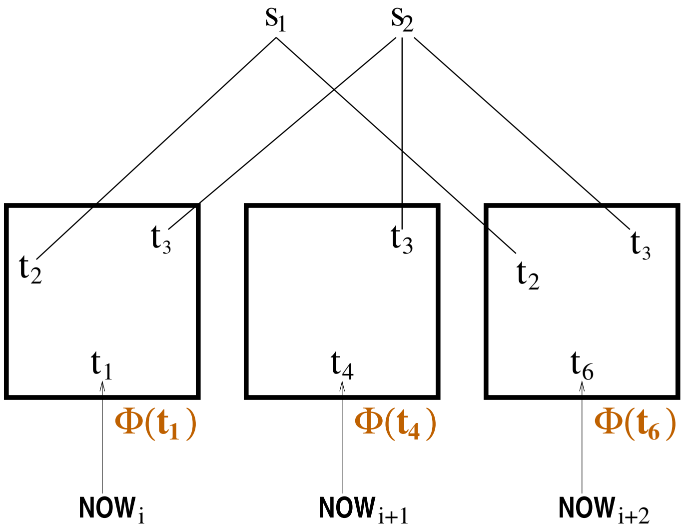

Figure 5.7: An impossible situation; t_2 must be in $\Phi(t_4)$ by the convexity of intervals.

5.2.1 Consistency over Time

As an agent acting and reasoning in a dynamic world, Cassie needs to be capable of handling failure. Failure manifests itself in two ways. First, Cassie may fail to perform some action or achieve some goal. This issue is investigated elsewhere (Ismail and Shapiro, 2000a) and will not be discussed in any detail here. Second, which is what concerns us, Cassie may fail to *reason* correctly about the domain. This happens when Cassie's belief space, β, is inconsistent, i.e., when it contains contradictory beliefs. In the theory presented here, there will be times when Cassie would make default assumptions about the persistence of states (see, in particular, Sections 7.5 and 8.1). It is possible that these assumptions are simply false, and should Cassie be aware of that, β would be inconsistent. Belief revision within the SNePS tradition has come a long way (Martins and Shapiro, 1988; Johnson and Shapiro, 2000a; Johnson and Shapiro, 2000b; Johnson and Shapiro, 2000c), and inconsistency may, in fact, be handled appropriately. However, belief revision in a theory like the one presented here is more complicated. In particular, certain meta-theoretical constraints are imposed on the representation of time and its progression (for example, see Axioms 5.2, 5.3, and 5.4). Whatever belief revision might do to resolve a contradiction, it should do it while observing these constraints. Currently, the belief revision system is not thus integrated with the theory of time. Therefore, should β become inconsistent, we would not be able to verify that the principles and axioms constraining the theory are not violated. It should be noted, however, that this would only be the case if the inconsistency involves beliefs about time, and, more specifically, if it involves beliefs about states holding in the present. Therefore, in what follows, I will make the assumption that beliefs about current states are never contradicted. More precisely, I will assume that, at no time, do pairs of propositions of the form $\mathsf{Holds}(s,{*}\mathsf{NOW})$ and $\neg\mathsf{Holds}(s,{*}\mathsf{NOW})$, explicitly or implicitly, co-exist in β. This would allow us to present results that we would not otherwise be able to formally prove unless a complete theory of temporal belief revision is presented. Note that these results are not *wrong*; it is just that some of them presume consistency.

Another kind of inconsistency, one that the theory tolerates and, indeed, endorses, exists at the non-logical, meta-theoretical level. This kind of inconsistency is symptomatic of the need to move NOW. It involves violations of the principles and constraints of the theory that are, not the result of inconsistencies in β, but a natural side-effect of the progression of time. For example, Theorem 8.3 in Section 8.2.3 states that whenever Cassie is perceiving some state holding, then she believes

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t)\}$.
3. `move_NOW`.
4. $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \sqsubset t\}$.

Figure 5.8: What happens when Cassie senses that the state *s* holds.

that it does. Naturally, there will be short periods of time when this theorem does not apply. In particular, those are the times it takes to update Cassie's belief space (and other components of the system) in order to reflect the newly-perceived information. I will, therefore, assume that those are 0-periods—that it takes no time, at the theoretical level, to perform any required updates and *re-stabilize* the system.

5.2.2 A Note on the Frame Problem

Algorithm `move_NOW` outlined in Section 5.1.2 merely represents how temporal progression is modeled. However, it does not express everything that happens when NOW moves. Whenever NOW moves, a new MTF is constructed, namely $\Phi(*\mathsf{NOW})$. What do we know about this MTF, and how can it be constructed? Suppose the new MTF is the i^{th} NOW-MTF. One thing we learned in Section 5.1.4 is that $\Phi(*\mathsf{NOW}_i)$ is not epistemically equivalent to $\Phi(*\mathsf{NOW}_{i-1})$. If NOW moves due to some state's ceasing, then the state interval, in $\Phi(*\mathsf{NOW}_{i-1})$, associated with that state should not be in $\Phi(*\mathsf{NOW}_i)$. On the other hand, if NOW moves due to some state's starting to hold, then a new state interval associated with that state should be in $\Phi(*\mathsf{NOW}_i)$. The algorithm in Figure 5.8 outlines what happens when Cassie senses that some state, *s*, holds. Cessation will be discussed below. But, in this section, I am concerned with a slightly different issue.

It is often the case that many states holding in $\Phi(*\mathsf{NOW}_{i-1})$ continue to hold in $\Phi(*\mathsf{NOW}_i)$, and there should be some way of incorporating them into the new MTF. Of course what is lurking between the lines here is the frame problem (McCarthy and Hayes, 1969). I will not delve into a long discussion of, nor propose any ingenious solutions to, the notorious problem; the literature is full of such discussions and proposals (Shoham and McDermott, 1988; Kautz, 1986; Shoham, 1986; Pylyshyn, 1986; Brown, 1987; Reiter, 1991; Morgenstern, 1996; Shanahan, 1997, to mention a few). Rather, I am going to make some informal remarks. Within the theoretical framework that has been developed in the previous sections, the frame problem manifests itself in determining which states that held in $\Phi(*\mathsf{NOW}_{i-1})$ continue to hold in the new $\Phi(*\mathsf{NOW}_i)$. First of all, I do not believe that there is a single, albeit elusive, solution to the problem; depending on the type of state, there may be different ways of determining whether it holds in $\Phi(*\mathsf{NOW}_i)$.

Armed with a theory of belief revision (Martins and Shapiro, 1988; Johnson and Shapiro, 2000c), we may adopt an off-the-shelf monotonic approach to solving the frame problem (Reiter, 1991; Thielscher, 1999, for example). These approaches provide logical solutions to the problem. That is, they involve Cassie's *reasoning* about whether a certain state continues to hold. In many cases, this is reasonable. In particular, the frame problem is usually discussed in one of two contexts:

1. the context of planning, where an agent needs to predict the state of the environment following any of its planned actions, and

2. the context of reasoning about (what have confusingly come to be known as) narratives. Basically, a narrative is a world description in the form of a sequence of events and some states that hold at various points. Given a narrative, a reasoning system is asked to make predictions about which states hold (or not) at various points in the history of the world (the notorious Yale Shooting scenario is a typical example (Hanks and McDermott, 1987)).

In such settings, abstract reasoning is the only way to account for the persistence of states, and almost all the work that has been done on the frame problem is concerned with the kinds of axioms and reasoning systems required for a robust and efficient account of persistence that allows only *natural* predictions. The situation that we have here is neither one of planning nor of reasoning about narratives. We have Cassie out there, in the world, reasoning, acting, and perceiving the environment while maintaining a sense of time. I am not primarily interested in reasoning about past or future states. (This is not to say that these are trivial or unimportant issues; they certainly are not.) Rather, my main concern is to naturally account for Cassie's awareness of *present* states. Granted, reasoning is still needed for this task—for projecting states from $\Phi(*\mathsf{NOW}_{i-1})$ onto the new MTF $\Phi(*\mathsf{NOW}_i)$. Nevertheless, there is at least a subset of the states holding in $\Phi(*\mathsf{NOW}_{i-1})$ that should be incorporated in the new MTF without any reasoning on the part of Cassie. Let us motivate this with an example. The following (successor state) axiom appears in (Reiter, 1998, p.

553):

$$going(l,l',do(a,s)) \equiv (\exists t)a = startGo(l,l',t) \vee going(l,l',s) \wedge \neg(\exists t)a = endGo(l,l',t).$$

The axiom states conditions under which it could be concluded that the state of going from one location to another holds in a particular situation. Now, suppose that this is the *current* situation (or, in our, terminology, the current NOW-MTF). The crucial question here is *who is going?* There is some agent that does not appear in the axiom but is certainly implied. The axiom itself represents some useful piece of knowledge that an agent, Cassie for example, may use to reason about states of "going". But using this axiom to conclude that some agent is *still* "going" makes sense only when the implicit agent is *not* Cassie, the reasoning agent. If Cassie is the one who is "going", then she can conclude that she is still "going", not because of her general knowledge of causes and effects, nor because it is reasonable to *assume* so, but because *she* is actually "going". Agents do not need to reason about the *current* states of their own bodies, or those of the perceived environment, they have first-person direct access to those states through perception and proprioception.

What I am suggesting here is that, for states of Cassie's own body, or of the perceived environment, whether they hold in the new MTF should be taken care of, not at the KL, but at the PML, as part of temporal progression routines. How may that be done? This is the topic of the next section.

5.2.3 Modalities

To account for the persistence of bodily or perceived states, a set, $\mathcal{M}$, of meta-logical variables is employed, corresponding to various agent modalities. (As shall be seen below, I interpret "modality" in a very broad sense.) The set $\mathcal{M}$ is partitioned into two sets: $\mathcal{M}_{\text{prop}}$ for proprioception, and $\mathcal{M}_{\text{per}}$ for perception. Each variable in $\mathcal{M}_{\text{prop}}$ contains a proposition representing the state currently occupying the corresponding modality. On the other hand, $\mathcal{M}_{\text{per}}$ represents what each modality conveys to Cassie about the external world. For example, a particular incarnation of Cassie may have the following set of proprioception modality variables:

$$\mathcal{M}_{\text{prop}} = \{\mathsf{VISION}_{\text{prop}}, \mathsf{AUDITION}_{\text{prop}}, \mathsf{LOCOMOTION}_{\text{prop}}, \mathsf{HANDLING}_{\text{prop}}\}$$

Such a Cassie has visual and auditory capabilities in addition to, maybe, wheels for movement and an arm for manipulating objects. Different processes performed by Cassie, or states of her body, occupy different modalities. For example, $\mathsf{VISION}_{\text{prop}}$ may contain the proposition "I am looking toward the green light" or "My eyes are shut" (where the first person refers to Cassie), $\mathsf{LOCOMOTION}_{\text{prop}}$ may contain the proposition "I am moving from l to l'" or "I am standing still", $\mathsf{AUDITION}_{\text{prop}}$ may contain "I am talking to Stu", and so on. Thus, $\mathcal{M}_{\text{prop}}$ represents what each modality is being *used for*.

The same Cassie may have a similar set of perception modality variables:[13]

$$\mathcal{M}_{\text{per}} = \{\mathsf{VISION}_{\text{per}}, \mathsf{AUDITION}_{\text{per}}\}$$

For example, $\mathsf{VISION}_{\text{per}}$ may contain the proposition "The green light is on" or "The block is on the table" and $\mathsf{AUDITION}_{\text{per}}$ may contain the proposition "The alarm is sounding" or "The radio is on".

In general, modality variables point to elements of β. If the modality corresponding to the variable $\mu \in \mathcal{M}_{\text{prop}}$ is occupied by some state, $[\![s]\!]$, then $*\mu = \mathsf{MHolds}(s,t)$, where t is the state interval associated with s, such that $^*\mathsf{NOW} \sqsubset t$. A similar idea underlies perception modality variables with the following provision. One main difference between perception and proprioception modality variables is that the former, but not the latter, take as values *sets* of propositions. Thus, for every $\mu \in \mathcal{M}_{\text{per}}$, there is a *set* of states and associated state intervals such that $*\mu = \{\mathsf{MHolds}(s_i,t_i) : {}^*\mathsf{NOW} \sqsubset t_i\}$. For example, Cassie not only sees that the block is on the table, but also that it is, say, red, and that she is close enough to pick it up. Thus, in general, a single perception modality may simultaneously convey more than one piece of information about the external world. On the other hand, a proprioception modality may be occupied by only *one* bodily state or process. This idea is stated precisely by the following axiom.[14]

Axiom 5.6 *For every $i \in \mathbb{N}$ ($i > 0$), and $\mu \in \mathcal{M}_{\text{prop}}$, at $[\![{}^*\mathsf{NOW}_i]\!]$, there is one, and only one, $s \in$*TEMP*, such that $[\![s]\!]$ occupies the modality corresponding to μ.*

In the above statement, $[\![{}^*\mathsf{NOW}_i]\!]$ refers to the interval of *real* time during which the value of NOW is $^*\mathsf{NOW}_i$. The important thing to note about this axiom is that it does *not* state that, at any time, a proprioception modality *variable* is occupied by a unique $\mathcal{FOCS}$-term; this is something that the theory has to ensure. The axiom merely states the *semantic* counterpart of this assertion. Note

[13]Typically, there are fewer perception modalities than there are proprioception modalities. For example, the locomotion system does not convey anything about the external environment, only its own state, and likewise for the handling system (unless we consider tactile perception).

[14]Note that I am, reasonably, assuming that bodily and perceived states may only be temporary.

also that, not only do we require any proprioceptual modality to be occupied by a single state, but that there is *always* some state occupying any given proprioceptual modality. This might seem strange given that there may be times when Cassie is not *using* some, or all, of her modalities. However, as mentioned above, I interpret the notion of modality in a very broad sense; passive states such as "I am standing still" or "My eyes are shut" *occupy* modalities in our theory (in this case, $\mathsf{LOCOMOTION}_{\mathrm{prop}}$ and $\mathsf{VISION}_{\mathrm{prop}}$, respectively). Thus, even if Cassie is not using its sensors and effectors (in the traditional sense of "using"), states corresponding to these resources being idle (or available) are legitimate modality occupiers.

In what follows, I make a couple of assumptions about modalities and states that occupy, or are perceived through, them. First, I assume that, for every bodily state, there is a unique set of proprioception modalities that it occupies if and when it holds. Formally, there is a function, $\mathrm{Mod}_{\mathrm{prop}}$, from TEMP to the power set of $\mathcal{M}_{\mathrm{prop}}$ mapping each state to the proprioception modality variables corresponding to modalities it occupies when it holds. Note that, for non-bodily states, the value of $\mathrm{Mod}_{\mathrm{prop}}$ is simply the empty set. Thus, bodily states may be formally identified as those members of TEMP for which the value of $\mathrm{Mod}_{\mathrm{prop}}$ is a non-empty set.

An important property of $\mathrm{Mod}_{\mathrm{prop}}$ is stated by the following axiom.

Axiom 5.7 *For every $s \in$ TEMP, $\mu \in \mathcal{M}_{\mathrm{prop}}$, and $i \in \mathbb{N}$ ($i > 0$), at $[\![{*}\mathsf{NOW}_i]\!]$, if $[\![s]\!]$ occupies the modality corresponding to μ, then $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$ and, for every $\mu' \in \mathrm{Mod}_{\mathrm{prop}}(s)$, $[\![s]\!]$ occupies the modality corresponding to μ'.*

For example, if $\mathrm{Mod}_{\mathrm{prop}}(s) = \{\mathsf{VISION}_{\mathrm{prop}}, \mathsf{LOCOMOTION}_{\mathrm{prop}}\}$, then, whenever s holds, it occupies *both*, and *only*, $\mathsf{VISION}_{\mathrm{prop}}$ and $\mathsf{LOCOMOTION}_{\mathrm{prop}}$. The following axiom is also needed to reflect, at the KL, what Axioms 5.6 and 5.7 require at the PML.

Axiom 5.8 *For every $s_1, s_2 \in$ TEMP, if $s_1 \neq s_2$ and $\mathrm{Mod}_{\mathrm{prop}}(s_1) \cap \mathrm{Mod}_{\mathrm{prop}}(s_2) \neq \{\}$, then, for all $t \in \Psi(\mathcal{T})$, if $\beta \vdash \mathsf{Holds}(s_1, t)$, then $\beta \vdash \mathsf{Holds}(\neg s_2, t)$.*

Basically, Axiom 5.8 is a constraint on the domain theory, a principle that the knowledge engineer should adopt in axiomatizing the domain.

Second, I assume that perception modalities are mutually exclusive. That is, the same state cannot be perceived via two distinct perception modalities; each modality presents a unique perspective of the external environment. Note, however, that if two distinct states, s and s', are perceived via two distinct perception modalities, μ_{per} and μ'_{per}, it could still be the case that the fact that s holds entails that s' holds (or vice versa), or that both states entail that some third state, s'', holds. For example, looking at the alarm clock, Cassie may *see* that it is 7 a.m. At the same time, she may also *hear* the alarm, which, given other background knowledge, would entail what she visually perceives. The main point is that, if a state is perceivable, then it is perceivable via one, and only one, modality.[15] Formally, we introduce a *partial* function, Mod_{per}, from TEMP to $\mathcal{M}_{\text{per}}$, such that $[\![s]\!]$ may only be perceived via the modality corresponding to $\text{Mod}_{\text{per}}(s)$.

Axiom 5.9 *For every $s \in$ TEMP, $\mu \in \mathcal{M}_{\text{per}}$, and $i \in \mathbb{N}$ $(i > 0)$, at $[\![*\mathsf{NOW}_i]\!]$, if $[\![s]\!]$ is perceived via the modality corresponding to μ, then $\mu = \text{Mod}_{\text{per}}(s)$.*

In addition, I assume that bodily states cannot be perceived. That is, the set of bodily states and the set of perceivable states are disjoint.

Axiom 5.10 *For every $s \in$ TEMP, if $\text{Mod}_{\text{prop}}(s) \neq \{\}$, then $\text{Mod}_{\text{per}}(s)$ is undefined.*

Modality variables are set at the PML when bodily or perceived states start or cease to hold. Thus, in the algorithm of Figure 5.8, if s is one of these states, step (2) is followed by pointing all the appropriate modality variables (those corresponding to modalities occupied by, or perceiving, s) to the new proposition $\mathsf{MHolds}(s,t)$. Constructing the new MTF may thus be characterized by the algorithm in Figure 5.9. The crucial thing here is that currently holding bodily or perceived states are smoothly incorporated in the new MTF without the need for any reasoning, reflecting Cassie's continuous sense of the states of her body and the perceived environment. Of course, the algorithm in Figure 5.9 presupposes the correct setting of modality variables. As pointed out above, this should be taken care of by a revised version of the algorithm in Figure 5.8 (and a similar account for cessation). To proceed and introduce such a revision, however, we need to carefully examine the different issues involved in state transitions.

5.2.4 Transitions

Before delving into what exactly is involved in a transition from one value of NOW to another (or from one NOW-MTF to the next), we need to be precise and explicit about the technical use of

[15] Note that, ultimately, this might just be a simplifying assumption of the theory.

Algorithm `setup_new_MTF`

1. `move_NOW`
2. For all $\mu \in \mathcal{M}_{\text{prop}}$
    3. If there are s and t such that $*\mu = \mathsf{MHolds}(s,t)$ then $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \sqsubset t\}$.
4. For all $\mu \in \mathcal{M}_{\text{per}}$
    5. For all s and t such that $\mathsf{MHolds}(s,t) \in *\mu$
        6. $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \sqsubset t\}$.

Figure 5.9: Algorithm `setup_new_MTF`.

some English expressions that shall recur henceforth. These are expressions that allude to Cassie's knowledge of state transitions over time.

Definition 5.13 *For every $s \in$ TEMP$\cup$PERM and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that** s **holds at** $*\mathsf{NOW}_i$ *if, at* $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$ *and, at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.

Typically, Cassie determines that s holds at $*\mathsf{NOW}_i$ if the transition from $*\mathsf{NOW}_{i-1}$ to $*\mathsf{NOW}_i$ involves Cassie's coming to believe that s holds. For example, suppose that s is the state of the walk-light being on. Further, suppose that the time is $[\![*\mathsf{NOW}_{i-1}]\!]$ and that Cassie sees (and, therefore, believes) that the walk-light is *not* on. In this case, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.[16] Now, if the walk-light turns on, then NOW moves to $*\mathsf{NOW}_i$ and $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$. This situation can be described by saying that Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$ (see Figure 5.10).

The reader should be careful about the intuition behind the above definition. In particular, note that the qualifications "at $[\![*\mathsf{NOW}_i]\!]$" and "at $[\![*\mathsf{NOW}_{i-1}]\!]$" are very crucial and cannot be dropped. Consider the situation depicted in Figure 5.11. At $[\![*\mathsf{NOW}_{i-1}]\!]$ Cassie is *not* looking toward the walk-light and, thus, does not know whether it is on (i.e., at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$). Now, Cassie turns to the walk-light, and two things happen: NOW moves to $*\mathsf{NOW}_i$ (due to the change in Cassie's orientation), and Cassie sees that the walk-light is on. In this case, as in the above one, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$, and we can describe this situation by saying that Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$. Unlike the above situation, however, Cassie may

[16]In fact, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. I am assuming consistency of β throughout the discussion.

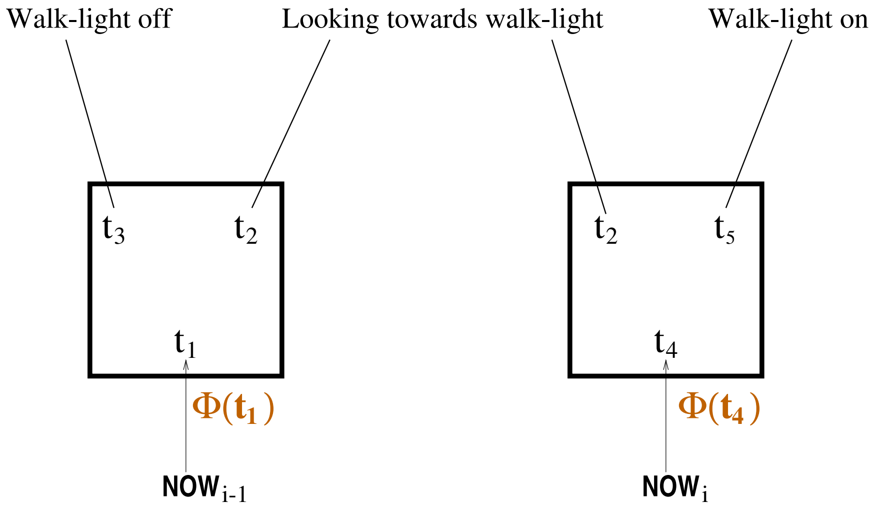

Figure 5.10: Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$.

later come to believe that the walk-light was actually on at $[\![*\mathsf{NOW}_{i-1}]\!]$, when she was not looking. This does not change anything; it is still the case that Cassie determines that the walk-light is on at $[\![*\mathsf{NOW}_i]\!]$, according to Definition 5.13. The important thing is not that $\beta \nvdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$, but that this is the case at $[\![*\mathsf{NOW}_{i-1}]\!]$, the time at which $*\mathsf{NOW}_{i-1}$ is itself $*\mathsf{NOW}$.

To distinguish the above two scenarios, we need a couple of definitions that are more specific than Definition 5.13

Definition 5.14 *For every $s \in$* TEMP∪PERM *and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that s starts to hold at** $*\mathsf{NOW}_i$ *if, at* $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s, *\mathsf{NOW}_i)$ *and, at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$.

The above definition applies to the situation in Figure 5.10, but not that in Figure 5.11. The following definition singles out the situation in Figure 5.11.

Definition 5.15 *For every $s \in$* TEMP∪PERM *and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that s persists at** $*\mathsf{NOW}_i$ *if, at* $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s, *\mathsf{NOW}_i)$ *and, at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$ *and* $\beta \nvdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$.[17]

The use of "persists" in the above definition will be clear below (Section 7.5.1). The intuition is that, when one observes a state holding (not starting to hold), one assumes that it *was* holding and will

[17] *Note that this definition may, in fact, be generalized to all sorts of states, not just* TEMP *and* PERM.

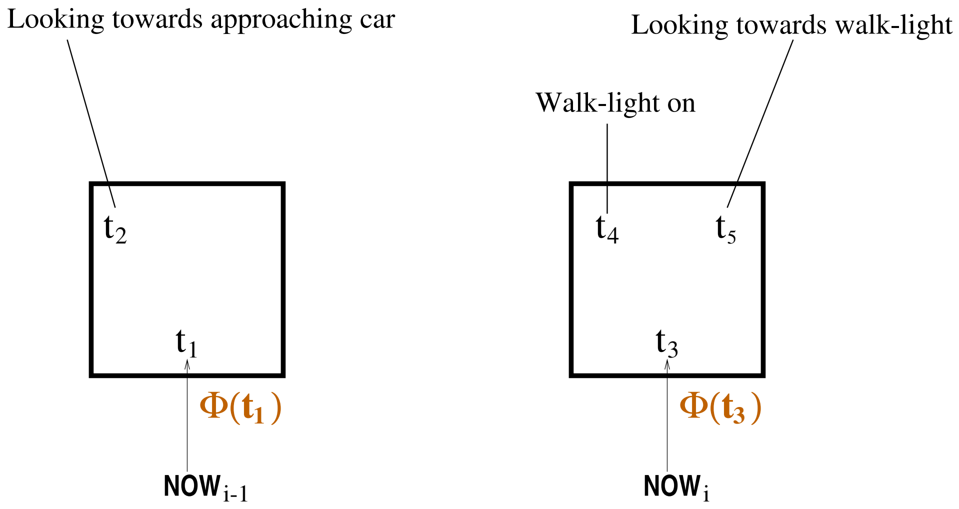

Figure 5.11: Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$. She does not know whether the walk-light is on at $[\![*\mathsf{NOW}_{i-1}]\!]$

continue to hold for a while. That is, one assumes that the state persists (cf. (McDermott, 1982, pp. 122–123)). Note, however, that the scenario of Figure 5.11 presents a peculiar instance of Cassie's determining the persistence of a state, namely an instance in which she *starts to perceive* the state.

Definition 5.16 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ $(i > 1)$*, we say that* **Cassie starts to perceive** s **at** $*\mathsf{NOW}_i$ *if*

1. *at* $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ *is perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$,

2. *at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $[\![s]\!]$ *is not perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$*, and*

3. *at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \not\vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.

The above definition is required since the case of starting to perceive a state has a particular significance for the state transition algorithms to be presented in Section 5.2.6. In particular, it involves updating elements of $\mathcal{M}_{\mathrm{per}}$. The first two conditions above capture the transition from not perceiving to perceiving the state s. The third condition merely rules out the possibility of Cassie's determining that s starts through perception, rather than merely starting to perceive it. One important difference between the two cases is that Cassie may indeed believe that s holds over $*\mathsf{NOW}_{i-1}$, for example by being told so. In such a case the transition described by Definition 5.16 is only one of the state of Cassie's perceptual modalities.

Definitions 5.14 and 5.15 describe the different situations in which Cassie determines that a state, s, holds. These may be viewed as involving a transition that Cassie undergoes from one belief state to another. In particular, with respect to s, Cassie undergoes a transition from some sort of a negative belief state to a positive one. More specifically, note that what is common among Definitions 5.13, 5.14, and 5.15, is that Cassie explicitly believes that s holds at the *later* value of NOW. Intuitively, there are similar situations where the opposite is the case.

Definition 5.17 *For every state* $s \in$ TEMP *and for every* $i \in \mathbb{N}$ $(i > 1)$*, we say that* **Cassie determines that** s **ceases to hold at** $^*\mathsf{NOW}_i$ *if, at* $[\![^*\mathsf{NOW}_i]\!]$, $\beta \vdash \neg\mathsf{Holds}(s, ^*\mathsf{NOW}_i)$ *and, at* $[\![^*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \mathsf{Holds}(s, ^*\mathsf{NOW}_{i-1})$.[18]

For example, in the situation represented by Figure 5.10, Cassie determines that the state of the walk-light's being off ceases at $^*\mathsf{NOW}_i$. Similarly, we can give a definition similar to Definition 5.15. This would capture a situation in which Cassie suddenly moves from having an explicit belief that s holds to having no way of telling whether it does. For reasons to be discussed in Section 8.1, I do not believe that this is a very realistic situation. But even if it is, such a situation does not have any major role to play in the theory and, thus, there is no need to introduce an expression describing it. But a situation similar to that embodied in Definition 5.16 is indeed significant.

Definition 5.18 *For every* $s \in$ TEMP *and* $i \in \mathbb{N}$ $(i > 1)$*, we say that* **Cassie ceases to perceive** s **at** $^*\mathsf{NOW}_i$ *if*

1. *at* $[\![^*\mathsf{NOW}_{i-1}]\!]$, $[\![s]\!]$ *is perceived via the modality corresponding to* $\mathrm{Mod}_{per}(s)$,

2. *at* $[\![^*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ *is not perceived via the modality corresponding to* $\mathrm{Mod}_{per}(s)$, *and*

3. *at* $[\![^*\mathsf{NOW}_i]\!]$, $\beta \nvdash \neg\mathsf{Holds}(s, ^*\mathsf{NOW}_i)$.

In the situation depicted in Figure 5.11, Cassie ceases to perceive that a car is approaching at $^*\mathsf{NOW}_i$. Note that this does not mean that Cassie ceases to *believe* that the car is approaching; she only ceases to perceive it, but there might be other reasons for her to believe that it is still there. Indeed, the third condition in Definition 5.18 allows this possibility.

As pointed out in Section 1.5.4, Cassie may determine that a state holds through various means: perception, proprioception, direct assertion, or inference. We should now define how the theory

[18] *Note that, in this case, s must be a* TEMP *state, since* PERM *states do not cease.*

reflects each of these modes of determination. First, let us define the following convenient notion of *directly* determining that a state holds.

Definition 5.19 *For every* $s \in$TEMP$\cup$PERM *and* $i \in \mathbb{N}$ *(*$i > 1$*), we say that* **Cassie directly determines that** s **holds at** $^*\mathsf{NOW}_i$ *if*

1. *Cassie determines that* s *holds at* $^*\mathsf{NOW}_i$*, and*

2. *there is some* $P \subset$ ETERNAL *such that*

   (a) *at* $[\![^*\mathsf{NOW}_{i-1}]\!]$*, for every* $p \in P$*,* $\beta \not\vdash p$*,*

   (b) *at* $[\![^*\mathsf{NOW}_i]\!]$*,* $P \subseteq \beta$*, and*

   (c) $P \vdash \mathsf{Holds}(s, {}^*\mathsf{NOW}_i)$.

Thus, Cassie directly determines that s holds if she acquires some *new* pieces of information which, by themselves, allow her to infer that s holds. That is, Cassie does not need to use any background knowledge to determines that s holds; the new information is sufficient. Note that this means that, at a time when β is empty, Cassie can determine that s holds only directly.

Given this definition, we may now be more precise about what it means to determine that a state holds through perception, proprioception, or direct assertion.

Definition 5.20 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ *(*$i > 1$*), we say that* **Cassie determines that** s **holds at** $^*\mathsf{NOW}_i$ **through perception** *if*

1. *Cassie directly determines that* s *holds at* $^*\mathsf{NOW}_i$ *and*

2. *at* $[\![^*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ *is perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$.

Definition 5.21 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ *(*$i > 1$*), we say that* **Cassie determines that** s **holds at** $^*\mathsf{NOW}_i$ **through proprioception** *if*

1. *Cassie directly determines that* s *holds at* $^*\mathsf{NOW}_i$ *and*

2. *at* $[\![^*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ *occupies the modalities corresponding to elements of* $\mathrm{Mod}_{\mathrm{prop}}(s)$.

Definition 5.22 *For every* $s \in$TEMP$\cup$PERM *and* $i \in \mathbb{N}$ *(*$i > 1$*), we say that* **Cassie determines that** s **holds at** $^*\mathsf{NOW}_i$ **through direct assertion** *if*

1. *Cassie directly determines that s holds at* $^{*}\mathsf{NOW}_i$,

2. *at* $[\![{}^{*}\mathsf{NOW}_i]\!]$, $[\![s]\!]$ *is not perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$*, and*

3. *at* $[\![{}^{*}\mathsf{NOW}_i]\!]$, $[\![s]\!]$ *does not occupy any of the modalities corresponding to elements of* $\mathrm{Mod}_{\mathrm{prop}}(s)$.

Note that, given the above definitions, direct assertion is simply the direct determination of a state holding by means other than perception and proprioception. A similar relation holds between determination through inference and direct determination.

Definition 5.23 *For every* $s \in$ TEMP$\cup$PERM *and* $i \in \mathbb{N}$ $(i > 1)$*, we say that* **Cassie determines that** s **holds at** $^{*}\mathsf{NOW}_i$ **through inference** *if*

1. *Cassie determines that s holds at* $^{*}\mathsf{NOW}_i$ *and*

2. *Cassie does not directly determine that s holds at* $^{*}\mathsf{NOW}_i$.

The crucial difference between this and Definition 5.19 is that the new information is not sufficient for Cassie to determine that s holds; she must also make use of *old* background information. Similar to Definitions 5.20 through 5.23, we can provide definitions for determining persistence, onsets, and cessations.

5.2.5 The Many Scenarios of Change

In this section, some insights into the nature of state change are outlined. This is done through a careful examination of the different situations in which Cassie determines that a change has taken place. Those insights will be presented as a set of principles amending the two principles of change of Section 5.1.4 (Axioms 5.4 and 5.5). It should be noted, however, that these principles are not part of the formal theory (and, hence, will not be presented as axioms); they are general guidelines that will be formally crystallized in a number of algorithms to be presented in the next section. The first of these principles follows directly from Axiom 5.6.

The Third Principle of Change. *For every bodily state* $[\![s]\!]$, $[\![s]\!]$ *ceases to hold when and only when some other bodily state,* $[\![s']\!]$*, starts to hold.*

The above principle underlines our previous discussion of how it is that there is always some bodily state occupying a given proprioceptual modality. For example, suppose that Cassie is moving, and that this occupies the $\mathsf{LOCOMOTION}_{\mathrm{prop}}$ modality. If Cassie stops moving, the

$\mathsf{LOCOMOTION}_{\mathrm{prop}}$ modality gets occupied by the state of her standing still. Similarly, if $\mathsf{VISION}_{\mathrm{prop}}$ is occupied by the state of Cassie's looking toward the walk-light, a cessation of the aforementioned state results in $\mathsf{VISION}_{\mathrm{prop}}$ being occupied by some other state, say Cassie's looking toward an approaching car. The important thing here is that, for bodily states, we have a single unified account for onsets and cessations, since they always occur together. That is, we only need to consider what happens when a bodily state starts to hold; a bodily state's ceasing is only one of the things that happen as another bodily state starts (or the other way around). This shall be reflected in the construction of state transition algorithms.

The Fourth Principle of Change. *For every $s \in$*TEMP *and for every $i \in \mathbb{N}$ ($i > 1$), if Cassie determines that s ceases to hold at $^*\mathsf{NOW}_i$, then this happens through inference or direct assertion.*

The basic insight behind the above principle is that one only perceives or feels (through proprioception) the presence of a state, not its absence. Our bodies always provide *positive* information about states, never negative ones. Thus, one does not directly perceive that the walk-light is not on; rather, one perceives that it is off, and *infers* that it is not on. Alternately, one may come to believe that the walk-light is not on if one is explicitly told so. According to Definition 5.17, determining cessation of a state hinges on coming to believe that it does not hold, which, as I noted, cannot happen through perception or proprioception. What this means is that we only need to account for cessation in cases of inference and direct assertion. Again, this will be embodied in state transition algorithms.

The Fifth Principle of Change. *For every $i \in \mathbb{N}$ ($i > 1$), Cassie may determine that more than one state holds at $^*\mathsf{NOW}_i$.*

The above principle embodies our discussion in Section 5.1.4 that, though generally stable, the world might present perceptually-crude agents with multiple simultaneous changes. But even for agents with fine-grained perception of time (humans, for example) the above principle still holds. In particular, it is often characteristic of the patterns of change exhibited by bodily states. As per the third principle of change, the cessation of a bodily state s is (simultaneous to) the onset of some other bodily state s'. If s occupies more than one modality, then its cessation would typically correspond to the onset of multiple states, each occupying one of those modalities. That is, as a bodily state ceases, multiple states rush in, occupying the modalities just made available. For example, let s be the state of Cassie's searching for a block. In a possible implementation, we have:

$$*\mathsf{VISION}_{\text{prop}} = *\mathsf{LOCOMOTION}_{\text{prop}} = \text{“I am searching for a block”}$$

As s ceases, we get the following situation.

$$*\mathsf{VISION}_{\text{prop}} = \text{“I am looking at a block”}$$

$$*\mathsf{LOCOMOTION}_{\text{prop}} = \text{“I am standing still”}$$

These two states (looking at a block and standing still) would virtually start simultaneously as the searching activity comes to an end.

The Sixth Principle of Change. *For every bodily state s, and for every* $i \in \mathbb{N}$ $(i > 1)$*, if Cassie determines that s holds at* $*\mathsf{NOW}_i$ *through proprioception, then Cassie determines that s starts to hold at* $*\mathsf{NOW}_i$.

The above principle simply states that the only possible situation in which Cassie determines that a state holds through proprioception is when she determines that it *starts* to hold. In particular, Cassie cannot proprioceptually determine that a state persists (as per Definition 5.15). This is reasonable since Cassie does not, all of a sudden, sense a bodily state holding; she has to, first, sense it starting to hold, given the continuous proprioceptual feedback. Thus, determining that a state persists is exclusive to non-bodily states—those that can be *discovered* in the midst of holding. In fact, the above principle embodies a major property that distinguishes bodily states from non-bodily states.[19]

But note the wording of the above principle. I was careful enough to require that the principle applies only if a bodily state is determined to be holding through *proprioception*. Is there another way by which an agent may determine that a state of its own body holds? Theoretically, the agent may be *told* anything, including information about its bodily states. For humans, such information would only be accepted if it conforms with proprioceptual information (at least in normal situations). For robots, however, I believe that this is an implementation decision. One may design a robot that accepts information about its own body from outsiders. Alternatively, the robot may simply reject any such information or choose some intermediate route, following the lead of humans. I will discuss what such design decisions may entail in Section 8.2.2. For now, however, I will adopt the more lenient stance; I will develop the *theory* so that Cassie is allowed to believe assertions about

[19] Of course, humans sometimes violate this principle—for example, when they lose consciousness and get back to their senses to find certain bodily states persisting. But perhaps something similar to the above principle is why such unfortunate incidents are very perplexing.

her bodily states holding, even if she does not *feel* those states. Note that, given Axiom 5.8, this would typically involve belief revision that favors information coming from one source over another (Johnson and Shapiro, 2000a). A particular implementation may choose to disable this mechanism (I will show below how this could be done within the current theory).

The Seventh Principle of Change. *For every $s \in$* TEMP$\cup$PERM *and for every $i \in \mathbb{N}$ ($i > 1$), if Cassie determines that s persists at* $*\mathsf{NOW}_i$*, then there is some $s' \in$* TEMP *such that Cassie determines that s' starts, or ceases, at* $*\mathsf{NOW}_i$.

The gist of the above statement is that determining persistence cannot, by itself, be responsible for the movement of NOW. I have informally stated that NOW moves only if Cassie determines that a state starts or ceases to hold. But the second principle of change (Axiom 5.5) makes a weaker claim: two consecutive NOW-MTFs are not epistemically-equivalent. Theoretically, this may be satisfied in a situation where Cassie determines that some state persists (not starts or ceases). Practically, however, this cannot be the case. In particular, consider what is involved in determining that a state persists. First, there are only three means by which Cassie may determine that a state, s, persists: inference, perception, and direct assertion (proprioception being ruled out by the sixth principle of change). Now, let us take a careful look at each of these.

- **Inference.** Suppose that Cassie determines that s persists at $*\mathsf{NOW}_i$ through inference. According to Definition 5.15, this means that, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$ and, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$ and $\beta \nvdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. Note that, not only do we require that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, Cassie does not have any explicit beliefs about whether s holds, but that she *cannot* infer whether it does, given everything that she knows. How, then, can such an inference be possible at $*\mathsf{NOW}_i$? Something must have changed. In particular, some set of propositions, P, must have been added to β (assuming monotonicity) so that what was not possible at $*\mathsf{NOW}_{i-1}$ is possible at $*\mathsf{NOW}_i$. Naturally, P must have been added to β by some means other than inference. For if it was merely inferred, then it cannot be the case that at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. Therefore, P must originate from perception, proprioception, or direct assertion. Now, if P includes a proposition of the form $\mathsf{Holds}(s',*\mathsf{NOW}_i)$, for some $s' \in \Psi(\text{TEMP})$, then Cassie determines that s' holds at $*\mathsf{NOW}_i$. If this is a determination of s' starting to hold, then the seventh principle of change is satisfied; if it is a determination of the mere persistence of s', then it must have been achieved through perception or direct assertion—the two cases I review below. If, on the other hand, P does not include any propositions of the form $\mathsf{Holds}(s', {}^*\mathsf{NOW}_i)$, then elements of P can only get to β through direct assertion (since one can only perceive or feel what is present). Since such propositions cannot be reporting a (current) change in the environment, I assume that their assertion is not responsible for the movement of NOW (more on this below).

- **Perception.** Suppose that Cassie determines that s persists at ${}^*\mathsf{NOW}_i$ through perception. As per Definition 5.15, at ${}^*\mathsf{NOW}_{i-1}$, Cassie does not have any beliefs (implicit or explicit) about whether s holds, whereas, at ${}^*\mathsf{NOW}_i$, she can actually perceive it. But, for this to happen, something else must have changed. In particular, there are two general scenarios. The first involves a change in one or more of Cassie's bodily states. In this scenario, Cassie changes the location, or direction, of her sensors in such a way that they come to have access to whichever environmental phenomena constitute the state s. In the situation of Figure 5.11, for instance, not only does Cassie perceive that the walk-light is on, but, crucially, she turns her head toward the walk-light. If s is the state of the walk-light's being on, then, in that situation, Cassie determines that s persists at ${}^*\mathsf{NOW}_i$ through (visual) perception. In addition, she determines that s' starts at ${}^*\mathsf{NOW}_i$, where s' is the bodily state of Cassie's looking toward the walk-light. In fact, this change in s' is causally-responsible for Cassie's determining that s persists. It is hard to come up with similar examples for auditory perception. One possibility, however, is Cassie's coming to perceive a certain sound as a result of her *moving* closer to its source. The second scenario involves the removal of a barrier, thereby allowing Cassie's sensors to have access to the portion of the environment including what Cassie can recognize as the state s holding. An example is when Cassie opens the door of a room to see the room lights on, or to hear the radio therein playing. Note that, in this scenario, determining that a state s persists through perception is accompanied by (indeed, causally-dependent on) determining, again through perception, that some state s' starts (or ceases, depending on the agent's perspective).[20] The above discussion indicates that, in general, when Cassie determines that a state

[20]In addition to the two scenarios discussed here, there is actually a third one whereby Cassie determines that a state persists through perception. This is when Cassie starts to *attend* to a particular state of the environment that is already accessible by her sensors. For example, Cassie may be looking towards the walk-light, with no barriers in between, but only realizes that the walk-light is on when she starts paying attention to it. Being attentive to a particular state is itself a temporary state and it is a change of that state that allows Cassie to determine the persistence of other states. Since we have not formalized this notion of attending in our theory, we do not discuss this issue any further.

starts (or ceases), it may also be the case that she determines that one or more states persist. In particular, this happens in two cases corresponding to the two scenarios discussed above. First, the state change determined by Cassie involves the removal of a barrier to perception. Second, the state change is that of a bodily state that uses a modality capable of conveying perceptual information. Technically, this is a modality to which there are corresponding variables in both $\mathcal{M}_{\text{prop}}$ and $\mathcal{M}_{\text{per}}$.

- **Direct assertion.** In the cases of inference and perception, I have argued, from an empirical point of view, for the validity of the seventh principle of change. As the reader must have noted, the discussion of these two cases, in many respects, resembled a proof of the statement of the seventh principle of change. The same could be done here; I can argue that, empirically, the seventh principle of change is a plausible statement for the case of direct assertion of persistence. However, the situation is less idealistic than in the cases of inference and perception; our *theory* of direct assertion is still not developed enough to reflect empirical intuitions.[21] In particular, our account of direct assertion treats it as an inspiration-like phenomenon, a mysterious activity resulting in new beliefs appearing, out of the blue, in Cassie's mind. Actually, this is fine and is indeed the traditional model of adding information to a knowledge base. However, empirically, this is not how it exactly works. An assertion does not make its way magically into one's head; rather, there is an assertion *event* that results in one's acquiring a new belief. Forgetting aside, an agent is aware of such an assertion event, when it happened, and who its agent was. The theory, however, does not include an account of these assertion events.[22] That is, when told something, Cassie does not have any beliefs about the state changes associated with that *telling*, she is only aware of the asserted proposition, not of the assertion event proper. To see why this is problematic for a proof-like justification of the seventh principle of change, consider the following example. Suppose Cassie determines, through direct assertion, that the walk-light is on. In other words, someone, Stu for example, says: "Cassie, the walk-light is now on". Assuming that nothing else changes in

[21]The reader should, nevertheless, note that linguistic forms of communication with Cassie, particularly in natural language, have been thoroughly studied in the past (Shapiro, 1989; Ali and Shapiro, 1993; Shapiro and Rapaport, 1995; Shapiro, 1996; Shapiro, 2000, for instance).

[22]Though, currently, the point of making Cassie aware of the *source* of each assertion is under investigation (Johnson and Shapiro, 2000a).

the environment, one thing that has certainly changed is that a new sentence has been uttered (or typed, if you will). More precisely, a state of Stu's uttering something starts, persists for a while, and then ceases, resulting in Cassie's determining the persistence of the walk-light's on-state. Thus, the seventh principle of change readily holds; for Cassie's determining that s (the walk-light's being on) persists is accompanied by (and, again, causally-dependent on) the cessation of some s' (Stu's uttering something). However, since the theory does not account for assertion events, the above scenario cannot be used to empirically justify the seventh principle of change for the case of direct assertion of persistence. Rather, the principle has to be interpreted as a constraint that is imposed, by fiat, on the system. The constraint simply entails that asserting persistence does not, in itself, constitute a change. But note what this means. Suppose that Cassie determines that s persists at $*\mathsf{NOW}_i$ through direct assertion. Further, suppose that $*\mathsf{NOW}_i$ is the value of NOW between the *real* clock-times 12 p.m. and 12:10 p.m., and that the proposition $\mathsf{Holds}(s,*\mathsf{NOW})$ is asserted at 12:05 p.m. To Cassie, $[\![s]\!]$ holds throughout $[\![*\mathsf{NOW}_i]\!]$, but, to us, this might not be the case; $[\![s]\!]$ may have started to hold at 12:04 p.m. To some, this may look like a major problem; for, apparently, Cassie simply has a false belief. Nevertheless, I do not think that there are any major problems, and, even if there is some problem, nothing major seems at stake. To Cassie, the only reason why s's not holding between 12 and 12:04 would matter is if clock times have any significance to her. If they do not, then there is no problem. If they do, then she would have to be aware of them, and NOW will necessarily move with every tick of the clock. In that case, however, the problem disappears, since $*\mathsf{NOW}_i$ cannot extend over a period in which there are salient changes. Note that this follows from the first principle of change (Axiom 5.4). In general, *any* assertion that does not result in Cassie's determining that a state starts (or ceases) also does not result in NOW moving. Note that, were we to actually provide an account of assertion events, direct assertion would cease to stand out as an independent method by which Cassie may acquire new beliefs. In particular, the assertion event itself would be *perceived* (through audition, for example) just as any other external event. What would be needed then is an axiom to the effect that whenever an assertion event takes place, Cassie should believe the content of the assertion (more on this in Section 8.2.2).

Similar to the seventh principle of change, we have the following principle.

The Eighth Principle of Change. *For every $s \in$*TEMP *and for every $i \in \mathbb{N}$ ($i > 1$), if Cassie starts, or ceases, to perceive s at* $*\mathsf{NOW}_i$*, then there is some $s' \in$*TEMP *such that Cassie determines that s' starts, or ceases, to hold at* $*\mathsf{NOW}_i$.

Justifying this principle follows the same line of thought introduced above for the perception case of the seventh principle of change. Cassie starts to perceive *s* due to the removal of a barrier or the repositioning, or re-orienting, of her sensors. Similarly, if Cassie ceases to perceive *s*, then this must be the result of either the introduction of a barrier or of Cassie's changing the position or orientation of her sensors.

Whether the seventh and eighth principles of change are empirically-justified or theoretically-motivated, I uphold them as guiding principles: mere determination of persistence, or change in perception, does not move NOW. This being said, let us now turn to what exactly *is* involved in the movement of NOW.

5.2.6 Mechanizing Belief Update

With the principles discussed above in mind, algorithm `state_change` in Figure 5.12 presents a general outline of the circumstances surrounding the movement of NOW. Of course, every step of this algorithm needs to be carefully explained. First, however, we need to be explicit about *when* the algorithm is executed and what the arguments, $S^{\uparrow}$ and $S^{\downarrow}$, signify. Algorithm `state_change` gets executed whenever Cassie's body and/or sensors detect a change regarding some state (actually, a set thereof). The types of *change* referred to here will be made explicit. But note that I have excluded information about states reported by direct assertion. This is because of the peculiar features of direct assertion discussed in Section 5.2.5. I choose to have a separate set of algorithms for handling direct assertion; these will be presented below. Thus, algorithm `state_change` is only concerned with changes in perception or proprioception. The types of change are reflected on the nature of the arguments, $S^{\uparrow}$ and $S^{\downarrow}$, which is totally revealed by the preconditions. Both arguments are *sets* of TEMP states (the only sort of state that may be perceived or proprioceived). Members of $S^{\uparrow}$ are states that *start* to be detected as holding by Cassie's perceptual and proprioceptual systems. More precisely, these are states that either start to hold or just start to be perceived. This is reflected by the first and second pre-conditions. The first precondition states that, just prior to executing the algorithm, modality variables do not reflect anything about members of $S^{\uparrow}$. The second pre-condition,

PRE-CONDITIONS:

1. *For every $s \in S^{\uparrow}$, there is no $t \in \Psi(\mathcal{T})$ such that, for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s,t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s,t) \in *\mu$.*

2. *For every $s \in S^{\uparrow}$, there is some $\mu \in \mathcal{M}_{\text{prop}}$ such that $[\![s]\!]$ occupies the modality corresponding to μ, or there is some $\mu \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ.*

3. *For every $s \in S^{\downarrow}$, there is some $t \in \Psi(\mathcal{T})$ such that either, for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s,t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s,t) \in *\mu$.*

4. *For every $s \in S^{\downarrow}$, there is no $\mu \in \mathcal{M}_{\text{prop}}$ such that $[\![s]\!]$ occupies the modality corresponding to μ and there is no $\mu \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ.*

Algorithm `state_change`($S^{\uparrow} \subseteq$ TEMP, $S^{\downarrow} \subseteq$ TEMP)

1. $P_{\text{new}} \longleftarrow \{\}$.

2. Pick some $t_{\text{tr}} \in \mathcal{T}$, such that $t_{\text{tr}} \notin \Psi(\mathcal{T})$.

3. $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \prec t_{\text{tr}}\}$.

4. For all $s_i \in S^{\uparrow}$

    5. If $\beta \vdash \mathsf{Holds}(s_i, *\mathsf{NOW})$ then `start_ceive`(s_i, t_i), where t_i is the state interval associated with s_i such that $\beta \vdash *\mathsf{NOW} \sqsubset t_i$

    6. else

        6a. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
        6b. If $\beta \vdash \neg\mathsf{Holds}(s_i, *\mathsf{NOW})$ then `state_start`(s_i, t_i, t_{tr}).
        6c. Else, `state_persist`(s_i, t_i).
        6d. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{Mholds}(s_i, t_i)\}$.

7. For all $s_i \in S^{\downarrow}$, `cease_perceive`(s_i).

8. `setup_new_MTF`.

9. `Forward`(P_{new}).

Figure 5.12: Algorithm `state_change`.

on the other hand, states that the states denoted by members of $S^{\uparrow}$ actually occupy proprioceptual modalities or are perceived via perceptual modalities. Part of the function of the algorithm is to update modality variables to reflect the new situation (see Theorem 5.5 below). It should be noted that, given the fourth, seventh, and eighth principles of change, there must be some $s \in S^{\uparrow}$ such that, prior to executing the algorithm, $\beta \vdash \neg\mathsf{Holds}(s, {*}\mathsf{NOW})$. That is, some state in $S^{\uparrow}$ must actually be starting to hold. Members of $S^{\downarrow}$ are states that *cease* to be detected as holding by Cassie's perceptual and proprioceptual systems. This may happen if a state ceases or just ceases to be perceived. Again, this is reflected by the third and fourth pre-conditions. I will, henceforth, assume that, whenever there are states satisfying its pre-conditions, algorithm `state_change` gets executed.

Let us now take a careful look at each step of the algorithm. Step 1 initializes a variable, P_{new}, that is used to collect newly introduced propositions about states holding (see step 6d). Steps 2 and 3 introduce a new interval, t_{tr}, and assert that it follows $*\mathsf{NOW}$. What is the significance of t_{tr}? This will only be made precise when algorithm `state_start` (step 6b) is discussed. However, some informal explanation is due at this point. t_{tr} is introduced to serve as the transition interval associated with all the state transitions involved in the movement of NOW effected by algorithm `state_change`. As pointed out above, there is at least one member of $S^{\uparrow}$ that has just started to hold. There could be more, and there could be members of $S^{\downarrow}$ that have just ceased. Since all of these transitions are (to Cassie) simultaneous, then a single transition interval needs to be associated with them.

Step 5 checks if any of the members of $S^{\uparrow}$ is a state that is already believed to hold. If this is the case, then such a state must have made it into $S^{\uparrow}$ because it started to be perceived or proprioceived.[23] This initiates algorithm `start_ceive`, shown in Figure 5.13.[24] The first step of `start_ceive` initiates algorithm `start_proprioceive` (see Figure 5.14) which updates proprioception modality variables. Note that the algorithm does not check if its argument, s, actually occupies any modalities; it merely checks if it is a bodily state. The reason is that, given precondition 2, $[\![s]\!]$ is either perceived or proprioceived. By Axiom 5.10, a bodily state cannot be perceived. Therefore, if s is a bodily state, then it must be proprioceived. If the state is *not* a bodily

[23] Recall our discussion in Section 5.2.5 regarding how Cassie may have a belief about a bodily state that she does not feel.

[24] Here, I am alluding to the term "ception" used in (Talmy, 2000). Though Talmy uses the term to cover perception, proprioception, and conception; I only use it to cover the first two.

Algorithm `start_ceive`(s,t)

1. `start_proprioceive`(s,t).
2. `start_perceive`(s,t).

Figure 5.13: Algorithm `start_ceive`.

Algorithm `start_proprioceive`(s,t)

1. If $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$, then, for all $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$,

   $\mu \longleftarrow \mathsf{MHolds}(s,t)$.

Figure 5.14: Algorithm `start_proprioceive`.

Algorithm `start_perceive`(s,t)

1. If $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, then

   $\mathrm{Mod}_{\mathrm{per}}(s) \longleftarrow *\mathrm{Mod}_{\mathrm{per}}(s) \cup \{\mathsf{MHolds}(s,t)\}$.

Figure 5.15: Algorithm `start_perceive`.

state, then it must be perceived via some perception modality. Updating the appropriate perception modality variable is the function of algorithm `start_perceive` of Figure 5.15. Note that, were we to opt for an implementation in which Cassie may *not* hold beliefs about bodily states that she does not feel, step 5 of algorithm `state_change` should initiate algorithm `start_perceive` directly.

It should also be noted that, unlike perception modality variables, the contents of proprioception modality variables are *overwritten*. Thus, the state previously occupying a proprioceptual modality no longer has a proposition in the corresponding modality variable (see the third principle of change in Section 5.2.5). This means that algorithm `setup_new_MTF` (Figure 5.9) would not include such a state in the new MTF. In fact, given Axiom 5.6 and the pre-conditions, all the states displaced by the newly starting bodily states are the bodily states in $S^{\downarrow}$.

If, in step 5 of `state_change`, the state is not already believed to be holding, then there are two possibilities: (i) the state has just started to hold or (ii) it has just started to be perceived, with

Algorithm `state_start`(s,t,t_{tr})

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t)\}$.
2. `event_start`(s,t,t_{tr}).
3. `start_ceive`(s,t).

Figure 5.16: Algorithm `state_start`.

Algorithm `event_start`(s,t,t_{tr})

1. Pick some $\dot{e} \in \dot{\mathcal{E}}$, such that $\dot{e} \notin \Psi(\dot{\mathcal{E}})$.
2. $\beta \longleftarrow \beta \cup \{\mathsf{Cat}(\dot{e}, \uparrow s), \mathsf{Occurs}(\dot{e}, t_{\text{tr}}), t_{\text{tr}} \supset\subset t\}$.

Figure 5.17: Algorithm `event_start`.

Cassie having no beliefs about it. In both cases, a new state interval needs to be associated with the state. This is achieved by step 6a. If the state is believed to be *not* holding, then it must have just started; otherwise, it has just been perceived to persist. In the first case, algorithm `state_start` gets initiated. This is shown in Figure 5.16. The algorithm performs three main tasks. First, it adds a new belief about the state, s, holding (step 1). Second, it introduces a new onset of s by initiating algorithm `event_start` (see figure 5.17). The algorithm introduces an event of category $\uparrow s$ reflecting Cassie's conceiving of the onset of s. Note that the transition interval, t_{tr}, introduced by step 2 of `state_change` is asserted to be the time of occurrence of that onset and to abut the state interval associated with the starting state (thus observing **AOC3**). Finally, `state_start` updates modality variables by initiating algorithm `start_ceive` in step 3.

If an element of $S^{\uparrow}$ has just started to be perceived with Cassie having no beliefs about whether it holds, then algorithm `state_persist` gets initiated by step 6c of `state_change`. This algorithm is shown in Figure 5.18. It is very similar to algorithm `state_start` with two important difference.

1. A new transition interval, t' is generated locally and passed to algorithm `state_start`. The reason t_{tr} is not used here is that t_{tr} is asserted to follow $\mathsf{*NOW}$ (step 3 of `state_change`) whereas the state s is not known to have just started.

Algorithm `state_persist`(s,t)

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t)\}$.
2. Pick some $t' \in \mathcal{T}$, such that $t' \notin \Psi(\mathcal{T})$.
3. `event_start`(s,t,t').
4. `start_perceive`(s,t).

Figure 5.18: Algorithm `state_persist`.

Algorithm `cease_perceive`(s)

1. If $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, then

$$\mathrm{Mod}_{\mathrm{per}}(s) \longleftarrow *\mathrm{Mod}_{\mathrm{per}}(s) \setminus \{\mathsf{MHolds}(s,t)\}$$

Figure 5.19: Algorithm `cease_perceive`.

2. Given the sixth principle of change, the algorithm assumes that its state-argument does not occupy any proprioceptual modalities and, thus, initiated algorithm `start_perceive`, rather than `start_ceive`, in step 3.

Step 7 of `state_change` processes members of the set $S^{\downarrow}$. Regardless of whether some $s_i \in S^{\downarrow}$ has actually ceased or simply ceased to be perceived, perception modality variables must be updated to reflect the lack of perceptual information about s_i (as indicated by pre-condition 4). This is the function of algorithm `cease_perceive` of Figure 5.19. Step 8 initiates algorithm `setup_new_MTF` of Figure 5.9 which is responsible for the actual movement of NOW and the construction of the new NOW-MTF.[25] The final step of the algorithm initiates forward inference on all of the new assertions of states holding. See Section 1.5.4 for the reasoning behind this step.

In addition to its pre-conditions, algorithm `state_change` has a number of significant post-conditions.

Theorem 5.5 *For every $s \in$ TEMP and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![*\mathsf{NOW}_{i-1}]\!]$, algorithm* `state_change` *gets initiated with $s \in S^{\uparrow}$, then, following the execution of the algorithm, there is some $t \in \Psi(\mathcal{T})$*

[25] In Section 8.1, a revised version of algorithm `setup_new_MTF` will be introduced which guarantees that, if any of the members of $S^{\downarrow}$ has actually ceased, then Cassie's beliefs would reflect this fact and a cessation event would be introduced.

such that $\beta \vdash \mathsf{MHolds}(s,t)$*, and either, for some* $\mu \in \mathcal{M}_{\text{prop}}$*,* $*\mu = \mathsf{MHolds}(s,t)$ *or, for some* $\mu \in \mathcal{M}_{\text{per}}$*,* $\mathsf{MHolds}(s,t) \in *\mu$*.*

Proof. Let s be an arbitrary member of $S^{\uparrow}$. At $[\![*\mathsf{NOW}_{i-1}]\!]$, it is either the case that (i) $\beta \vdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$, (ii) $\beta \vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$, or (iii) $\beta \not\vdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$ and $\beta \not\vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$. Consider each case.

(i) Suppose that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$. By step 5 and Axiom 5.2, there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s,t)$. This proves the first conjunct in the consequent of the statement of the theorem. Step 5 initiates algorithm `start_ceive` with arguments s and t. By pre-condition 2, either there is some $\mu \in \mathcal{M}_{\text{prop}}$ such that $[\![s]\!]$ occupies the modality corresponding to μ, or there is some $\mu \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ. In the first case, following the execution of algorithm `start_proprioceive`, there is some $\mu \in \mathcal{M}_{\text{prop}}$ such that $*\mu = \mathsf{MHolds}(s,t)$. By Axiom 5.6, there is no $s' \in S^{\uparrow}$ ($s' \neq s$) occupying the modality corresponding to μ. Therefore, subsequent initiations of `start_proprioceive` through the end of `state_change` do not change the value of μ. In the second case, following the execution of algorithm `start_perceive`, there is some $\mu \in \mathcal{M}_{\text{per}}$ (namely Mod_{per}, as per Axiom 5.9) such that $\mathsf{MHolds}(s,t) \in *\mu$. The only place in `state_change` where the proposition $\mathsf{MHolds}(s,t)$ may be removed from $*\mu$ is algorithm `cease_perceive` in step 7. But, given the pre-conditions, $S^{\uparrow}$ and $S^{\downarrow}$ are disjoint, and algorithm `cease_perceive` never gets applied to s. Therefore $\mathsf{MHolds}(s,t)$ continues to be a member of $*\mu$ through the end of `state_change`.

(ii) Suppose that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$. By step 6b, algorithm `state_start` gets initiated with s and t as arguments, where t is the new interval introduced in step 6a. By step 1 of `state_start`, $\beta \vdash \mathsf{MHolds}(s,t)$, which proves the first conjunct of the consequent of the theorem. The proof of the second conjunct follows that of part (i) above.

(iii) Suppose that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \not\vdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$ and $\beta \not\vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$. By step 6c, algorithm `state_persist` gets initiated with s and t as arguments, where t is the new interval introduced in step 6a. By step 1 of `state_persist`, $\beta \vdash \mathsf{MHolds}(s,t)$, which proves the first conjunct of the consequent of the theorem. Step 4 of algorithm `state_persist` initiates algorithm `start_perceive` with s and t as arguments. By the sixth principle of change, $[\![s]\!]$ cannot be occupying a proprioceptual modality. Therefore, given pre-condition 2, there is some $\mu \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ. Thus, following the execution of `start_perceive`, there is some $\mu \in \mathcal{M}_{\text{per}}$ such that $\mathsf{MHolds}(s,t) \in *\mu$. Following the proof of part (i), $\mathsf{MHolds}(s,t)$ continues to be a member of $*\mu$ through the end of `state_change`. $\square$

A corresponding result holds for members of $S^{\downarrow}$.

Theorem 5.6 *For every $s \in$*TEMP *and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![{}^*\mathsf{NOW}_{i-1}]\!]$, algorithm* `state_change` *gets initiated with $s \in S^{\downarrow}$, then, following the execution of the algorithm, there is no $t \in \mathcal{T}$ such that for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s,t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s,t) \in *\mu$.*

Proof. Let s be an arbitrary member of $S^{\downarrow}$. Given pre-condition 3, suppose that there is some $t \in \mathcal{T}$ and $\mu \in \mathcal{M}_{\text{prop}}$ such that $*\mu = \mathsf{MHolds}(s,t)$. By pre-condition 4 and the third principle of change, there is some $s' \in S^{\uparrow}$ that occupies the modality corresponding to μ. By the proof of Theorem 5.5, a proposition $\mathsf{MHolds}(s',t')$ (for some $t' \in \Psi(\mathcal{T})$) overwrites the contents of μ. On the other hand, suppose that there is some $t \in \mathcal{T}$ and $\mu \in \mathcal{M}_{\text{per}}$ such that $\mathsf{MHolds}(s,t) \in *\mu$. By Axiom 5.9, $\mu = \text{Mod}_{\text{per}}(s)$. Following the execution of algorithm `cease_perceive` in step 7 with s as an argument, the proposition $\mathsf{MHolds}(s,t)$ gets removed from $*\text{Mod}_{\text{per}}(s)$. Since $S^{\downarrow}$ and $S^{\uparrow}$ are disjoint (given the pre-conditions), then the result of Theorem 5.5 does not apply to any of the members of $S^{\downarrow}$. Therefore, following the execution of `state_change` there is no $t \in \mathcal{T}$ such that, for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s,t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s,t) \in *\mu$. $\square$

Theorem 5.7 *For every $s \in$*TEMP *and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![{}^*\mathsf{NOW}_{i-1}]\!]$, algorithm* `state_change` *gets initiated with $s \in S^{\uparrow}$, then, following the execution of the algorithm, $\beta \vdash \mathsf{Holds}(s,{}^*\mathsf{NOW}_i)$.*

Proof. Let s be an arbitrary member of $S^{\uparrow}$. Given the proof of Theorem 5.5, there is some $t \in \mathcal{T}$ such that, just before executing step 8 of algorithm `state_change`, for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s,t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s,t) \in *\mu$. Therefore, by executing algorithm `setup_new_MTF`, NOW moves from ${}^*\mathsf{NOW}_{i-1}$ to ${}^*\mathsf{NOW}_i$ and either step 3 or step 6 adds the proposition ${}^*\mathsf{NOW}_i \sqsubset t$ to β. Since, by Theorem 5.5, $\beta \vdash \mathsf{MHolds}(s,t)$, then, given **AS2** and **AS3**, $\beta \vdash \mathsf{Holds}(s,{}^*\mathsf{NOW}_i)$. Since s is arbitrary, then the result follows for all $s \in S^{\uparrow}$. $\square$

Given the above theorem, two corollaries readily follow.

Corollary 5.2 *For every $s \in \Psi($*TEMP*$)$ and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![{}^*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,{}^*\mathsf{NOW}_{i-1})$ and algorithm* `state_change` *gets initiated with $s \in S^{\uparrow}$, then Cassie determines that s starts to hold at ${}^*\mathsf{NOW}_i$.*

Proof. Follows directly from Definition 5.14 and Theorem 5.7. □

Corollary 5.3 *For every s ∈*TEMP *and $i \in \mathbb{N}$ ($i > 1$), if, at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$, $\beta \nvdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$, *and algorithm* `state_change` *gets initiated with $s \in S^{\uparrow}$, then Cassie determines that s persists at* $*\mathsf{NOW}_i$.

Proof. Follows directly from Definition 5.15 and Theorem 5.7. □

The above corollaries show that, as a side-effect, the execution of algorithm `state_change` results in Cassie's determining that states start or persist. In what follows, I will make the following complementary assumption: if, through perception or proprioception, Cassie determines that a state starts or persists, then this may only be an effect of executing algorithm `state_change`. In other words, algorithm `state_change` is the *only* component of the system responsible for Cassie's determining the start or persistence of states through perception or proprioception.

As pointed out above, separate algorithms are responsible for the direct assertion of state change. Direct assertion of persistence is the responsibility of algorithm `assert_persist` of Figure 5.20. Asserting onsets and cessations are done through the two algorithms shown in Figures 5.21 and 5.22, respectively.[26] The algorithms should be self-evident. The only thing to note is that algorithms `assert_persist` and `assert_start` are quite similar. The crucial differences (other than the pre-conditions, that is) are (i) how each algorithm introduce onset events and (ii) the fact that executing `assert_start` results in NOW moving while executing `assert_persist` does not. Also note the introduction of a cessation event in algorithm `assert_cease`. The algorithm uses algorithm `event_cease` of Figure 5.23. `event_cease` introduces a new event token that is asserted to be of category $\downarrow s$, for the ceasing state s. Note the temporal location of this event: it precedes the new *NOW (step 5 of `assert_cease`) and is abutted by the state interval associated with s (step 2 of `event_cease`).

Given these algorithms we can easily prove the following results.

Theorem 5.8 *For every s ∈*TEMP∪PERM *and $i \in \mathbb{N}$ ($i > 1$), if algorithm* `assert_persist` *is initiated at* $[\![*\mathsf{NOW}_i]\!]$ *with $s \in S$, then Cassie determines that s persists at* $*\mathsf{NOW}_i$.

[26] If, in a particular implementation, it is decided that Cassie should filter out any direct assertions about states of her body, then these algorithms constitute the locations of such filters.

PRE-CONDITIONS:

1. *If* $*\mathsf{NOW} = *\mathsf{NOW}_i$ *for some* $i > 1$*, then, for every* $s \in S$*,* $\beta \nvdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$.

2. *If* $*\mathsf{NOW} = *\mathsf{NOW}_i$ *for some* $i > 1$*, then, for every* $s \in S$*,* $\beta \nvdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$.

Algorithm `assert_persist`($S \subseteq$ TEMP $\cup$ PERM)

1. $P_{\text{new}} \longleftarrow \{\}$.
2. For all $s_i \in S$
    3. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
    4. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i, t_i), *\mathsf{NOW} \sqsubset t_i\}$.
    5. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
    6. Pick some $t_i' \in \mathcal{T}$, such that $t_i' \notin \Psi(\mathcal{T})$.
    7. `event_start`(s_i, t_i, t_i').
8. `Forward`(P_{new}).

Figure 5.20: Algorithm `assert_persist`.

PRE-CONDITIONS:

1. *For every* $s \in S$*,* $\beta \vdash \neg\mathsf{Holds}(s, *\mathsf{NOW})$.

Algorithm `assert_start`($S \subseteq$TEMP $\cup$ PERM)

1. Pick some $t_{\text{tr}} \in \mathcal{T}$, such that $t_{\text{tr}} \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \prec t_{\text{tr}}\}$.
3. `setup_new_MTF`.
4. $P_{\text{new}} \longleftarrow \{\}$.
5. For all $s_i \in S$
    6. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
    7. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i, t_i), *\mathsf{NOW} \sqsubset t_i\}$.
    8. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
    9. `event_start`$(s_i, t_i, t_{\text{tr}})$.
10. `Forward`(P_{new}).

Figure 5.21: Algorithm `assert_start`.

PRE-CONDITIONS:

1. *For every* $s \in S$, $\beta \vdash \mathsf{Holds}(s, {*}\mathsf{NOW})$.

Algorithm `assert_cease`($S \subseteq$ TEMP)

1. $P_{\text{new}} \longleftarrow \{\}$.
2. *old-now* $\longleftarrow {*}\mathsf{NOW}$.
3. `setup_new_MTF`.
4. Pick some $t_{\text{tr}} \in \mathcal{T}$, such that $t_{\text{tr}} \notin \Psi(\mathcal{T})$.
5. $\beta \longleftarrow \beta \cup \{t_{\text{tr}} \prec {*}\mathsf{NOW}\}$.
6. For all $s_i \in S$
    7. $\beta \longleftarrow \beta \cup \{\neg\mathsf{Holds}(s_i, {*}\mathsf{NOW})\}$.
    8. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\neg\mathsf{Holds}(s_i, {*}\mathsf{NOW})\}$.
    9. `event_cease`(s_i, t_i, t_{tr}),
    where t_i is a state interval associated with s_i such that $\beta \vdash$ *old-now* $\sqsubset t_i$.
10. `Forward`(P_{new}).

Figure 5.22: Algorithm `assert_cease`.

Algorithm `event_cease`(s, t, t_{tr})

1. Pick some $\dot{e} \in \dot{\mathcal{E}}$, such that $\dot{e} \notin \Psi(\dot{\mathcal{E}})$.
2. $\beta \longleftarrow \beta \cup \{\mathsf{Cat}(\dot{e}, \downarrow s), \mathsf{Occurs}(\dot{e}, t_{\text{tr}}), t \supset\subset t_{\text{tr}}\}$.

Figure 5.23: Algorithm `event_cease`.

Proof. For every $s \in S$, by step 4 of `assert_persist`, **AS2**, and **AS3**, $\beta \vdash \mathsf{Holds}(s, {}^*\mathsf{NOW}_i)$ at $[\![{}^*\mathsf{NOW}_i]\!]$, following the execution of the algorithm. Given the pre-conditions of `assert_persist` and Definition 5.15, Cassie determines that s persists at ${}^*\mathsf{NOW}_i$. □

Theorem 5.9 *For every $s \in$* TEMP∪PERM *and $i \in \mathbb{N}$ $(i > 1)$, if algorithm* `assert_start` *is initiated at $[\![{}^*\mathsf{NOW}_{i-1}]\!]$ with $s \in S$, then Cassie determines that s starts to hold at ${}^*\mathsf{NOW}_i$.*

Proof. Given the pre-condition of `assert_start` and Definition 5.14, the proof follows that of Theorem 5.8. □

Theorem 5.10 *For every $s \in$* TEMP *and $i \in \mathbb{N}$ $(i > 1)$, if algorithm* `assert_cease` *is initiated at $[\![{}^*\mathsf{NOW}_{i-1}]\!]$ with $s \in S$, then Cassie determines that s ceases to hold at ${}^*\mathsf{NOW}_i$.*

Proof. For every $s \in S$, by step 7 of `assert_cease`, $\beta \vdash \neg\mathsf{Holds}(s, {}^*\mathsf{NOW}_i)$ at $[\![{}^*\mathsf{NOW}_i]\!]$, following the execution of the algorithm. Given the pre-condition of `assert_cease` and Definition 5.17, Cassie determines that s ceases to hold at ${}^*\mathsf{NOW}_i$. □

Similar to the assumption we made regarding algorithm `state_change`, we assume that algorithms `assert_persist`, `assert_start`, and `assert_cease` provide the sole means by which Cassie may determine, through direct assertion, that a state persists, starts, or ceases, respectively. Such assumptions are needed to be able to prove further results about the system (see Sections 7.5.4 and 8.2).

Before concluding this section, one final algorithm needs to be introduced. This algorithm takes care of all the initialization procedures required by the system. Figure 5.24 shows the algorithm which basically sets up the first NOW-MTF. Step 2 initiates algorithm `initialize_NOW` of Figure 5.3. The rest of the steps update the belief space and modality variables to reflect states of Cassie's body or the perceived environment as Cassie starts operating.

5.3 Conclusions

In Chapters 3 and 4, a logical theory axiomatizing the domains of time, states, and events was laid out. As pointed out in those chapters, the logical theory presents a perspective of time and situations from the point of view of an acting agent that reasons about situations as time passes by. In this

PRE-CONDITIONS:

1. *For every $s \in S$, there is some $\mu \in \mathcal{M}_{\text{prop}}$ such that $[\![s]\!]$ occupies the modality corresponding to μ, or there is some $\mu \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ.*

Algorithm `initialize`($S \in$ TEMP)

1. $P_{\text{new}} \longleftarrow \{\}$.
2. `initialize_NOW`.
3. For all $s_i \in S$
    4. Pick some $t_i \in \mathcal{T}$ such that $t_i \notin \Psi(\mathcal{T})$.
    5. Pick some $t_i' \in \mathcal{T}$, such that $t_i' \notin \Psi(\mathcal{T})$.
    6. `state_start`(s_i, t_i, t_i').
    7. $\beta \longleftarrow \beta \cup \{{*}\mathsf{NOW} \sqsubset t_i\}$.
    8. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
9. `Forward`(P_{new}).

Figure 5.24: Algorithm `initialize`.

chapter, we have investigated in detail issues of the passage of time. Temporal progression, as presented in this chapter, is a meta-theoretical phenomenon that takes place at the PML. It has been observed that the domain of time intervals is covered by a set of lattices, MTFs, each corresponding to all states that hold at any atomic interval of time. Of particular interest, are NOW-MTFs—MTFs whose smallest elements are NOW-intervals, representing points of direct experience in Cassie's temporal career. Differences among consecutive NOW-MTFs mirror changes that take place as time passes by. A number of general principles were presented that constrain transitions from one NOW-MTF to the next, how they happen, and what sorts of phenomena are responsible for them. These principles were embodied in a number of algorithms that precisely characterize how a new NOW-MTF is constructed given the previous NOW-MTF and the changes responsible for the transition. In this respect, a mechanism for projecting bodily and perceived states from one NOW-MTF to the next was worked out. The mechanism, using a set of PML variables whose values reflect what each agent modality is being used for, provides a simple, reasoning-free method for modeling the persistence of states about which information is continuously available through bodily-feedback.

In the next two chapters, we will look at two problems that emerge as a result of the interaction between the logical theory of time and meta-theoretical temporal progression system. Basically, these are problems with reasoning about "now".

Chapter 6

The Problem of the Unmentionable Now

Endowing an agent with a subjective sense of temporal progression directly impacts the inner workings of reasoning processes. The problem presented in this short chapter is a rather simple one. Nevertheless, it sheds light on what reasoning *in* time implies: time must be built into the inference and acting systems; it is not just the subject-matter of reasoning. In the following, I will try to keep the discussion at a minimum level of formality. This will be done mainly to emphasize the conceptual basis of the problem (which is the important thing) without being drawn into unnecessary technical details. It is also important that the reader appreciate the generality of the problem and its independence of any particular formalism.

6.1 The Deictic Nature of Acting

In Chapter 3, a four-way classification of states (into eternal, permanent, co-permanent, and temporary) has been proposed. In this section, let us focus only on two sorts of states: eternal and temporary. How does introducing the eternal-temporary distinction affect the reasoning process? Consider the following sentence schema:

(6.1) IF *ant* THEN *cq*

Schema (6.1) means that if Cassie believes that *ant* holds, then she may also believe that *cq* holds ("may", because the rule might not fire, even if Cassie believes that *ant* holds). This works fine if *ant* and *cq* denote eternal states (for example, "IF Mammals(whales) THEN GiveBirth(whales)"). However, if, instead, they denote temporary states, we need to quantify over time; the temporary state-denoting terms by themselves do not say anything about the states holding over time. Schema (6.2) captures the intended meaning: if Cassie believes that *ant* holds over time t, then she may also believe that *cq* holds over t.[1]

(6.2) $\forall t$ IF $\mathsf{Holds}(ant, t)$ THEN $\mathsf{Holds}(cq, t)$

Schemata (6.1) and (6.2) represent sentences for *pure* reasoning. Doing reasoning in the service of acting requires sentences for *practical* reasoning. In particular, let us concentrate on one kind of acting rule: rules about *when* to act.[2] Imagine Cassie operating in a factory. One reasonable belief that she may have is that, when the fire-alarm sounds, she should leave the building. The underlying schema for such a belief is represented in (3) (Kumar and Shapiro, 1994a).

(6.3) WHEN *cond* DO *act*

The intended interpretation of (6.3) is that, when Cassie comes to believe that the condition *cond* holds, she should perform the act *act*. Again, this is fine so long as *cond* denotes an eternal state. If forward inference causes both *cond* and (6.3) to be asserted in Cassie's belief space, she will perform *act*. What if *cond* denotes a temporary state? Obviously, we need to somehow introduce time, since assertions about temporary states holding essentially involve reference to time. Following (6.2), one may propose the following representation.

(6.4) $\forall t$ WHEN $\mathsf{Holds}(cond, t)$ DO *act*

Asserting $\mathsf{Holds}(cond, t_1)$, for some particular time t_1, (6.4) would be matched and Cassie would perform *act*. On the face of it, (6.4) looks very innocent and a straightforward extrapolation of (6.2). However, a closer look shows that this is by no means the case. Using quantification over time works well for inference, since the consequent is a proposition that may just happen to be *about* time. Acting, on the other hand, takes place *in* time, resulting in an interesting problem. Table 6.1 represents a timed sequence of assertions entered into Cassie's belief space. The left column shows

[1] In $\mathcal{FOCS}$, (6.2) may be represented in either of the following equivalent ways:

(i) $\forall Tv[\mathsf{Holds}(ant, Tv) \Rightarrow \mathsf{Holds}(cq, Tv)]$

(ii) $\forall Tv \mathsf{Holds}([ant \Rightarrow cq], Tv)$

[2] By "rule", I mean a domain rule, expressed in the logical language, which Cassie might come to believe as a result of being told it in natural language. I do *not* mean a rule of inference, which would be implemented in the inference engine of the knowledge representation and reasoning system.

Assertion	Assertion Time
(6.5) $\forall t$ $\mathsf{WHENHolds}(\mathsf{Sounds}(\text{alarm}),t)$ $\mathsf{DOLeave}(\text{building})$	t_1
(6.6) $\mathsf{Holds}(\mathsf{Sounds}(\text{alarm}),t_0)$	t_2

Table 6.1: A timed sequence of assertions for the fire-alarm problem.

the assertion, and the right column shows Cassie's term for the time of the assertion (technically, this is value of NOW at the time of the assertion). The problem is that t_0 in (6.6) may refer to a time preceding t_2 (or even t_1). That is, (6.6) could be an assertion about the alarm sounding at some time in the past, something that we should be capable of asserting. Nevertheless, (6.6) matches (6.5) and Cassie leaves the building —at t_2— even though there is no danger!

One problem with (6.5) (and generally (6.4)) is that nothing relates the time of performing the act to the time at which the state holds. We may attempt to revise (6.4) by tying the action to that time:

(6.7) $\forall t[\mathsf{WHEN}\ \mathsf{Hold}(cond,t)\ \mathsf{DO}\ \mathsf{Perform}(act,t)]$

Where $\mathsf{Perform}(act,t)$ is intended to mean that Cassie should perform *act* at time t. However, this alleged semantics of Perform is certainly ill-defined; acts may only be performed *NOW, in the present. Cassie cannot travel in time to perform *act* in the past, at a time over which (she believes that) *cond* held. The basic problem seems to be quantifying over *all* times.

What we really want to say is that when the state holds *NOW, perform the act. That is,

(6.8) $\mathsf{WHEN}\ \mathsf{Hold}(cond,{*}\mathsf{NOW})\ \mathsf{DO}\ act$

However, we cannot mention "*NOW"; it is not itself a term in the logic. If we replace (6.5) in Table 6.1 with the appropriate variant of (6.8), "*NOW" in the left column would be just a shorthand for the term appearing in the right column, namely t_1. The assertion would, therefore, be very different from what we intended it to be. In particular, it would represent a rule to leave the building if the agent ever comes to believe that the fire-alarm sounded at t_1!

The problem resides in the deictic nature of acting. Acts can only be performed in the present, and, therefore, the only time that is relevant to an acting rule along the lines of (6.4) must be the present time. But note which "present time" is relevant; not the time of asserting the rule, but the time of using it—a totally unknown and unpredictable one.

There is also another factor contributing to the problem: the choice to not represent "now" in the object language. Note that, for example, the system of (Lespérance and Levesque, 1995) may indeed be able to represent our acting rule. Adopting their suggestion of introducing a new constant to represent each new "now" and equating it to the indexical object-language **now**, the rule may be represented as follows.

(6.9) $\forall t$ WHEN $[\mathsf{Holds}(cond, t) \ \wedge \ t = \mathbf{now}]$ DO act

However, for reasons discussed in Chapter 2, I choose not to introduce any notion of "now" in the object language. Our problem then is how to represent acting rules along the lines of (6.3) without mentioning "now" in any object-language sentence. You might think of this as an attempt to see how far one can go without introducing indexicals into the object language. Before presenting the proposed solution to the problem, we first need to discuss an approach that might seem to solve it. I will show that, although it may appear to eliminate the problem, it actually introduces more drastic ones.

6.2 The Time-Stamping Alternative

Inspecting the second row of Table 6.1, one may think that part of the problem is the inequality of the time appearing in the right column (t_2) to that in the left column (t_0). Indeed, if somehow we can ensure that these two times are identical, the problem may be solved. (Kamp, 1971) proposes an ingenious mechanism for correctly modeling a particularly interesting compositional property of the English "now" that was briefly mentioned in Chapter 2 and that is relevant to our problem here. To recapitulate, an occurrence of "now", though logically-redundent in a present-tense sentence, is not always vacuous. For example, it should be clear that the following two sentences are logically equivalent (at least as far as truth-theoretic semantics goes):

(6.10) It is raining.

(6.11) It is now raining.

That is, inserting "now" in (6.10) does not have any impact on its truth-theoretic semantics. But consider the following pair of sentences, due to (Kamp, 1971, p. 229):

(6.12) I learned last week that there would be an earthquake.

(6.13) I learned last week that there would now be an earthquake.

Evidently, (6.12) and (6.13) have different truth conditions. In particular, according to (6.13), the earthquake that I learned about is one that is supposed to happen at the time of uttering (6.13), whereas, according to (6.12), this does not have to be the case. Informally, an occurrence of "now" always refers to the time of the utterence (in ordinary discourse) even when deeply nested in a sequence-of-tense construction. As pointed out by Kamp, this property of "now" proves problematic for the standard model-theoretic semantics of tense logics, which assigns truth values to sentences at a given temporal index.[3] To overcome this difficulty, Kamp introduces a second temporal index of evaluation. The two temporal indices may be thought of as the Reichenbachian event and speech times (Reichenbach, 1947). As Kamp puts it: "The concept we ought to analyze is not simply 'the truth-value of φ at t', but rather 'the truth-value of φ at t when part of an utterence made at t''." (Kamp, 1971, p. 238)

How is all this relevant to our problem? To see this, first consider how the temporal logic adopted here manages to represent the simple dependence of truth values on time. Instead of relativizing semantic interpretation to a semantic temporal index, we move this index into the syntactic realm. The function of "Holds" is to build the temporality of formulas into the logic, by explicitly associating a formula with a time at which it is true. Now, it should be evident that part of the difficulty we are having with acting rules like (6.4) is that the rule should fire only if its condition is asserted to hold in the present. In other words, it should fire only if the condition is asserted to "hold at t", where t is the very same time of the assertion. Kamp's proposal suggests that a temporal index corresponding to the time of the utterence (or, in general, the time of the assertion) is needed to properly account for the semantics of "now". Thus, the same reasoning underlying the introduction of "Holds" may be extrapolated to include Kamp's second temporal index into the syntax. In particular, we need to formally *stamp* each assertion with the time at which it was made. In the $\mathcal{FOCS}$ system, we may introduce a function symbol, Asserted, that denotes a function from propositions

[3]Technically, if φ is a tense-logical formula and v is a semantic valuation function, then $v(\varphi)$ is, not a truth value, but a function from times to truth values.

and times to propositions. Thus, "$\mathsf{Asserted}(p,t_a)$" denotes the proposition that, at $[\![t_a]\!]$, Cassie came to believe $[\![p]\!]$. We then replace (6.4) by (6.14):

(6.14) $\forall t$ WHEN $\mathsf{Asserted}(\mathsf{Holds}(cond,t),t)$ DO act

That is, Cassie would only perform *act* when she comes to believe that *cond* holds, at a time at which it actually holds. This will indeed not match any assertions about past times and apparently solves the problem. However, there are at least two major problems with this proposal.

1. The rule in (6.14) would only work if its condition is directly introduced into Cassie's belief space, through perception, proprioception, or direct assertion. However, it would not work if *cond* is *inferred* to be holding. The reason is that propositions of the form $\mathsf{Asserted}(\mathsf{Holds}(cond,t),t_a)$ can never be inferred. In particular, introducing the assertion time results in complications with simple implications as that in (6.1), through which inferences are made. Due to its semantics, the assertion time of the antecedent need not be that of the consequent; Cassie may come to believe in *ant* at t_1 and infer *cq* later at t_2. The problem is that the time at which the inference is made cannot be known in advance. Essentially, this is the same problem that we started with; we only know that the inference will be made at some unmentionable future *NOW.

2. So far, we have only discussed the problem in the context of forward chaining. The same problem also emerges in some cases of backward reasoning in the service of acting. For example, Cassie might have a plan for crossing the street. Part of the plan may include a conditional act: *"If the walk-light is on, then cross the street"*. Note that this is a conditional act, one that involves two things: (i) trying to deduce whether the walk-light is on, and (ii) crossing the street or doing nothing, depending on the result of the deduction process. Evidently, to formalize the act, we have the same difficulty that we have with (6.4). Using the assertion time proposal, an assertion mentioning the act might be represented as shown in (6.15), where "ActIf" is a conditional-act forming operator, such that its second argument, an act, is performed if the first argument, a state, holds:

   (6.15) $\forall t[\ldots\ \mathsf{ActIf}(\mathsf{Asserted}(\mathsf{Holds}(\mathsf{On}(\text{walk-light}),t),t),\mathsf{Cross}(\text{street}))\ \ldots]$

However, attempting to deduce $\mathsf{Asserted}(\mathsf{Holds}(\mathsf{On}(\text{walk-light}), t), t)$ will succeed even if t matches some past time, t_0, at which it was asserted that the walk-light is on. The problem is that a proposition of the form $\mathsf{Asserted}(\mathsf{Holds}(cond, t), t)$ merely states that, at time $[\![t]\!]$, an assertion about the present was made. However, it does not, in itself, say anything about whether t is *now* the present time—the elusive piece of information we are trying to capture. That (6.15) works (ignoring the diffculty discussed in 1, that is) is not due to any feature of the representation proper, but because it is assumed that assertions are made with forward inferece. Thus, (6.15) would only fire due to an assertion that was just made, and it is because of this meta-logical assumption that the time mentioned in the assertion is indeed the current time. As illustrated above, once the situation is one of backward, rather than forward, reasoning, introducing the assertion time into the picture fails to capture the deictic nature of acting.

6.3 A Solution

What is the problem? At its core, the problem is that we need to make some assertions about future acts that necessarily refer to unmentionable future *NOWs at which conditions initiating, or enabling, those acts are required to hold. Those *NOWs would only be known at the time of acting. Even their being *future* is not something absolute that we know about them; they are only future with respect to the assertion time. We somehow need to introduce *NOW only when it is known—at the time of acting. The proposal is to eliminate reference to time in rules like (6.4) (or acts like that in (6.15) for that matter) and let the inference and acting system introduce *NOW when it is using these rules (or executing these acts). Thus, instead of (6.4), we shall use (6.3) (repeated below for convenience) for both cases where *cond* is eternal or temporary and build temporality into the underlying inference and acting system.

(6.3) WHEN *cond* DO *act*

Figures 6.1 and 6.2 outline modified forward and backward chaining procedures. The input to these procedures is a set of states S. Note that *NOW is inserted by the procedures themselves at the time of reasoning. This guarantees picking up the appropriate *NOW.

Algorithm `Forward`$(S \subseteq \Psi(\mathcal{S}))$

1. For all $s \in S$
   2. Perform usual forward chaining on s.
   3. If there is $s' \in$ TEMP $\cup$ PERM such that $s =$ Holds(s',*NOW) then `Forward`$(\{s'\})$.

Figure 6.1: Modified forward chaining procedure.

Algorithm `Backward`$(S \subseteq \Psi(\mathcal{S}))$

1. For all $s \in S$
   2. Perform usual backward chaining on s.
   3. If $s \in$ TEMP $\cup$ PERM then `Backward`$(\{$Holds(s,*NOW)$\})$.

Figure 6.2: Modified backward chaining procedure.

Assertion	**Assertion Time**
(6.16) WHEN Sounds(alarm) DO Leave(building)	t_1
(6.17) Holds(Sounds(alarm), t_0)	t_2
(6.18) Holds(Sounds(alarm), t_3)	t_3

Table 6.2: Fire-alarm scenario for the modified chaining procedures.

Going back to the fire-alarm example, consider the timed sequence of assertions in Table 6.2 (which is a variant of Table 6.1). First, note that (6.17) does not match (6.16) and hence the act of leaving the building will not be activated by step 2 of the `Forward` procedure. Second, since t_0 is not identical to *NOW (t_2), the recursive call to `Forward` in step 3 will not be performed. Thus, Cassie will, correctly, not leave the building just because she is informed that the fire-alarm sounded in the past. On the other hand, at t_3 the fire-alarm actually sounds. Still, (6.18) does not match (6.16). However, since t_3 is itself *NOW, step 3 results in `Forward` being called with "{Sounds(alarm)}" (which matches s') as an argument. By step 1, this will match (6.16) resulting in Cassie, correctly, leaving the building.

Similarly, we may replace the representation of the conditional act in (6.15) by (6.19):

(6.19) ActIf On(walk-light) THEN Cross(street).

If Cassie is told to perform this conditional act at t_1, the procedure `Backward` would be called with "{On(walk-light)}" as an argument. Since On(walk-light) is a temporary state, backward chaining will be performed on "Holds(On(walk-light),*NOW)" (step 3), thus querying the knowledge base about whether the walk-light is on at t_1, the time we are interested in.

6.4 On Complex States

The proposal outlined above basically dictates that the representation of acting rules (along the lines of schema (6.3)) and conditional acts (as in (6.17)) should not include any mention of the time at which the conditions are required to hold. The appropriate time is introduced as needs be by the inference procedures in Figures 6.1 and 6.2. The assumption, though, is that time can always be discharged. Is this assumption valid? Without loss of generality, let us concentrate on (6.3) for the rest of the discussion. In the fire-alarm example presented above, discharging time poses no problems as illustrated by tracing the algorithms with the scenario in Table 6.2. But that example is particularly simple, and its simplicity stems from the fact that the condition, Sounds(alarm), is a simple state (see Definition 3.1). More specifically, (6.16) is expected to fire whenever the fire-alarm is asserted to be sounding. Since the fire-alarm's sounding denotes a TEMP state, such an assertion is of the form

(6.20) $\mathsf{Holds(Sounds(alarm),{}^*NOW)}$[4]

With such an assertion, discharging *NOW gives us the simple state Sounds(alarm), which matches the condition part of (6.16) resulting in appropriate action. But suppose that we want Cassie to follow a slightly more complicated rule ("more detailed" would be a better qualification). Here is an English example:

(6.21) When the fire-alarm sounds and you are in the building, then leave the building.

Logically, the condition in this case is a conjunction:

(6.22) $\mathsf{Holds(Sounds(alarm),{}^*NOW)} \wedge \mathsf{Holds(In(I,building),{}^*NOW)}$

Note that each conjunct may in fact be asserted separately, but both are required for the rule to fire. The question now is how to discharge *NOW from a conjunction like the above one and, more importantly, how to represent the rule so that asserting such a conjunction fires it. It is mainly because of these questions that I spent some time in Sections 3.5.4 and 3.5.5 extending the use of the standard propositional connectives and quantifiers to the entire domain of states. In particular, given **AS5**, we have:

(6.23) $$[\mathsf{Holds(Sounds(alarm),{}^*NOW)} \wedge \mathsf{Holds(In(I,building),{}^*NOW)}] \Leftrightarrow \mathsf{Holds([Sounds(alarm)} \wedge \mathsf{In(I,building)],{}^*NOW)}$$

Thus, rather than forward chaining into the rule using (6.22), we can use the right-hand side of (6.23) from which *NOW may be discharged similar to (6.20). The rule may, therefore, be represented as follows, resulting in the correct behavior given the the NOW-sensitive forward chaining algorithm of Figure 6.1:

(6.24) $\mathsf{WHEN\ [Sounds(alarm)} \wedge \mathsf{In(I,building)]\ DO\ Leave(building)}$

However, things are not that simple if the condition part of an acting rule involves negation. For example, once Cassie leaves the building as a result of hearing the fire-alarm, it is reasonable for her to adopt the following rule:

[4] Note that the assertion cannot just be "Sounds(alarm)". This is not just because β is defined to be a set of propositions rather than one of states; we can always change the definition. The reason is that, not only do we want Cassie to act upon knowledge of the fire-alarm's sounding, but we also want her to *remember* when it did.

(6.25) When the fire-alarm is not sounding, return to the building.

This rule may be represented as follows, but the representation needs some justification:

(6.26) WHEN $\neg$Sounds(alarm) DO ReturnTo(building)

The problem is that, in the case of negation, we do not have a nice bi-directional implication similar to **AS5**. In particular, the following is not correct (see Section 3.5.4):

(6.27) $\neg\mathsf{Holds}(Sv, Tv) \Rightarrow \mathsf{Holds}(\neg Sv, Tv)$

Since we cannot move "$\neg$" into the scope of "Holds", asserting "$\neg$Holds(Sounds(alarm),*NOW)" would not result in firing (6.26). Nevertheless, I maintain that (6.26) is the only possible representation of the rule, which means that assertions where negation falls outside the scope of Holds would not be effective. This is not as bad as it may seem; there are a couple of points to note:

1. By **TS3**, a state fails to hold over some interval if and only if its complement holds over a sub-interval thereof. Thus, unless $[\![\neg\mathsf{Sounds(alarm)}]\!]$ holds throughout $[\![t]\!]$, "$\neg$Holds(Sounds(alarm), t)" might mean that $[\![\mathsf{Sounds(alarm)}]\!]$ holds over some sub-interval of $[\![t]\!]$ (see Figure 3.7). Clearly, this is not the intuition behind (6.26). In general, whenever the condition of an acting rule involves negation, *strong* negation, with "$\neg$" having the narrowest possible scope, is intended. Thus, strictly speaking, "$\neg$Holds(Sounds(alarm),*NOW)" should *not* result in firing (6.26); rather, asserting "Holds($\neg$Sounds(alarm),*NOW)" should, and, in fact, would so result, given algorithm `Forward`.

2. We are interested only in assertions about a state failing to hold "now". In the meta-theory, NOW-intervals are atomic and cannot be divided into two sub-intervals such that the state holds over one and its complement over the other. Thus, given **TS3**, as far as Cassie is concerned,

   $$\neg\mathsf{Holds(Sounds(alarm),*NOW)} \Leftrightarrow \mathsf{Holds(\neg Sounds(alarm),*NOW)}$$

   The problem, of course, is that this is only in the meta-theory; Cassie cannot infer it. However, practiccally, it means that, for any state, its failing to hold "now" must mean that its complement holds, which leads to the next point.

3. Where does negative information about states come from? All information comes from the four sources discussed before: perception, proprioception, direct assertion, and inference. However, as pointed out in the discussion of the fourth principle of change (Section 5.2.5), perception and proprioception only convey positive information about states. Thus, Cassie would believe that "$\neg$Holds(s,*NOW)" if she determines through perception or proprioception that some other state, s', holds. In a sense, s and s' are contradictory states. However, it is hard to imagine examples where determining that s' holds would imply that s does not hold but would not imply that $\neg s$ holds. What I am suggesting here is that it is the responsibitlity of the knowledge engineer to make sure that the right thing is inferred: "Holds($\neg s$,*NOW)" versus "$\neg$Holds(s,*NOW)" (Axiom 5.8 indeed mandates inferring the "right thing" for bodily states). In commonsense reasoning, the former, strong notion of negation is almost always what people have in mind (witness how unnatural it is to think of the situation depicted in Figure 3.7).

   Now, this takes care of inferences about states not holding initiated by perception or proprioception, but what about direct assertion? It is not possible to control what someone may tell Cassie. That is, someone may, for example, assert "$\neg$Holds(Sounds(alarm),*NOW)" expecting this to cause (6.26) to fire, when it actually would not. There are two replies to this. First, as pointed out in point 1 above, (6.26) should indeed not fire in that case, and it is the responsibility of whoever chooses to directly assert statements into Cassie's belief space to make sure that they are precise enough and say the right thing. Second, as I repeatedly said, I envision direct assertion to be done in natural language. In that case, the parsing system should translate English sentences into the appropriate $\mathcal{FOCS}$ expressions. In particular, an English sentence of the form "*s does not hold now*" should be directly translated into "Holds($\neg s$,*NOW)", which is justified by point 2 above.

Since we only need to account for $\wedge$ and $\neg$ to generate arbitrarily complex states, we can represent complex rules along the lines of (6.24) and (6.26). For example, (6.28) may be represented by the formal rule in (6.29):[5]

(6.28) When the fire-alarm sounds and no children are in the building, then leave the building.

[5]In Chapter 9, I present a more precise characterization of the WHEN ... DO ... construct within the $\mathcal{FOCS}$ system.

(6.29) $\mathsf{WHEN}\ [\mathsf{Sounds}(\mathsf{alarm}) \land \forall Xv[\mathsf{Child}(Xv) \Rightarrow \neg\mathsf{In}(Xv, \mathsf{building})]]\ \mathsf{DO}\ \mathsf{Leave}(\mathsf{building})$

6.5 Conclusions

Time matters. And it matters not only as an object of reasoning for acting agents, but as a factor that the inference procedures should take into account. This stems from two features, one of time itself, namely that time inevitably passes as reasoning and acting take place, and one of the particular approach to time that I adopt. It has been argued before (Lespérance and Levesque, 1994; Lespérance and Levesque, 1995) that cognitive agents need to have some notion of indexical time (in particular, a notion of "now"). This fundamental idea is one that underlies the general approach to cognitive robotics presented here. However, the particular approach to implementing indexicality in the system is different from that adopted by (Lespérance and Levesque, 1994; Lespérance and Levesque, 1995) (also see (Schmiedel, 1990; Artale and Franconi, 1998; Dix et al., 2001)). In those systems, "now" is represented by a term in the knowledge-representation language, and semantics is given relative to time indices. As noted in Chapter 2, the problem with that approach is that it essentially involves unnatural knowledge-base updates if the system embodies a notion of the passage of time. The approach presented here introduces "now" only in the meta-theory, not in the knowledge-representation language. The problem of the unmentionable now reveals the cost of doing this: "now" has to be built into the reasoning system. One way to view this is that "now" only appears in the *semantics* of the knowledge representation language, not in the syntax. This, in fact, is probably better since the indexical "now" of, for example, (Lespérance and Levesque, 1995) seems to function primarily as a *place-holder* and does not have the distinctive semantic properties of the English "now" (in particular, the property discussed in Section 6.2).

Chapter 7

The Problem of the Fleeting Now

This chapter discusses another temporal progression problem, *the problem of the fleeting now*. The problem was first pointed out by (Ismail and Shapiro, 2000b). The gist of the problem is that reasoning takes time. It emerges when the agent is reasoning about "now", when the very process of reasoning results in "now" moving, and thereby fleeting from the agent's *mental grip*.[1] A solution to the problem involves, not only *now-sensitive* inference procedures as with the problem of the unmentionable now, but also endowing Cassie with a *feel* for how much time has passed. In the first three sections, the problem and the proposed solution are presented in informal and general terms. The rest of the chapter investigates the issues involved within the formal framework of Chapter 5.

7.1 The Vagueness of "now"

> Salesman:*Good morning Madame; may I interest you in a fine . . . ?*
>
> Woman (annoyed): *Not now.*
>
> Salesman: . . . *Now?*

What is funny about the above joke is probably the realization that the salesman's reply, though obviously silly, is in some sense valid. The woman's *"not now"* is primarily a show of disinterest but may also be taken to include an implicit invitation for the salesman to approach her later, at some different "now". The salesman's response is silly because our intuition is that it came at the

[1] Interestingly, this is the gist of the uncertainty principle in quantum mechanics.

same “now”. Nevertheless, there is also a sense in which the response is appropriate, for, strictly speaking, whenever time passes there is a new “now”.

What this all means is that we maintain different notions of “now” at the same time. Specifically, there is a set of time intervals, linearly-ordered by the subinterval relation, representing the concept of “now” at different levels of granularity. What an occurrence of “now” refers to is some element of this set; which one in particular is determined by contextual and pragmatic factors. In the above joke, the salesman conveniently interprets “now” at a level of granularity that is finer than the one picked by the woman. He does that by, probably, choosing to ignore some salient pragmatic factors.[2]

A similar problem raised by (Quine, 1960) may help drive the point home. Quine delves into a long discussion of *radical translation*: “translation of a language of a hitherto untouched people” (Quine, 1960, p. 28). Quine discusses a certain difficulty that a linguist encounters when an accompanying native informant points in the direction of a rabbit and says: “Gavagai”. The difficulty is that the mere act of pointing toward a rabbit does not necessarily mean that “Gavagai” is the translation of “Rabbit”. For as far as the linguist can tell, the informant may be pointing, not to the rabbit as a whole, but to a part thereof (Quine, 1960, pp. 51–54). The pointing action is ambiguous as to which of a set of successively larger regions of space it is supposed to cover. Similarly, an utterance of “now” *points* to the present time, and is vague since there is always a nest of intervals containing the utterance event.

This vagueness of “now” has seldom been noted in the literature. (Allen, 1983, pp. 840–841) too argues for a hierarchy of intervals to represent the present. However, his argument is primarily motivated by the desire to enhance the computational efficiency of a temporal reasoning system, rather than by genuine reasoning problems. It might seem that even though “now” is vague, this vagueness does not pose serious problems to language understanding. One reply is to witness the silliness of the salesman’s response above, which stems from an inappropriate resolution of

[2]The issue is not just one of jokes. A real-life example has been reported by Bill Rapaport (in personal communication). He documented the following December 1999 interaction which he had with his, then 4-year-old, son Michael on their way to McDonald’s.

> Michael:*Are we almost there?*
> Bill:*We’re about half way.*
> Michael:*How about now?*
> Bill:*We’re still half way.*
> Michael:*How about now?*

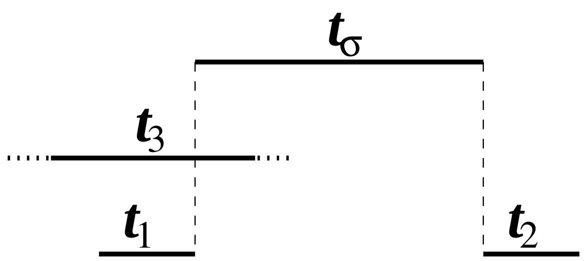

Figure 7.1: The temporal structure of the fleeting now problem.

the vagueness. However, I am not primarily concerned with linguistic problems here. The main problem that I address is one that faces a reasoning and acting agent that maintains a sense of the present time. This problem, dubbed *the problem of the fleeting now*, is a direct result of ignoring the vagueness of "now".

7.2 The Fleeting Now

The time is t_1. Consider an agent that needs to know whether some state, s, holds *now*. To achieve its goal, the agent performs some sequence of actions, σ. Performing σ results in the agent observing that s indeed holds *now*. Naturally, σ takes time, and the observation is made at some time, t_2, that is later than t_1. Thus, strictly speaking, the agent observes that "s holds at t_2", whereas its original concern is whether "s holds at t_1". The problem, of course, is that, for the temporally-fanatic agent, the question is about one *now* (henceforth now_1) while the answer is about another (now_2). Just as in the case of the salesman, the agent picks the wrong level of granularity at which to interpret now_1. What is needed is a level relative to which now_1 and now_2 are *indistinguishable* (Hobbs, 1985). That is, now_1 should be interpreted as a coarser interval, t_3, which includes both t_1 and t_2. Thus, now_1 would be the same as now_2.

Figure 7.1 depicts the situation, where t_σ is the time interval over which σ is performed. As the figure shows, the only thing that we definitely know about t_3 is that it is a super-interval of t_1, i.e., it is a possible interpretation of now_1. For t_3 to also be a super-interval of t_2, it must survive t_σ. Since the nest of "now"s is potentially infinite, one can always find a *now* that is large enough to persist beyond t_σ. However, there is another restriction on the possible candidates for now_1. This is where knowledge about the state s comes into play. People have intuitions about the typical lengths of intervals over which various states persist. For example, walk-lights are on for about 15 to 30 seconds, meals are typically eaten in 15 to 30 minutes, conferences are held for several days, vacations last for a few weeks, and so on.[3] However now$_1$ is interpreted, it should be interpreted at a level of granularity that is neither too fine nor too coarse for s. In metrical terms, t_3 should be neither too short, nor too long, for the typical duration of s. If it is too short, the fleeting-now problem emerges. On the other hand, if it is too long, then observing s at t_2 cannot be used (at least defeasibly) to conclude that it holds throughout t_3.

The following spatial analogy may help fix the idea. Think of the nested intervals that are candidates for the interpretation of now$_1$ as some sort of *cognitive ropes* with various lengths. These ropes are fixed from one end (the left end, assuming time flows from left to right) at t_1. They are also somehow elastic, so that they could be stretched, within a certain limit, beyond their minimum lengths. When asked whether s holds *now*, the agent mentally picks one of these ropes. In particular, it picks one whose length is comparable (maybe within a half order of magnitude (Hobbs, 2000)) to the typical duration of s. The agent then moves to the right, holding the loose end of the rope, until at some point, t_2, it sees that s holds. If the agent is still holding the rope, it may answer the question affirmatively. If, on the other hand, it has run out of rope along the way from t_1 to t_2, it would not be able to conclude whether s holds at now$_1$. Thus, for the answer to be "yes", the agent needs to be still holding the rope when it sees s. In addition, it should have picked the *right* rope in the first place.

For example, suppose that at 12:15 p.m. sharp an agent is asked: *"Is John having lunch now?"*. To answer the question, the agent walks to John's office, where it sees John busy munching something, with a sandwich in his hand. Assume that this perception event takes place at 12:17. Since now$_1$ is interpreted at a level of temporal granularity appropriate for "having lunch", its approximate duration is certainly longer than the 2-minute span of t_σ (the time it took the agent to walk to John's office). Thus, the agent may go back to the questioner and answer affirmatively. On the other hand, suppose that John is not in his office, but unbeknown to the questioner, is at home. The agent reasons that if it drives to John's house, it will arrive there at a *now* that is different from now$_1$, even when interpreted at the coarse lunch-granularity level (i.e., it will run out of rope). In this case, the agent has no way of answering the question, at least as long as it insists on adopting

[3] See (Hobbs, 2000) for an interesting discussion of the mathematical properties of typical measures.

the driving-to-John's plan.

The above intuitions may be easy to state. But how would one formalize them for cognitive robotics applications? How can we give an embodied agent, not only a sense of the passage of time, but also a *feel* of how much time has passed? The rest of the chapter attempts to answer these questions.

7.3 Life without a Clock

How do clocks keep track of time? "A clock is basically an instrument for creating a periodic process and counting the periods" (Dowden, 1998).[4] Thus, the amount of time elapsed is a function of the number of periods counted. In particular, since clock-periods have equal lengths, the amount of time elapsed is simply the number of counted periods multiplied by some constant period length. But agents do not always have access to clocks. For some people, forgetting to wear their watch before leaving to work in the morning may become a really disturbing experience. Throughout the day, they would have a feeling of loss and disorientation because of their inability to precisely keep track of time. Yet, despite these feelings, people do not totally lose their sense of time. They can still behave appropriately, estimate how much time it took to type a report, how long a meeting has lasted, and whether it is time to go home. How is this possible without a clock? Are there other means by which people perceive the passage of time? Certainly; the psychology literature provides various models of time perception (see (Friedman, 1990) for a general discussion). These are categorized into two major categories: biological models that hypothesize the existence of some internal clock or what is referred to as the *pacemaker* (see (Pastor and Artieda, 1996)), and cognitive models that assume its absence (see (Levin and Zakay, 1989)).

Computationally, a pacemaker is simply a process that generates equally-spaced *ticks*. (Ladkin, 1986) discusses a system that uses the workstation clock for representing the progression of time, where *now* is interpreted as the time indicated by the clock. Similar methods may also be adopted. It should be noted, however, that, in general, what is needed is some method to *measure* the amount of time elapsed; the exact relation between ticks of the pacemaker and successive values of *now* need not be a one-to-one correspondence. The ticks of the pacemaker are best thought of as providing,

[4] Of course, this only applies to the familiar (analog or digital) clocks, not to hour-glasses, for example.

not absolute dates, but a feel of the temporal proportions carved out by different events.

However, pacemaker-theories do not explain most of our everyday experiences of time.[5] If there is no internal clock, how else do we get a feel of the duration of what happens? Basically, we do it by knowledge of what *else* happens. For example, (Poynter, 1989) claims that the number of sensory and mental events filling a time interval is one major factor that provides the feel for its duration. He further argues that "whether an event turns out to be a useful marker of time passage depends on the length of time which is to be remembered" (Poynter, 1989, p. 312). For example, if one is interested in a duration on the scale of hours, then knowledge of events happening on the same scale is more helpful than knowledge of those happening on the scales of seconds or years. What this suggests is that it is not just the mere *number* of events filling an interval that matters, but that intuitions about typical durations of events provide rough metrics for measuring time.

The theory of time developed in Chapter 5 and in the following sections is flexible enough to accommodate either of the above views. In Section 7.4, I shall illustrate how the model may be used in conjunction with the cognitive theory of duration judgment. For technical reasons, however, I will present a fuller, more elaborate discussion of a pacemaker-based theory. Nevertheless, it should be noted that this is a tactical, rather than a strategic, decision.

7.4 The Perception of Time

The moving NOW provides Cassie with some sense of temporal progression. Nevertheless, the perception of time is not only confined to distinguishing a present moment or a chain of NOW-MTFs; it crucially involves a feel of the *amount* of time taken up by different intervals. In this section, I will examine alternative means by which we can model this aspect of time perception.

7.4.1 Amounts and Durations

If Cassie is to have beliefs and reason about durations of time intervals and typical durations of states, the $\mathcal{FOCS}$ system must be extended to accommodate such notions. In particular, a new sort with a suitable suite of relations is introduced.

[5]See (Friedman, 1990, ch. 2) for a lengthy discussion and references.

Sort. There is a sort, $\mathcal{Q}$, where $\mathcal{D}_\mathcal{Q}$ is a set of amounts. I use q, q', and q_n ($n \in \mathbb{N}$) as meta-variables ranging over elements of $\mathcal{Q}$.

Constants. Qc, Qc', Qc_n ($n \in \mathbb{N}$) are constant symbols in $\mathcal{Q}$.

Variables. Qv, Qv', Qv_n ($n \in \mathbb{N}$) are variables in $\mathcal{Q}$.

Function Symbols. Only three functions are needed for our purposes.

- SDur: $\textsc{Temp} \times \mathcal{Q} \longrightarrow \mathcal{P}$, where $[\![\mathsf{SDur}(s,q)]\!]$ is the proposition that the typical duration of intervals over which $[\![s]\!]$ *maximally* holds is $[\![q]\!]$.
- Dur: $\mathcal{T} \times \mathcal{Q} \longrightarrow \mathcal{P}$, where $[\![\mathsf{Dur}(t,q)]\!]$ is the proposition that the duration of interval $[\![t]\!]$ is $[\![q]\!]$.
- $<_\mathcal{Q}$: $\mathcal{Q} \times \mathcal{Q} \longrightarrow \mathcal{P}$, where $[\![q <_\mathcal{Q} q']\!]$ is the proposition that amount $[\![q]\!]$ is less than amount $[\![q']\!]$.

Axiom. Without stating them, I assume the existence of axioms defining $<_\mathcal{Q}$ as a strict linear order.

7.4.2 Knowing and Feeling

Before setting out to discuss alternative models, we need to make a distinction between two concepts: the *knowledge* of a duration, and the *feel* thereof. Cassie's knowledge of the duration of some time interval, t, is represented by a proposition, $\mathsf{Dur}(t,q)$, in β, where q is the amount of time representing the duration of t. Such knowledge may come from any of the three sources of information: direct assertion, inference, and bodily-feedback (i.e., perception and proprioception). There are no restrictions on whether t is in a NOW-MTF or a non-NOW-MTF. Thus, Cassie may have knowledge of the durations of events that she did not witness, but merely heard of from a third person. On the other hand, a feel of the duration of an interval requires a first-person experience of that interval. Formally, Cassie has a feel of the duration of an interval t only if every MTF in Span(t) is a NOW-MTF,[6] that is, if every piece of t is in Cassie's chain of consciousness.

[6]But see below for a minor qualification of this statement.

If knowledge of durations is represented by conscious beliefs, what is a *feel*, and how could it be represented? An analogy from the good old domain of color cognition may help. What do we know about colors? Two things. First, we have a mental entity that may be associated with a *name* for the color: *"red"*, *"blue"*, *"green"*, etc. Second, if we have *seen* an object of that color, we would have a *feel*, or perceptual experience, associated with the mental entity. In case we do not have a name for a color, the only possible way to express it is to use variations of "*the color of this or that object*". In our formal framework, the name of the color resides at the KL, and the feel, or perceptual experience, exists at the PML. Similarly, we have two corresponding notions for time intervals: mental entities, represented by symbols in $\mathcal{T}$; and perceptual experiences, represented by PML structures.[7] Associations between KL and PML structures are represented by an *alignment set* (Shapiro, 1998). Essentially, the alignment set, $\mathcal{A}$, is a set of pairs, $\langle \tau, \pi(\tau) \rangle$, where τ is some KL term and $\pi(\tau)$ is the PML structure representing the perceptual experience, or feel, associated with τ. Thus, whereas knowledge of the duration of an interval, t, is represented by a belief in β, a feel for that duration is represented by a pair $\langle t, \pi(t) \rangle$ in $\mathcal{A}$. Of course, Cassie may have both a feel, $\langle t, \pi(t) \rangle$, and a belief, $\mathsf{Dur}(t, q)$, about the duration of t. It is the duty of PML recognition processes to generate the appropriate q, given some $\pi(t)$. I assume the existence of a mapping, ρ, that would map $\pi(t)$ to the appropriate term in $\mathcal{Q}$ (see Figure 7.2, where δ maps a time interval to the duration associated with it in β). The exact characterization of ρ depends on the nature of PML representations and the interpretation of elements of $\mathcal{Q}$.

7.4.3 The Contents of an Interval

The basic claim of the cognitive theories of duration perception (Levin and Zakay, 1989, for instance) is that the contents of an interval provide the feel for its duration. The more events happen within an interval, the longer it feels. The framework developed in Chapter 5 readily provides such a measure; the larger $|\text{Span}(t)|$, the longer t feels (see Definition 5.8). Thus, $|\text{Span}(t)|$ seems to provide a pretty good measure for the duration of t as dictated by cognitive theories. Nevertheless, there are two problems with such an approach. First, the contents of a duration in cognitive theories include both external events as well as internal mental events. According to the Second Principle

[7] I have to explicitly state that, by using terms like "feel" or "feeling" in the sense given above, I am not trying to make any claims about the phenomena traditionally associated with those terms in philosophy and psychology. As far as this work goes, to *feel* something is to have a PML representation of it.

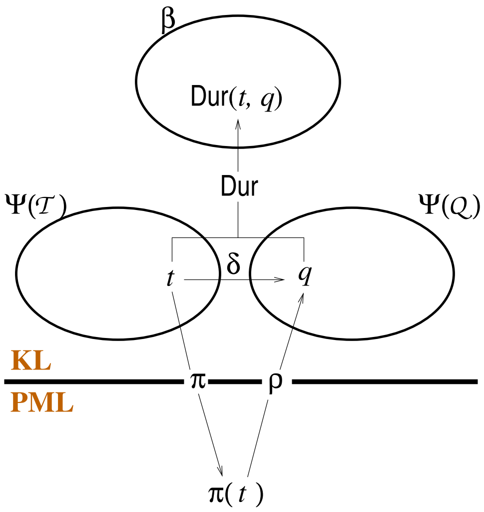

Figure 7.2: Mappings across the KL-PML interface.

of Change, NOW moves only when a change that Cassie is *aware* of takes place. As pointed out in Section 5.1.2, Cassie is not aware of internal changes caused by mental events, such as inference, only of external environmental events. Thus, $|\text{Span}(t)|$ is a very rough and inaccurate measure of duration. Second, even if Cassie is aware of mental events, unlike humans, an artificial agent is not always busy-minded. Unless it is occupied with a particular reasoning task, typically triggered by some external event (including a query by a human operator), an artificial agent is "unconscious" most of the time. Thus, $\text{Span}(t)$ would still be rather empty, and an inaccurate measure of the duration of t.

Even more important, note that relative durations are not maintained by $|\text{Span}(t)|$. Suppose that an interval t_1 is actually (i.e., for us, humans) much longer than an interval t_2. Using spans as measures for duration may make Cassie come to believe that t_2 is longer than t_1 just because she became aware of more events during the former. Now one may argue that this is actually reasonable. For what we want to represent is *Cassie's* sense of time, not her knowledge of ours. Thus, for Cassie, t_2 is *actually* longer than t_1. There are however two responses. First, one thing that we would like to represent is temporal regularity: that all occurrences of some event take almost the same amount of time to happen. This kind of regularity is lost if the feel for the durations of those occurrences is dependent on something as irregular as the changes that Cassie becomes aware of during them. Second, artificial agents typically operate in environments where they interact with humans. For this interaction to be effective, something as basic as the sense of time, should be unified. If Cassie were to operate in an environment inhabited by agents with similar perceptual capacities, and where *all* the relevant changes are perceivable, then $|\text{Span}(t)|$ may be a good measure of durations. However, as long as artificial agents operate among us, humans, they need to be adapted to *our* environment.

7.4.4 The Pacemaker

One problem with using the span of an interval as a measure for its duration, is that spans only represent the number of changes that happened during the interval; they do not carry any information about the amounts of time between those changes. It seems that if there is some means by which we can characterize those amounts, then the duration of an interval is simply the sum of all amounts of time between changes that happened within it. Note that the amount of time between two changes is essentially the duration of an MTF or the smallest element thereof.

Thus, the base case is the feel of durations of atomic intervals. In particular, the durations of NOW-intervals, since one can only have a feel of a duration if all the MTFs in its span are NOW-MTFs (see Section 7.4.2). The question then is how to account for this base case. What is it that may provide the feel for the duration of a NOW-interval? Note that one cannot resort to the number of changes, since NOW-intervals are atomic. We are, therefore, left with only one option: some sort of a pacemaker. Our pacemaker is a PML process, essentially a counter that starts counting once the agent comes to life (i.e., starts operating) and is reset every time NOW moves. More specifically, a PML process periodically increments the integer value of a meta-theoretical (PML) variable, COUNT, and a revised version of algorithm `move_NOW` (see Figure 7.3) resets it. The revised `move_NOW` aligns *NOW with the number of ticks produced by the pacemaker since the last time NOW moved (step 3), thus providing a feel for its duration. Figure 7.4 shows a similarly revised version of algorithm `initialize_NOW`. There are a couple of things to note.

1. Ticks of the pacemaker at the PML, do not correspond to NOW-intervals at the KL. The dynamics of NOW is still governed by Axioms 5.4 and 5.5. The pacemaker merely provides the feel for the duration of NOW-intervals.

2. The rate at which COUNT is incremented is not significant so long as it is (i) constant, and (ii) fast enough to provide different feels for intervals whose durations need to be distinguished

Algorithm `move_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{\text{*NOW} \prec t, \mathsf{Dur}(\text{*NOW}, \rho(\text{*COUNT}))\}$.
3. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle \text{*NOW}, \text{*COUNT} \rangle\}$.
4. $\mathsf{COUNT} \longleftarrow 0$.
5. $\mathsf{NOW} \longleftarrow t$.

Figure 7.3: A revised version of `move_NOW` that aligns NOW-intervals with PML structures representing Cassie's sense of their duration.

Algorithm `initialize_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\mathsf{COUNT} \longleftarrow 0$.
3. $\mathsf{NOW} \longleftarrow t$.

Figure 7.4: A revised version of `initialize_NOW` that resets the PML variable COUNT.

as dictated by the domain of application.

Now, let us take a careful look at step 3 of algorithm `move_NOW`. A new pair is added to the alignments set, $\mathcal{A}$, thereby extending it. Recall that a pair $\langle x, y \rangle$ in $\mathcal{A}$ represents an association between a KL term and a PML structure representing its perceptual, or bodily, experience. I assume that, arguably, if $\langle x, y_1 \rangle \in \mathcal{A}$ and $\langle x, y_2 \rangle \in \mathcal{A}$, then $y_1 = y_2$. Thus, the set $\mathcal{A}$ may be thought of as an extensional representation of a function—a partial function from KL terms to PML structures. In this case, we can use $\mathcal{A}(\tau)$ to refer to the PML structure associated with the term τ in $\mathcal{A}$. But recall that, in Section 7.4.2 (also see Figure 7.2), I have indicated that pairs in $\mathcal{A}$ are of the form $\langle \tau, \pi(\tau) \rangle$, where π is a function that maps τ into its perceptual experience. Is there a difference between π and $\mathcal{A}$ (conceived as a function)? For most sorts of KL terms, π and $\mathcal{A}$ are indeed identical, but for terms in $\mathcal{T}$ there is a subtle difference between them. Following standard notation, let $\pi|_{\mathcal{T}}$ and $\mathcal{A}|_{\mathcal{T}}$ be the restrictions of π and $\mathcal{A}$ to $\mathcal{T}$, respectively. Note that $\mathcal{A}|_{\mathcal{T}}$ (and, in general, $\mathcal{A}$) is extensionally-defined, that is, $\mathcal{A}|_{\mathcal{T}}(t)$ is defined *only* if a pair $\langle t, \pi|_{\mathcal{T}}(t) \rangle$ is in $\mathcal{A}$. $\pi|_{\mathcal{T}}$, on the other hand, is intensionally-defined; there is an effective procedure for computing $\pi|_{\mathcal{T}}(t)$, whenever

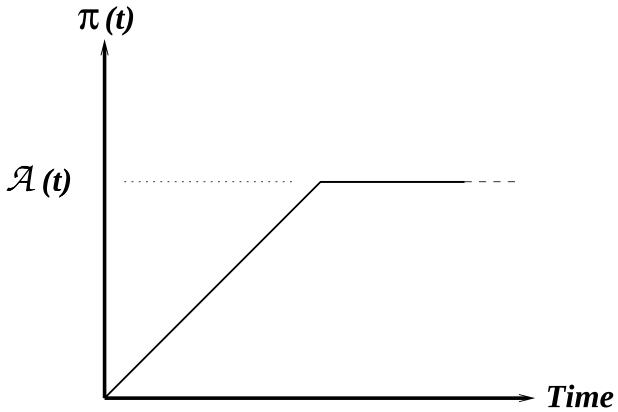

Figure 7.5: The behavior of $\pi(t)$ over time.

possible.[8] Thus, for some terms, there are occasions in which $\pi|_{\mathcal{T}}$ is defined and $\mathcal{A}|_{\mathcal{T}}$ is not.

To see what those *occasions* may be, note that the perceptual experience of an interval, the feel for its duration, evolves over time, as long as the interval has not moved into the past. Thus, $\pi|_{\mathcal{T}}(t)$ increases with time until *NOW is no longer a subinterval of t. At that point, $\pi|_{\mathcal{T}}(t)$ reaches a steady-state value which is the one that gets permanently associated with t. That is, $\mathcal{A}|_{\mathcal{T}}(t)$ *is* the steady state value of $\pi|_{\mathcal{T}}(t)$ (see Figure 7.5). Note that, just like atomicity, spans, and MTFs, $\pi|_{\mathcal{T}}$ is time-dependent. Actually, the same applies to $\mathcal{A}$, for, with time, the domain of $\mathcal{A}$ gets broader as more entities are perceived. But the dependency of $\pi|_{\mathcal{T}}$ on time is more radical; for a particular $t \in \mathcal{T}$, $\pi(t)$ changes with time as per Figure 7.5. Note that the horizontal coordinate of the origin in Figure 7.5 does not correspond to zero-time, the time Cassie comes to life. Rather, it corresponds to the time at which t is first conceived of, moving into $\Psi(\mathcal{T})$. The steady state value is reached once t expires and moves into the past. Note that $\pi(*\mathrm{NOW})$ *is* *COUNT which increases linearly with time due to the constant rate of the pacemaker. Step (3) in Figure 7.3 simply associates *NOW with the steady state value of $\pi(*\mathrm{NOW})$. For a non-atomic interval, t, $\pi(t)$ is defined according to the following equation.

$$\pi(t) = \sum_{\Phi(t_i) \in \text{ Span}(t)} \pi(t_i) \tag{7.1}$$

As long as *NOW is a subinterval of t, the value computed according to the above equation monotonically-

[8] By "whenever possible", I am stressing the partiality of $\pi|_{\mathcal{T}}$; recall that $\pi|_{\mathcal{T}}$ is defined for t only if all the members of Span(t) are NOW-MTFs.

Algorithm $\pi(\tau)$

1. If $\tau \notin \mathcal{T}$, then

   1a. If $\mathcal{A}(\tau)$ is defined, then return $\mathcal{A}(\tau)$.

   1b. FAIL.

2. If $\tau =$*NOW, then return *COUNT.

3. If $\tau =$*NOW$_i$, for some $i \in \mathbb{N}$, then return $\mathcal{A}(\tau)$.

4. If τ is a transition interval the return 0.

5. If τ is an atomic interval, then FAIL.

6. Return $\sum_{\Phi(t_i)\in \text{ Span}(t)} \pi(t_i)$

Figure 7.6: Definition of the function π.

increases with time, reflecting the increase in both $|\text{Span}(t)|$ and $\pi(\text{*NOW})$. Once t has moved into the past $\pi(t)$ reaches its steady state value, $\mathcal{A}(t)$.

Figure 7.6 sums up these ideas (and adds a couple more) in an algorithm that computes the function π for a given term τ. Note that π is necessarily a partial function, since it need not be defined for all terms. This is demonstrated in steps 1b and 5 by the algorithm's failure to compute the value of π. Note when this happens: when τ is a non-$\mathcal{T}$ term for which $\mathcal{A}$ is undefined, or when it is an atomic interval that is neither a NOW interval nor a transition interval. Also note that the value of π for a transition interval is 0, an even stronger reflection of the punctuality of transitions. Given this algorithm, it should be clear that Cassie may only have a feel for the duration of a non-atomic interval if its span includes only transition-MTFs (see Definition 5.11) and NOW-MTFs. However, it should be clear that only the NOW-MTFs count, since the duration of a transition interval is 0 according to step 4 of the algorithm.

Having pointed out the various properties of $\pi|_{\mathcal{T}}$, the mapping from KL terms to PML structures, we now direct our attention to the other direction, the recognition mapping ρ.

7.4.5 From Feeling to Knowing

First, let us examine what properties of ρ are reasonable for what we take it to represent. For one thing, it is reasonable to assume that ρ is a function; that is, it maps an element of its domain to one and only one element of its range. This is the least we can require of a *recognition* mapping. The question now is whether it is one-to-one or many-to-one. Being one-to-one implies a very *sharp* recognition function, one that allows Cassie to *consciously* distinguish between two durations no matter how similar they feel. Typically, however, conscious knowledge is much coarser than perceptual experience. For example, we use the same word, *"red"*, to refer, not to a single sharp wave length, but to a band thereof. This is not to say that our perception does not distinguish different shades or hues of red; it is just that those distinctions have no impact on reasoning and communication. Therefore, I will take ρ to be many-to-one, with elements of $\mathcal{Q}$ corresponding to ranges of numbers at the PML. Note that this is not an innovation; it is a standard practice in qualitative physics (Forbus, 1984, for example). This choice is also grounded in the interpretation of elements of $\mathcal{Q}$ as representing intuitions about *typical* amounts. Because of their inherent vagueness, typical amounts are best viewed as ranges, rather than specific values.

Another question is whether ρ should be an onto function. This would imply that every term in $\mathcal{Q}$ corresponds to some perceptual experience. Although, in principle, this may happen to be the case, I opt for the more liberal interpretation, and do not require ρ to be onto. The reason is that Cassie may conceive of amounts that she may *never* experience, due to limitations on perception. For example, even though we can reason about nano-seconds, or construct thought experiments about the speed of light, we cannot hope (at least for a while) to be able to have a direct experience of how a nano-second feels, or how traveling at the speed of light may affect our sense of time.

To precisely establish the mapping ρ, we should decide on some partitioning of the set of natural numbers into intervals that correspond to elements of $\mathcal{Q}$. Admittedly, any partitioning would have to be arbitrary unless based on psychological evidence, which, as far as I know, does not exist. Nevertheless, we can still do better than picking some random partition; one possibility is to partition the natural numbers into *half-orders of magnitude* (Hobbs, 2000). According to Hobbs, half-orders of magnitudes (HOMs) partition the positive reals into intervals geometrically-centered around $\sqrt{10}$, i.e., intervals of the form $[\sqrt{10}^{h-\frac{1}{2}}, \sqrt{10}^{h+\frac{1}{2}}]$, where h is a natural number. Hobbs provides data and argues that HOMs seem to be the right level of granularity for representing typical measures.

> ... I have observed that people find it ... easy to come up with half-order-of-magnitude estimates and that these are very often just as informative as they need to be. ... This suggests that there is some cognitive basis for thinking in terms of half orders of mag-

Algorithm $\eta(n \in \mathbb{N})$

1. If $n = 0$ then, return 0.
2. Return $1 + \text{round}(\log_{\sqrt{10}}(n))$.

Figure 7.7: Definition of the function η.

> nitudes. … For scales that are isomorphic to the integers or the reals, precise values are often not available. We need coarser-grained structures on scales (Hobbs, 2000, p. 28).

Since each HOM corresponds to a unique natural number, h, HOMs are linearly-ordered. Naturally, elements of $\Psi(\mathcal{Q})$ should also be linearly-ordered to reflect the linear hierarchy of HOMs. Order over elements of $\mathcal{Q}$ is established by the function $<_{\mathcal{Q}}$ (see Section 7.4.1), and beliefs about the relative orders of various typical amounts is induced by the natural order on whole numbers at the PML. A function, η, maps the natural numbers onto half-orders of magnitude. The algorithm that computes η is shown in Figure 7.7. For some natural number, n, $\eta(n)$ evaluates to a natural number h, such that, if $n = 0$, $h = 0$. Otherwise, n is within the interval $[\sqrt{10}^{(h-1)-\frac{1}{2}}, \sqrt{10}^{(h-1)+\frac{1}{2}}]$. This identifies the HOM to which n belongs.[9] Elements of $\mathcal{Q}$ are associated in $\mathcal{A}$ with whole numbers corresponding to the HOMs they represent. Figure 7.8 shows the algorithm that computes the function ρ, given some some positive real number, n. Step (1) computes the HOM of n. Step (2) simply checks if a $\mathcal{Q}$-term corresponding to the HOM of n has already been introduced into $\Psi(\mathcal{Q})$. If not, step (3) introduces a new term. Steps (5) through (8) make sure that the new term is inserted in the appropriate position within the $<_{\mathcal{Q}}$-chain of elements in $\Psi(\mathcal{Q})$. Figure 7.9 depicts a commuting diagram for time intervals and their durations across the KL-PML interface. Note that, except for η, all the mappings depicted are partial.

The above algorithm guarantees that the image of ρ is linearly-ordered by $<_{\mathcal{Q}}$. That is, symbols denoting durations of intervals for which Cassie has a feel are linearly-ordered. What about intervals for which Cassie has no perceptual experience? For example, suppose that we tell Cassie, by direct assertion, that $\mathsf{Dur}(t,q)$, for some new t. Even though $\mathcal{A}(t)$ is not defined, Cassie may still know the position of q within the $<_{\mathcal{Q}}$-chain. First, q may already be in the chain, which would happen if

[9] The only reason why 1 is added in step 2 of the algorithm is to distinguish the cases where $n = 1$ and $n = 0$. Note that nothing much hangs on this; η merely generates natural numbers that uniquely correspond to distinct HOMs.

Algorithm $\rho(n \in \mathbb{N})$

1. $h \longleftarrow \eta(n)$.
2. If there is $q \in \Psi(\mathcal{Q})$ such that $\mathcal{A}(q) = h$, then return q.
3. Pick some $q \in \mathcal{Q}$, such that $q \notin \Psi(\mathcal{Q})$.
4. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle q, h \rangle\}$
5. $min \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h' < h\}$.
6. $max \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h < h']\}$.
7. If min is not empty, then $\beta \longleftarrow \beta \cup \{q_{gmin} <_{\mathcal{Q}} q\}$, where q_{gmin} is the greatest element of the linearly-ordered poset $\langle min, <_{\mathcal{Q}} \rangle$.
8. If max is not empty, then $\beta \longleftarrow \beta \cup \{q <_{\mathcal{Q}} q_{lmax}\}$, where q_{lmax} is the smallest element of the linearly-ordered poset $\langle max, <_{\mathcal{Q}} \rangle$.
9. return q.

Figure 7.8: The recognition function ρ.

$$
\begin{tikzcd}
\Psi(\mathcal{T}) \arrow[r, "\delta"] \arrow[d, "\mathcal{A}|_{\mathcal{T}}"'] & \Psi(\mathcal{Q}) \arrow[d, "\mathcal{A}|_{\mathcal{Q}}"] \\
\mathbb{N} \arrow[r, "\eta"] \arrow[ur, dashed, "\rho"] & \mathbb{N}
\end{tikzcd}
$$

Figure 7.9: Commuting diagram for time intervals and their durations across the KL-PML interface.

$\mathcal{A}(q)$ is defined. Second, there could be some q' in the $<_Q$-chain such that $\mathsf{Equiv}(q, q')$ is in β, in which case, both q and q' would occupy the same location in the chain. Otherwise, Cassie would not know the exact position of q unless she is explicitly told.

Now, there are two caveats that should be brought to the reader's attention.

1. Ideally, we should be able to make the strong assertion that the image of ρ is exactly the $<_Q$-chain. That is, Cassie may only know the location of an amount in the $<_Q$-chain if she has a perceptual experience of some interval, t, such that $\mathsf{Dur}(t, q)$. Nevertheless, I opt for a weaker assertion: I assume the image of ρ to be a *subset* of the $<_Q$-chain, while maintaining that only amounts for which $\mathcal{A}$ is defined are in the chain. This implies that there may be amounts for which $\mathcal{A}$ is defined and that, nonetheless, are not in the image of ρ (nor δ, for that matter). These are amounts representing the typical durations of states. More specifically, they are terms in the set $\{q|\exists s[\beta \vdash \mathsf{SDur}(s, q)]\}$. Admittedly, there is no deep theoretical motivation for this step. However, it does give us a technical benefit. The solution to the problem of the fleeting now (discussed in the next section) involves Cassie's comparison of the durations of intervals in the chain of NOW-MTFs to typical durations of states. This requires the relative orders of these durations to be known. In an ideal world, Cassie would be able to *learn* the typical durations of states by actually experiencing them. This is, indeed, attainable in our model (steps 3 through 8 in Figure 7.8). For practical purposes, however, we need to be able to *hardwire* associations between terms representing typical durations and HOMs in the alignments set $\mathcal{A}$.

2. For the purpose of this work, reasoning about amounts is confined to reasoning about their order. Complex arithmetic reasoning within the KL (i.e., not subconscious PML computations) may be useful for a planning agent, for example. However, such reasoning is essentially probabilistic (Hobbs, 2000, pp. 29–31) and is beyond the scope of this work.

7.5 Seizing the Fleeting Now

7.5.1 The Lattice of "now"s

Let us now go back to the problem of the fleeting now and try to precisely characterize it within our formal framework. Cassie wonders whether some temporary state, s, holds "now". This wondering may be initiated by various events, including a direct English query by a human operator. "now" (or, in general, the present tense) is interpreted as $*\mathsf{NOW}$. Thus, Cassie's wonder initiates a deductive process for $\mathsf{Holds}(s,*\mathsf{NOW})$. Assuming that NOW is pointing to some term, t_1, then the deduction is actually for $\mathsf{Holds}(s,t_1)$, since $*\mathsf{NOW}$ is but a meta-theoretical shorthand for whichever term NOW happens to be pointing to. In order to figure out whether s holds, Cassie performs some sequence of acts, σ. Since these are acts that Cassie herself performs, she is aware of changes corresponding to their onsets and cessations. Therefore, by the First Principle of Change, NOW moves to some different time, t_2. Cassie determines that s indeed persists. According to the algorithms developed in Section 5.2.6 (particularly, `assert_persist`, `state_persist`, and `setup_new_MTF`), this is recorded by means of two assertions: $\mathsf{MHolds}(s,t)$ and $*\mathsf{NOW} \sqsubset t$, for some newly-introduced t. By **AS2** and **AS3**, this results in Cassie's coming to believe that $\mathsf{Holds}(s,*\mathsf{NOW})$. This belief *looks* exactly like the one for which the deduction was initiated. Nevertheless, since $*\mathsf{NOW}$ is merely a shorthand for t_2, what is deduced is $\mathsf{Holds}(s,t_2)$, whereas the deduction is for $\mathsf{Holds}(s,t_1)$.

Where exactly does the problem lie? I believe that it lies in the interpretation of "now" as $*\mathsf{NOW}$, the most fine-grained representation of the present time.[10] As pointed out in Section 7.1, there is generally a collection of intervals representing the present at different levels of granularity. This collection is, in principle, infinite, and Cassie may have a belief to that effect. Nevertheless, at any point, there is only a finite number of those in $\Psi(\mathcal{T})$, and each is introduced by a specific linguistic or reasoning discourse. One such discourse is the query of whether s holds "now". As pointed out in Section 7.2, "now" should be interpreted as an interval whose length is neither too short, nor too long, for the typical duration of s. This is the basic intuition. How to translate that into a solution to the problem is the subject of the rest of this section. First, a piece of notation.

Definition 7.1 $\lceil . \rceil$ *is a function from the set of* NOW*-intervals to* $\Psi(\mathcal{T})$ *such that, for every* $i \in \mathbb{N}$*,*

[10] Others may believe that the problem lies elsewhere (see Section 7.6 below), but hopefully the proposed solution would satisfy everybody.

$\lceil$*NOW$_i\rceil$ *is the set of all reference intervals in* Φ(*NOW$_i$)*. That is,* $\lceil$*NOW$_i\rceil = \{t|t \in \Phi$(*NOW$_i$) *and* t *is a reference interval*$\}$.

We can immediately make the following observations.

Observation 7.1 *For every* $i \in \mathbb{N}$, $\lceil$*NOW$_i\rceil$ *is a temporal frame.*

Proof. Since, by definition, $\lceil$*NOW$_i\rceil$ is a subset of a temporal frame (namely Φ(*NOW$_i$)), then, by Definition 5.3 (temporal frame), $\lceil$*NOW$_i\rceil$ is a temporal frame. $\square$

Observation 7.2 *For every* $i \in \mathbb{N}$, *NOW$_i \in \lceil$*NOW$_i\rceil$.

Proof. Since *NOW$_i$ is a reference interval (by Axiom 5.3), then, by Definition 7.1, *NOW$_i \in \lceil$*NOW$_i\rceil$. $\square$

Observation 7.3 $\lceil . \rceil$ *is one-to-one.*

Proof. Let $i, j \in \mathbb{N}$ such that $\lceil$*NOW$_i\rceil = \lceil$*NOW$_j\rceil$. Assume that *NOW$_i$ and *NOW$_j$ are distinct NOW-intervals. From Theorems 5.1 and 5.2, Φ(*NOW$_i$) and Φ(*NOW$_j$) are distinct NOW-MTFs. By Theorem 5.3 and Proposition 5.2,

$$\{\text{*NOW}_i, \text{*NOW}_j\} \subseteq \Phi(\text{*NOW}_i) \triangle \Phi(\text{*NOW}_j).$$

But, by Definition 7.1, $\lceil$*NOW$_i\rceil \subseteq \Phi$(*NOW$_i$) and $\lceil$*NOW$_j\rceil \subseteq \Phi$(*NOW$_j$). Therefore, by Observation 7.2,

$$\{\text{*NOW}_i, \text{*NOW}_j\} \subseteq \lceil\text{*NOW}_i\rceil \triangle \lceil\text{*NOW}_j\rceil.$$

This means that $\lceil$*NOW$_i\rceil$ and $\lceil$*NOW$_j\rceil$ are not identical, which leads to a contradiction. Therefore, *NOW$_i$ =*NOW$_j$. Since i and j are arbitrary, then the result applies to all members of $\mathbb{N}$. Therefore, $\lceil . \rceil$ is one-to-one. $\square$

Observation 7.4 *For every* $i \in \mathbb{N}$ $(i > 0)$, $\langle\lceil$*NOW$_i\rceil, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)\rangle$ *is a meet semilattice.*

Proof. Given that $\sqsubset$ is a strict partial order, then, obviously, $\lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)$ is a partial order. Since, for every $t \in \lceil$*NOW$_i\rceil$, $\beta \vdash$*NOW$_i \sqsubset t \vee$*NOW$_i = t$, then every two elements of $\lceil$*NOW$_i\rceil$ have an infimum, namely *NOW$_i$. Therefore, $\langle\lceil$*NOW$_i\rceil, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)\rangle$ is a meet semilattice. $\square$

Algorithm $\texttt{state\_query}(s)$

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{{}^*\mathsf{NOW} \sqsubset t, \mathsf{Dur}(t,q)\}$, where $\beta \vdash \mathsf{SDur}(s,q)$.
3. $\texttt{Backward}(\{\mathsf{Holds}(s,t)\})$.

Figure 7.10: The algorithm `state_query`.

I will, henceforth, refer to $\lceil {}^*\mathsf{NOW}_i \rceil$ as *the i^{th} lattice of now's* (or simply *the lattice of now's* if the context is clear). It may, or may not, be linearly-ordered to form a *stack* of "now"s (Ismail and Shapiro, 2000b). The reason $\lceil {}^*\mathsf{NOW}_i \rceil$ is not necessarily linear is that different extended "present"s may just overlap and are not always nested. A collection of reference intervals are nested only if they are all introduced for the first time in the same $\lceil {}^*\mathsf{NOW}_i \rceil$. That is, if they represent the *same* present time at different levels of granularity.

To solve the problem of the fleeting now, two points need to be revised:

1. How queries about states holding "now" are represented, and

2. How beliefs about states holding "now" are recorded.

The first point has already been discussed in Section 7.2. Figure 7.10 shows the algorithm `state_query` which outlines the steps taken whenever Cassie wonders whether some state, s, holds "now". Basically, a new reference interval is introduced, and its duration is restricted to be within the same HOM as the typical duration of the state s (but see below for a revision of this statement). Note that this new interval is a member of $\lceil {}^*\mathsf{NOW} \rceil$.

According to the algorithms of Section 5.2.6, to record that some state, s, holds, two main assertions are made: $\mathsf{MHolds}(s,t)$, for some new t, and ${}^*\mathsf{NOW} \sqsubset t$. The point behind the second assertion is to make sure that Cassie believes that the state holds in the present. Given that we have different representations of the present, the second assertion should be replaced by a set of assertions for the *appropriate* elements of $\lceil {}^*\mathsf{NOW} \rceil$. In Section 7.2, I hinted that the appropriate elements of $\lceil {}^*\mathsf{NOW} \rceil$ are those whose durations are restricted to be less than, or within, the same HOM as the typical duration of s. Figure 7.11 outlines an algorithm that, given a state, s, and an interval, t, over which it holds, adds a set of assertions to β to make sure that t includes all of the appropriate "now"s. Basically, *t bubbles up* the lattice of "now"s as high as it could, incorporating

Algorithm `state_present`(s,t)

1. For every $t' \in \lceil$*NOW$\rceil$

   2. If $\beta \nvdash \neg\mathsf{Holds}(s,t')$
   and $\delta(t') = q$ or $\beta \vdash \delta(t') <_Q q$, where $\beta \vdash \mathsf{SDur}(s,q)$,
   then $\beta \longleftarrow \beta \cup \{t' \sqsubset t\}$.

Figure 7.11: The algorithm `state_present`.

all those "now"s that fall within its extent. Note that, given the first conjunct of the conditional in step 2, including coarse-grained "now"s within the extent of t is, technically, a default assumption. Also note that, because of the same conjunct, algorithm `state_present` has effect only in the case of determining that a state persists, not that it starts.

How exactly are algorithms `state_query` and `state_present` linked to the rest of the system? I will answer this question in Section 7.5.3. But, first, we need to consider a fine adjustment of the algorithms themselves.

7.5.2 The Backward Projection Factor

Algorithm `state_present` (Figure 7.11) is supposed to capture the intuition that, when one observes a state holding, they assume that it has persisted and will continue to persist for a while. In particular, whenever a state is observed to be holding, not only is it asserted to be holding over *NOW, the most fine-grained representation of the present, but over all representations of the present whose durations fall within the typical duration of the state. By doing so, the state is projected backward in time, which reflects the intuition that the state has not just started. This is very crucial for our proposed solution to the problem of the fleeting now, for it is this backward projection that justifies the agent's belief that the state held at some past query-time.

Now, the basic idea of backward projection is plausible, but the exact details of the process, embodied in algorithm `state_present`, may require some fine-tuning. Consider the example of the agent checking whether John is having lunch (see Section 7.2). The query takes place at 12:15 p.m. and let us say that the typical duration of having lunch is between 15 and 30 minutes. Suppose that, for some reason, the agent could only reach John's office at 12:40 p.m., when he sees John eating. Now, the time between the query and observation events is indeed within the 15-to-30-

Algorithm `state_present`(s,t)

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s,q)$.
2. For every $t' \in \lceil {*}\mathsf{NOW} \rceil$
   3. If $\beta \nvdash \neg\mathsf{Holds}(s,t')$
      and $\delta(t') = q'$ or $\beta \vdash \delta(t') <_Q q'$,
      then $\beta \longleftarrow \beta \cup \{t' \sqsubset t\}$.

Figure 7.12: The revised algorithm `state_present`.

minutes window of having lunch, but is it reasonable for the agent to assume, by perceiving John having lunch at 12:40, that he was having lunch at 12:15? My own intuition is: no. Had the walk to John's office taken 5 or 10 minutes, such an assumption would have been valid. Had it taken 15 minutes, the assumption would have still been possible but not as plausible. But 25 minutes, or even 20, is too long a period to safely make that assumption.

The problem is not whether backward projection per se is valid; I believe it certainly is. The problem is how far in the past one should project a state, how long one should assume that a state, observed to be holding, has been continuously holding. Algorithm `state_present` takes the position that one may project the state backward in time so long as the projection is within the typical duration of the state. Examples like the above, however, renders such a position questionable. Backward projection should be allowed only within some fraction of the typical duration of the state. This is shown in Figure 7.12, a revision of algorithm `state_present`.

The main difference between this version of the algorithm and that in Figure 7.11 is the introduction of a duration q', instead of q, the typical duration of the state s, as an upper limit on the span of backward projection. Step (1) introduces q'—obviously a function of s. The computation of q' works as follows:

1. Compute the geometric center of the HOM corresponding to q (given by $\sqrt{10}^{(\mathcal{A}(q)-1)}$).
2. Raise this value to the value of $bpf(s)$. Note that since HOMs are based on geometric means, this corresponds to "multiplying" the value computed in (1) by the factor $bpf(s)$. Thus, the old version of algorithm `state_present` (Figure 7.11) is a special case where the value of this factor is unity.

Algorithm `state_query`(s)

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s,q)$.
2. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
3. $\beta \longleftarrow \beta \cup \{{}^*\mathsf{NOW} \sqsubset t, \mathsf{Dur}(t,q')\}$.
4. `Backward`$(\{\mathsf{Holds}(s,t)\})$.

Figure 7.13: The algorithm `state_query` for the general case of backward projection factor that is not unity.

3. Return the Q-symbol, q', corresponding to the HOM to which the value computed in 2 belongs. Note that since we use the function ρ for this computation, q' may be a newly-introduced term.

Of course the question now is what the function *bpf* is. The function evaluates to a real number representing the *backward projection factor* of its argument state (hence the name). What do we know about this function? Not much; only that it evaluates to a positive real in the interval $(0,1]$. Actually, I am not even sure if it is necessarily a function of the state, or if it is constant for all states, and in the former case, I do not know exactly how it depends on the state. I believe that the exact definition of the *bpf* function is an empirical question that ought to be posed to psychologists: When observing a state holding, what are the biases of human subjects as to how long the state has been holding, and what are the factors determining those biases? As far as I know, the psychology literature is silent about these issues.

Modifying algorithm `state_present` as indicated above requires a similar modification to algorithm `state_query`. In particular, the reference interval introduced at the query time intuitively represents the period of time during which observing the state holding would be relevant to the query. As such, the duration of that reference interval should be restricted to q' (as computed above) rather than q, the typical duration of the state. This is shown in Figure 7.13. For generality, and since there is no decisive way to determine the innards of the *bpf* function, I will not commit myself to any precise definition of *bpf*. For practical purposes, however, we may take *bpf* to be the constant $\frac{2}{3}$. Note that this is just a working hypothesis that I will not build into the theory.

Algorithm `Backward`($S \subseteq \Psi(\mathcal{S})$)

1. For all $s \in S$
   2. Perform usual backward chaining on s.
   3. If $s \in$ TEMP $\cup$ PERM then `state_query`(s).

Figure 7.14: Modified backward chaining procedure: Version 2

7.5.3 Forward Inference Meets Backward Projection

Let us now return to the question of how algorithms `state_query` and `state_present` fit within the rest of the system. First, consider algorithm `state_query` since there is not much to be said about it. The algorithm is initiated whenever a query is issued about whether some state holds in the present. This can happen in a number of ways. First, it could be a direct query by some other agent (for example, a question posed in English). Second, it could be an *internal* query generated by, for example, the acting system in the process of performing some conditional act.[11] As outlined in Chapter 6, queries about whether a state holds in the present correspond to initiating algorithm `Backward` (see Figure 6.2) with the state among its set-of-states argument. Thus, algorithm `state_query` is simply a step in a modified `Backward`. This new version of `Backward` is shown in Figure 7.14. In what follows, I will indicate the initiation of `state_query` with argument s at $[\![*\mathsf{NOW}_i]\!]$ by saying that *a query is issued for s at* $[\![*\mathsf{NOW}_i]\!]$.

Now, let us turn to algorithm `state_present`. Intuitively, the algorithm should be initiated whenever Cassie determines that a state persists. Thus, one might propose, a call to `state_present` should be added at the appropriate points in algorithms `state_change` and `assert_persist`. Granted, executing these algorithms should also, somehow, result in executing algorithm `state_present`. Nevertheless, note that this only accounts for cases where Cassie determines the persistence of a state through perception, proprioception, or direct assertion; it does not cover the case of inference. To appreciate this point, consider the following argument dismissing it as a problem.

How may Cassie infer that a state, s_1, holds in the present? Typically, this would involve Cassie's

[11] By "conditional act", I mean an act that is performed only if some state holds; for example, crossing the street only if the walk-light is on (see Chapter 6).

having a belief along the following schema, where s_2 is some state different form s_1 and p is a proposition describing background conditions that need to be true for s_2 to entail s_1.

$$\forall Tv[[\mathsf{Holds}(s_2, Tv) \wedge p] \Rightarrow \mathsf{Holds}(s_1, Tv)]$$

The argument now goes like this. Determining that s_1 holds through inference, necessarily involves determining that s_2 holds. Ultimately, this has to be based on determining that some s_n holds through some way other than inference, i.e., perception, proprioception, or direct assertion. Since algorithms `state_change` and `assert_persist` initiate forward inference, then following the above schema, it would be inferred that s_1 holds over all the coarse-grained "now"s over which s_2 holds. Therefore, backward projection is applied to s_1, albeit not directly through algorithm `state_present`.

There are at least two flaws in the above argument.

1. The reason the argument works is that, by performing simple forward inference, s_1 is backward-projected by virtue of s_2's projection. The problem, however, is that, inspecting the above schema, it is very possible that the typical duration of s_1 is longer than that of s_2. Thus, s_1 would not be projected into the past as far as *its* typical duration allows, only as far as s_2's typical duration does.[12]

2. Even if we choose to ignore the issue raised in 1, there is a more basic problem. Given the above schema, Cassie may infer that s_1 holds, not only by determining that s_2 holds, but also by coming to believe the proposition p (or an instance of the schema itself, for that matter). Even more dramatic, although determining that s_1 holds through inference typically involves a belief along the above schema, this is only typical, not necessary. What seems necessary is that the variable Tv be mentioned in the antecedent. For example, a possible $\mathcal{FOCS}$-translation of (7.1) below appears in (7.2), where Tc_1 and Tc_2 denote "3" and "5", respectively.

   (7.1) Stu is home from 3 to 5.

   (7.2) $\forall Tv[[Tc_1 \prec Tv \wedge Tv \prec Tc_2] \Rightarrow \mathsf{Holds}(\mathsf{At}(\mathrm{STU}, \mathrm{HOME}), Tv)]$.

[12] This is similar to the problem of the persistence of derived information raised by (Myers and Smith, 1988) (also known as the problem of "dependent fluents" (Giunchiglia and Lifschitz, 1995)).

By telling Cassie (7.2), she may determine that Stu is now home if she already believes that "now" is between 3 and 5. Note that this does not involve Cassie's determining that any other state holds.[13] Thus, should algorithm `state_present` be initiated only within `state_change` and `assert_persist`, Cassie may determine that the state $\mathrm{At}(\mathrm{STU}, \mathrm{HOME})$ persists without the application of backward projection.

Therefore, it seems that backward projection should be built into the very process of inference. In particular, whenever Cassie infers that a state holds "now", backward projection should be applied. This may be achieved by introducing a time-sensitive version of the forward inference algorithm `Forward`. This is shown in Figure 7.15. The algorithm takes a set of propositions, P, as an argument. It does two main things: (i) it applies forward inference on members of P as per Chapter 6 (this is achieved by algorithm `Forward_old` in step 1 which is the `Forward` algorithm of Chapter 6) and (ii) it applies backward projection whenever appropriate. Step 1 initiates traditional forward inference on P. The set P_{inf} is the set of those propositions inferred in the process. The rest of the steps are responsible for backward projection. The conditional in step 3 filters in those members of $P \cup P_{\mathrm{inf}}$ that are about states holding "now". Basically, it then initiates algorithm `state_present`. The only catch is that a state interval is introduced (step 5a) if one is not associated with the state asserted to be holding "now". Note that, this way, adherence to Axiom 5.2 is built into the algorithms, and the theory builder need not worry about it. Also note that the filtering process in step 3 considers elements of both P_{inf} and P. Since new information is always introduced with forward inference (see Section 1.5.4), then backward projection gets applied to all states—those perceived, proprioceived, directly asserted, or inferred.

7.5.4 Pulling the Rope

By introducing coarse-grained reference intervals in both the querying and the assertion processes, the problem of the fleeting now may be readily solved. For whatever reference interval the querying process introduces would be available at the assertion time if its duration is long enough. Of course, what now remains is an account of how to carry this reference interval from one MTF to the next until it either expires, or the state is observed (*pulling the rope* à la Section 7.2). The first step

[13] It might be possible to rephrase (7.2) so that the antecedent involves Holds, but as (7.2) attests, it is also possible not to.

Algorithm `Forward`$(P \subseteq \Psi(\mathcal{P}))$

1. $P_{\text{inf}} \longleftarrow$ `Forward_old`(P).
2. For every $p \in P \cup P_{\text{inf}}$
   3. If $p = \mathsf{Holds}(s, {*}\mathsf{NOW})$, for some $s \in \Psi(\text{TEMP})$, then
      4. If there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s,t)$, then `state_present`$(s,t)$.
      5. Else
         5a. Pick some $t \in \mathcal{T}$ such that $t \notin \Psi(\mathcal{T})$.
         5b. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t), {*}\mathsf{NOW} \sqsubset t\}$.
         5c. `state_present`(s,t).

Figure 7.15: Algorithm `Forward`. Building backward projection into forward inference.

is to modify the `setup_new_MTF` algorithm (Figure 5.9) so that it takes elements of $\lceil {*}\mathsf{NOW} \rceil$ into account. Figure 7.16 outlines the modified algorithm. The algorithm assumes that the new MTF is the i^{th} NOW-MTF. Steps 7 through 9 incorporate members of $\lceil {*}\mathsf{NOW}_{i-1} \rceil$ into $\lceil {*}\mathsf{NOW}_i \rceil$ just in case the amount of time elapsed since they were introduced (represented by $\pi(t)$ which is computed according to Equation 7.1) is less than or within the same HOM as their projected durations (otherwise they move into the past as per step 9).

Given what has been presented so far, we can now formally prove the effectiveness of our solution to the problem of the fleeting now. To accomplish this, we need to introduce some notation that should help us proceed through the proofs more conveniently.

Definition 7.2 *For every $i, j \in \mathbb{N}$, the* **temporal distance** *between $\Phi({*}\mathsf{NOW}_i)$ and $\Phi({*}\mathsf{NOW}_j)$, denoted $d_t(\Phi({*}\mathsf{NOW}_i), \Phi({*}\mathsf{NOW}_j))$, is the amount of time that, Cassie feels, separates the start of $[\![{*}\mathsf{NOW}_i]\!]$ and the start of $[\![{*}\mathsf{NOW}_j]\!]$. More precisely,*

$$d_t(\Phi({*}\mathsf{NOW}_i), \Phi({*}\mathsf{NOW}_j)) = \sum_{k=\min(i,j)}^{\max(i,j)-1} \pi({*}\mathsf{NOW}_k),$$

where $\min(i,j)$ *and* $\max(i,j)$ *are the smaller and larger of i and j, respectively.*

It should be clear that d_t is a metric over the space of NOW-MTFs. I will not attempt to prove this, however, since the proof is obvious and not instructive in any relevant sense. What should be noted though is that the temporal distance is only defined for NOW-MTFs since the definition primarily

Algorithm `setup_new_MTF`

1. `move_NOW`
2. For all $\mu \in \mathcal{M}_{\text{prop}}$
   3. If there are s and t such that $*\mu = \mathsf{MHolds}(s,t)$
      then $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
4. For all $\mu \in \mathcal{M}_{\text{per}}$
   5. For all s and t such that $\mathsf{MHolds}(s,t) \in *\mu$
      6. $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
7. For all $t \in \lceil \text{*}\mathsf{NOW}_{i-1} \rceil \setminus \{\text{*}\mathsf{NOW}_{i-1}\}$
   8. If $\eta(\pi(t)) \leq \mathcal{A}(\delta(t))$, then $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
   9. Else $\beta \longleftarrow \beta \cup \{t \prec \text{*}\mathsf{NOW}\}$.

Figure 7.16: Algorithm `setup_new_MTF`: Version 2

depends on their linear order (cf. Theorem 5.4) and the fact that π is defined for NOW intervals but not necessarily for other atomic intervals.

To prove the main result, we need to first prove two lemmas. The first asserts that reference intervals are extended into the appropriate NOW-MTFs.

Lemma 7.1 *For every $t \in \Psi(\mathcal{T})$, $q \in \Psi(\mathcal{Q})$, and $i,n \in \mathbb{N}$ ($i > 0$), if*

1. *t is a reference interval,*
2. *$\beta \vdash \mathsf{Dur}(t,q)$,*
3. *$\mathcal{A}(q)$ is defined,*
4. *for every $\Phi \in \text{Span}(t)$, Φ is a NOW-MTF or a transition MTF,*
5. *$\Phi(\text{*}\mathsf{NOW}_i)$ is the smallest element of the poset $\langle \text{Span}(t), \text{precedes} \rangle$, and*
6. *$\eta(d_t(\Phi(\text{*}\mathsf{NOW}_i), \Phi(\text{*}\mathsf{NOW}_{i+n}))) \leq \mathcal{A}(q)$, then*
7. *$\Phi(\text{*}\mathsf{NOW}_{i+n}) \in \text{Span}(t)$.*

Proof. I use induction on n to prove the lemma.

Basis. Let $n = 0$. Since $\Phi(\text{*NOW}_i)$ is the smallest element of $\langle \text{Span}(t), \text{precedes} \rangle$, then, trivially, $\Phi(\text{*NOW}_i) \in \text{Span}(t)$.

Induction Hypothesis. Assume that, for every $t \in \Psi(\mathcal{T})$, $q \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ ($i > 0$), the conjunction of statements 1 through 6 implies statement 7, for some $n \in \mathbb{N}$.

Induction Step. We need to show that, for $n+1$, the conjunction of statements 1 through 6 implies statement 7. By statement 6,

$$\eta(d_t(\Phi(\text{*NOW}_i), \Phi(\text{*NOW}_{i+n+1}))) \leq \mathcal{A}(q).$$

By Definition 7.2, and since $i < i+n+1$,

$$\eta(\sum_{k=i}^{i+n} \pi(\text{*NOW}_k)) \leq \mathcal{A}(q).$$

Therefore,

$$\eta(\sum_{k=i}^{i+n-1} \pi(\text{*NOW}_k) + \pi(\text{*NOW}_{i+n})) \leq \mathcal{A}(q).$$

By Definition 7.2, and since $i \leq i+n$,

$$\eta(d_t(\Phi(\text{*NOW}_i), \Phi(\text{*NOW}_{i+n})) + \pi(\text{*NOW}_n)) \leq \mathcal{A}(q).$$

Since $\pi(\text{*NOW}_n)$ is a positive quantity, and since η is monotonic,[14] then

$$\eta(d_t(\Phi(\text{*NOW}_i), \Phi(\text{*NOW}_{i+n}))) \leq \mathcal{A}(q).$$

Therefore, using the induction hypothesis, $\Phi(\text{*NOW}_{i+n}) \in \text{Span}(t)$. By Definitions 5.8 (Span) and 7.1 ($\lceil . \rceil$), and statement 1, $t \in \lceil \text{*NOW}_{i+n} \rceil$. As NOW moves from *NOW_{i+n} to *NOW_{i+n+1}, algorithm `setup_new_MTF` gets executed and the conditional in step 8 is applied to t. Since, by statement 5, t is not atomic, then, by step 6 of algorithm π,

$$\pi(t) = \sum_{\Phi(t_i) \in \text{ Span}(t)} \pi(t_i)$$

Now, note that at the time of evaluation of the conditional, the greatest element of $\langle \text{Span}(t), \text{precedes} \rangle$ is $\Phi(\text{*NOW}_{i+n})$. Given that $\pi(t_{\text{tr}}) = 0$, for any transition interval t_{tr} (step 4 of algorithm π), then, by statement 4, only NOW-MTFs contribute to the value of $\pi(t)$. Since $\Phi(\text{*NOW}_i)$ is the smallest element thereof (statement 5) then

[14] Note that this needs to proved; the proof is obvious though.

$$\pi(t) = \sum_{\Phi(t_i) \in \text{ Span}(t)} \pi(t_i) \quad \leq \quad \sum_{k=i}^{i+n} \pi(*\mathsf{NOW}_k).^{15}$$

But since

$$\eta(\sum_{k=i}^{i+n} \pi(*\mathsf{NOW}_k)) = \eta(d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_{i+n+1}))) \leq \mathcal{A}(q),$$

then

$$\pi(t) \leq \mathcal{A}(q) = \mathcal{A}(\delta(t)).$$

Therefore, by step 8 of algorithm `setup_new_MTF`, $*\mathsf{NOW}_{i+n+1} \sqsubset t \in \beta$. Thus, by Definition 5.8 (Span), $\Phi(*\mathsf{NOW}_{i+n+1}) \in \text{Span}(t)$. $\square$

Inspecting the proof of Lemma 7.1, we can draw the following monotonicity result. I will not show the proof since it could be easily reconstructed from the proof the lemma.

Corollary 7.1 *For all $i, j, k \in \mathbb{N}$, if $i \leq j \leq k$, then* $d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_j)) \leq d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_k))$.

We now prove the second lemma needed to present the main result. The lemma asserts that, whenever Cassie determines that a state persists, she believes that it holds over all of the appropriate members of the lattice of "now"s.

Lemma 7.2 *For every $s \in \Psi(\text{TEMP})$, $t \in \Psi(\mathcal{T})$, $q_1, q_2 \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ ($i > 1$), if*

1. $q_1 = \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, *where* $\beta \vdash \mathsf{SDur}(s, q)$,
2. $\beta \vdash \mathsf{Dur}(t, q_2)$,
3. $q_2 = q_1$ *or* $\beta \vdash q_2 <_{\mathcal{Q}} q_1$,
4. $t \in \lceil *\mathsf{NOW}_i \rceil$,
5. $\beta \nvdash \neg\mathsf{Holds}(s, t)$, *and*
6. *Cassie determines that s persists at* $*\mathsf{NOW}_i$, *then*
7. $\beta \vdash \mathsf{Holds}(s, t)$.

[15] This is, in fact, an equality given the convexity of t (**AT9**).

Proof. By statement 6 and Definition 5.15, at $[\![^*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s, {}^*\mathsf{NOW}_{i-1})$ and $\beta \nvdash \neg\mathsf{Holds}(s, {}^*\mathsf{NOW}_{i-1})$ and, at $[\![^*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s, {}^*\mathsf{NOW}_i)$. For this to be the case, there must be some $P \subset \mathcal{P}$, such that, at $[\![^*\mathsf{NOW}_{i-1}]\!]$, for every $p \in P$, $\beta \nvdash p$, and, at $[\![^*\mathsf{NOW}_i]\!]$, $P \subseteq \beta$. The existence of such P is necessary since, at $[\![^*\mathsf{NOW}_{i-1}]\!]$, ${}^*\mathsf{NOW}_i$ was yet to be introduced. Since adding new information to β always initiates forward inference (see Section 1.5.4), then algorithm `Forward` of figure 7.15, gets executed, at $[\![^*\mathsf{NOW}_i]\!]$, with P as an argument. Since, at $[\![^*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s, {}^*\mathsf{NOW}_i)$, then either $\mathsf{Holds}(s, {}^*\mathsf{NOW}_i) \in P$ or $\mathsf{Holds}(s, {}^*\mathsf{NOW}_i) \in P_{\text{inf}}$ as computed by step 1 of `Forward`. Therefore, $\mathsf{Holds}(s, {}^*\mathsf{NOW}_i) \in P \cup P_{\text{inf}}$. Since, at this point, ${}^*\mathsf{NOW}_i = {}^*\mathsf{NOW}$, then algorithm `state_present` gets executed (by steps 4 or 5c) with s and t' as arguments, where t' is the state interval associated with s at $[\![^*\mathsf{NOW}_i]\!]$. By statements 1 through 5, step 5 of `state_present` adds the proposition $t \sqsubset t'$ to β. By **AS2** and **AS3**, $\beta \vdash \mathsf{Holds}(s, t)$. $\square$

Given the above results, we can now prove the following theorem which establishes the effectiveness of our solution to the problem of the fleeting now.

Theorem 7.1 *For every $s \in \Psi(\text{TEMP})$, $q \in \Psi(Q)$, and $i, n \in \mathbb{N}$ ($i > 0$), if*

1. $\beta \vdash \mathsf{SDur}(s, q)$,
2. $\mathcal{A}(q)$ *is defined,*
3. *at* $[\![^*\mathsf{NOW}_i]\!]$, *a query is issued for s,*
4. *n is the smallest integer such that Cassie determines that s persists at* ${}^*\mathsf{NOW}_{i+n}$,
5. $d_t(\Phi({}^*\mathsf{NOW}_i), \Phi({}^*\mathsf{NOW}_{i+n})) \leq \sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)}$, *and*
6. *for all m,* $0 \leq m \leq n$, $\beta \nvdash \neg\mathsf{Holds}(s, {}^*\mathsf{NOW}_{i+m})$, *then*
7. *for all m,* $0 \leq m \leq n$, $\beta \vdash \mathsf{Holds}(s, {}^*\mathsf{NOW}_{i+m})$.

Before stating the proof of the above theorem, I need to point out an issue that would render the proof not as formal as I would like it to be. It should be clear how the proof would proceed. Basically, we shall show that if the reference interval introduced at the time of the query (by algorithm `state_query`) survives until the time of determining the persistence of s, then, by Lemma 7.1, all the intervening "now"s are subintervals thereof. Using Lemma 7.2, we can then show that

s holds over all of those "now"s. The problem, however, lies in statements 4 of Lemma 7.1 and 5 of Lemma 7.2 which should hold for the proof to follow. Statement 4 of Lemma 7.1 requires *all* MTFs in the span of the reference interval to be NOW-MTFs or transition-MTFs. Unfortunately, the theory presented here does not explicitly, and formally, specify that this is the case. For all we know, a user of the system (maybe a human operator interacting with Cassie) may make an assertion involving the reference interval, thereby adding some other kind of MTF to its span. Nevertheless, it should be noted that a valid assumption is that this would not happen (unless somebody is going out of their way to mess with Cassie's mind). Why? Because reference intervals introduced as extended representations of the present (in particular, those introduced by algorithm `state_query`) are, in a sense, *private* symbols of Cassie's mental language. More specifically, users of the system do not know anything about these internally-generated terms. More importantly, I envision assertions to be made to the system in natural language, and there would be no way to thus refer to these reference intervals. Therefore, in the following proof, I will make the tacit assumption that reference intervals introduced as extended representations of the present contain only NOW-MTFs or transition-intervals in their spans. In fact, I will make a stronger (yet, still reasonable) assumption: the only assertions involving reference intervals introduced by algorithm `state_query` are those made by algorithms `state_query`, `state_present`, and `setup_new_MTF`. In particular, note in assertions of the form $t' \sqsubset t$, where t is a reference interval introduced by algorithm `state_query`, t' is always a NOW-interval. In addition, successive assertions of that form, involve successive values of NOW. Thus, the smallest element of $\langle \text{Span}(t), \text{precedes} \rangle$ is $\Phi(*\mathsf{NOW}_i)$, where $*\mathsf{NOW}_i$ is the value of NOW at the time algorithm `state_query` gets executed.

There is a catch though. Given the convexity of intervals (**AT10**), it may actually be inferred that for some non-NOW interval, t', $t' \sqsubset t$. However, this may only happen if t' falls between two NOW intervals. Following the same argument above (also see Section 5.1.4), assertions about NOW intervals cannot be directly made unless the interval is *NOW. The only situation where an interval is inserted between two NOW intervals is when it is a transition interval (see Algorithms `state_change`, `assert_start`, and `assert_cease`), which is consistent with the assumption that $\text{Span}(t)$ includes only NOW and transition MTFs.

Given this assumption, statement 5 of Lemma 7.2 follows from statement 6 of the theorem. Statement 5 of Lemma 7.2 requires that Cassie does not (implicitly or explicitly) believe that s

does not hold over the reference interval, t, introduced by `state_query`. The problem is that, without the above assumption, anybody can assert anything about reference intervals, and it would be impossible to prove statement 5. How does the above assumption save the situation? To answer this question, consider how Cassie may come to believe that s does not hold over t. There are two possibilities. First, someone directly asserts $\neg\mathsf{Holds}(s,t)$. This, however, is dismissed by our assumption. The only assertions involving reference intervals are those made by algorithms `state_query`, `state_present`, and `setup_new_MTF`. Evidently, these assertions are about temporal parthood; they do not mention any states. Second, Cassie believes that s does not hold over some sub-interval of t. However, given our assumption that the only sub-intervals of t that matter are NOW-intervals this possibility is also dismissed by statement 6 of the theorem.

Proof. By statement 3, algorithm `state_query` gets executed at $[\![{*}\mathsf{NOW}_i]\!]$ (see the discussion in Section 7.5.3). Step 2 introduces a new reference interval, t', into $\Psi(\mathcal{T})$. Step 3 makes the two assertions: ${*}\mathsf{NOW} \sqsubset t'$ and $\mathsf{Dur}(t', q')$, where $q' = \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$. Since, by statement 5,

$$d_t(\Phi({*}\mathsf{NOW}_i), \Phi({*}\mathsf{NOW}_{i+n})) \leq \sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)},$$

then by the monotonicity of η and the definition of ρ,

$$\eta(d_t(\Phi({*}\mathsf{NOW}_i), \Phi({*}\mathsf{NOW}_{i+n}))) \leq \mathcal{A}(q').$$

In fact, by Corollary 7.1 and the monotonicity of η, for every m, $0 \leq m \leq n$

$$\eta(d_t(\Phi({*}\mathsf{NOW}_i), \Phi({*}\mathsf{NOW}_{i+m}))) \leq \mathcal{A}(q').$$

Therefore, the following statements are true.

1. t' is a reference interval.

2. $\beta \vdash \mathsf{Dur}(t', q')$.

3. $\mathcal{A}(q') = \sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)}$ is defined.

4. For every $\Phi \in \text{Span}(t')$, Φ is a NOW-MTF or a transition-MTF (which follows from the assumption discussed above).

5. $\Phi({*}\mathsf{NOW}_i)$ is the smallest element of the poset $\langle \text{Span}(t'), \text{precedes} \rangle$ (which follows from the same assumption).

6. $d_t(\Phi(\text{*NOW}_i), \Phi(\text{*NOW}_{i+m})) \leq \mathcal{A}(q')$, for every m, $0 \leq m \leq n$.

Therefore, by Lemma 7.1, , for every m, $0 \leq m \leq n$, $\Phi(\text{*NOW}_{i+m}) \in \text{Span}(t')$. By Definition 5.8 (Span) and 7.1 ($\lceil . \rceil$), it follows that, for every m, $0 \leq m \leq n$, $t' \in \lceil \text{*NOW}_{i+m} \rceil$ and $\text{*NOW}_{i+m} \sqsubset t'$.

Now, by statement 6 of the theorem and the assumption discussed above, $\beta \nvdash \neg\mathsf{Holds}(s, t')$. Therefore, the following statements are true.

1. $q' = \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s, q)$.
2. $\beta \vdash \mathsf{Dur}(t', q')$.
3. $t' \in \lceil \text{*NOW}_{i+n} \rceil$.
4. $\beta \nvdash \neg\mathsf{Holds}(s, t')$.
5. Cassie determines that s persists at *NOW_{i+n} (by statement 4 of the theorem).

Therefore, by Lemma 7.2, $\beta \vdash \mathsf{Holds}(s, t')$. By the divisitivity of states (**AS2**), it follows that, for every m, $0 \leq m \leq n$, $\beta \vdash \mathsf{Holds}(s, \text{*NOW}_{i+m})$. □

7.6 Where Exactly Does the Problem Lie?

All of the above being said, we should now consider possible objections and suspicions that some might raise against the proposed solution—indeed, the proposed interpretation—of the problem of the fleeting now. As pointed out in Section 7.1, the problem, I believe, lies in an intrinsic vagueness of the concept of "now": "now" refers to an interval that does not have any well-defined, context-independent boundaries. The NOW-intervals of the theory represent the concept of "now" at the finest level of granularity and non-NOW reference intervals stand for coarser representations thereof. The solution to the problem of the fleeting now is based on interpreting the "now" of the query, not as *NOW, but as a coarser reference interval—a member of $\lceil \text{*NOW} \rceil$.

The above notwithstanding, it should be noted that, *technically*, nothing much hangs on this assumption. In particular, some may argue that such a move is not motivated, and that "now" is not vague, but always refers to the sharp instant of experience.[16] Even under that assumption,

[16]For example, Antony Galton, in personal communication.

the proposed solution would still work. This would only involve replacing step 3 in algorithm `state_query` (Figure 7.13) with "Initiate deduction for $\mathsf{Holds}(s,$*NOW$)$". More precisely, the solution would work as follows.

1. Introduce a new reference interval, t, and restrict its length to a factor of the typical duration of s as determined by the *bpf* function.

2. Assert that t is a super-interval of t_1, the current value of *NOW.

3. Initiate deduction for $\mathsf{Holds}(s,$*NOW$)$.

4. Suppose that s is determined to persist at t_2.

5. If t survives until t_2, then algorithm `state_present` guarantees that we get $\mathsf{Holds}(s,t)$.

6. Now, by the divisitivity of states, we also get $\mathsf{Holds}(s,t_1)$, and the query is answered positively.

Thus, the only difference between the two scenarios (other than the conceptual one) is the use of the divisitivity of states to *link* determining the persistence of s, at t_2, to the query at t_1. In this case, the only crucial notion is that of backward projection.

But, now, someone may argue that, if backward projection is the key, then the issue is more general. For example, if, at 2 p.m., someone says that John was having lunch at 12:15 p.m. Should we assume that John was having lunch sharply at 12:15? Or should we employ backward projection and project the state of having lunch several minutes prior to 12:15? In more abstract terms, suppose that, at $[\![*\mathsf{NOW}_i]\!]$, Cassie is told that a state, s, held at $*\mathsf{NOW}_{i-n}$. Shouldn't the backward projection mechanism be applied in this case too, asserting s to hold over all members of $\lceil *\mathsf{NOW}_{i-n} \rceil$ whose durations fall within bpf(s)? Intuitively, this seems reasonable. Therefore, backward projection is not peculiar to the present, it may also be applied to assertions about the past, and, thus, the problem of the fleeting now has nothing to do with "now" per se!

Although, technically, the above argument may have some merit, there are reasons why I believe that it does not provide an adequate explanation of the problem of the fleeting now. The main point is that, though it might appear that backward projection is all that is needed to account for the problem, I believe that this is only the case at the technical level, not at a more fundamental, conceptual level. This belief stems from my concern, not only with explaining the reasoning aspects of the problem of the fleeting now, but with what it reveals about our very conceptualization of the present, in particular, as revealed by our use of language.

Consider, in more detail, the example of John's lunch. Suppose that Stu asks Cassie:

(7.3) Is John having lunch (now)?

Cassie walks to John's office and finds him eating his lunch. How should she reply to Stu? Intuitively, she should say:

(7.4) Yes, he is.

What is interesting here is that *both* the question and the reply are in the present tense. How can that be, given that (7.3) and (7.4) are uttered, strictly speaking, at different times? The only way that this may be possible is if there is an interpretation of the present that encompasses the times of both utterances. This is exactly the gist of the proposed solution to the problem of the fleeting now: interpreting the "now" of the question as a coarse reference interval. If this interval persists until the time of the reply (where the persistence is determined by contextual factors), then both the question and the reply fall within the same "now" and may, thus, be both expressed in the present tense.

Someone might claim that the fact that both (7.3) and (7.4) are expressed in the present tense is a mere peculiarity of language. In particular, since (7.4) is a reply to (7.3), it uses the same tense. However, this is obviously not true. Consider the situation where Cassie walks to John's office to find that he has just finished his lunch. In this case, (7.4) would not be an appropriate answer to (7.3). Rather, (7.5) (or a variant thereof) seems to be the only reasonable thing to say. The important thing to note here is that an affirmative answer to (7.3) would have to be expressed in the past tense. (Note that simply "He *was*" with the appropriate intonation is equally plausible.)

(7.5) He was when you asked, but not any more.

Another objection may be that (7.4) is expressed in the present tense simply because John is indeed having lunch at the time of its utterance. But consider the situation where Cassie walks to John's office to find him just about to start eating his lunch. Possible reasonable replies may be:

(7.6) He has just started.

(7.7) He *is* now, but not when you asked.

Or even:

(7.8) He is *now* having lunch.

Although all of these replies are in the present tense, they all (implicitly or explicitly) make it clear that, at the time of the question, John was not having lunch. Our ability to make these distinctions and yet find (7.4) to be a plausible reply in the original situation (where Cassie finds John in the midst of having lunch) can only be explained by the vagueness of the present.

7.7 Conclusions

The problem of the fleeting now is a problem that faces agents interleaving reasoning and acting while maintaining a sense of the present time. Basically, the agent is interested in whether some state holds "now". However, since reasoning and sensory acts take time, whatever conclusion it makes will be, strictly speaking, about a different "now". The solution proposed in this chapter is based on the simple intuition that the concept of "now" is vague as to the size of the interval it represents. The agent wonders whether the state holds, not at the sharp moment of experience, but over a broader "now", an interval whose duration is comparable to the typical duration of the state. Such an interval may still be "now", relative to some coarse level of granularity, at the time of the conclusion. Whether this is the case depends on the amount of time it takes the agent to reach a conclusion. To formalize these intuitions, we developed a theory of temporal perception. The theory, motivated by psychological research, not only provides a cognitive agent with a sense of the passage of time, but also gives it a feel of how much time has passed and intuitions on the typical durations of various activities and states.

The agent's sense of how much time has passed is grounded in a subconscious internal clock whose ticks are aligned with terms in the (conscious-level) object language. The use of the internal clock follows one of two schools of thought in the psychology of time. It should be noted, however, that this choice is merely a tactic, not a strategic commitment; the formal theory can accommodate either of the two dominant views in the psychology of time perception. Another tactical decision is the choice of Hobbs's half orders of magnitude (Hobbs, 2000) to represent amounts of time. In this respect, the theory presented here raises a number of questions to psychologists of time:

1. What kinds of mental representations and processes are involved in reasoning about typical durations of states?

2. As regards the backward projection factor, what are the biases of human subjects as to how long a perceived state has been holding?

3. What are the factors determining those biases?

Through the use of reference intervals, knowledge of their durations, beliefs about the typical durations of states, and backward projection; we are able to provide an intuitive solution to the problem of the fleeting now.

Chapter 8

Persistence

8.1 Another Note on the Frame Problem

In Section 5.2.2, I proposed a mechanism for temporally-projecting perceived and bodily states. This proposal provides an elegant account of the *current* persistence of such states as time goes by—the variant of the frame problem that emerges in our theory. Nevertheless, the model of time perception developed in Section 7.4 has some further implications for the frame problem. Research on solving the frame problem is dominated by the important insight that, except for very few exceptions, the world is generally stable. Indeed, it must be, or otherwise we will not be able to function appropriately. (Shanahan, 1997) summarizes this view in what he calls "the common sense [sic] law of inertia": "Inertia is normal. Change is exceptional." (Shanahan, 1997, p. 18). It is this basic idea that underlies most of the solutions proposed to the frame problem. Now, as put by (Shanahan, 1997), the commonsense law of inertia indeed makes sense. However, what often lies behind most of the proposed solutions to the problem is a stronger version of the law of inertia. Typically, the assumption is that, unless it is known that something has happened which would cause some state to cease to hold, then the state still holds. Again, such an assumption might be reasonable as long as one is willing to be open-minded in interpreting "something has happened". In particular, there is something that is always happening, and yet has only been considered by a handful of authors (McDermott, 1982; Dean and Kanazawa, 1988; Lifschtiz and Rabinov, 1989; Kanazawa, 1992, for example) in the literature on the frame problem, namely *the passage of time*. Time is always moving and, as it does, it is reasonable to assume that some states have ceased to hold, based on intuitions about their typical durations. For example, if one sees a cat sitting on a mat, a green traffic-light, or a smiling face; it is not reasonable to assume that, an hour later, the cat is *still* on the mat, the light is *still* green, and the person is *still* smiling. In fact, (Stein, 1990, p. 372) advises that states "that are 'known' to change without warning" should be excluded from the scope of the commonsense law of inertia. What Stein fails to tell us is how such knowledge may come about.

Similar to our notion of the typical duration of a state, (McDermott, 1982) introduces the notion of a *life-time* associated with a *persistence*. Roughly, a persistence is the uninterrupted holding of a state over a period of time, which weakly corresponds to our state intervals (in McDermott's terminology, our "states" are "facts"). The life-time associated with a persistence is the duration of the period of time throughout which a state may be assumed to persist unless otherwise known (note the non-monotonicity involved).

(Dean and Kanazawa, 1988) (followed by (Kanazawa, 1992)) present similar notions within a probabilistic framework. Using a discrete time line, Dean and Kanazawa present a system to compute the probability of a given fluent holding at a given time (their system does much more than that, but this is the aspect that concerns us here). Given a time, t, and a fluent, P, if no knowledge is available about any events causing P to start or cease between t and $t+\Delta$, then the probability of P holding over $t+\Delta$ is given by the following formula ($\langle P, t\rangle$ is their way of saying that P holds at t).

$$p(\langle P, t+\Delta\rangle) = p(\langle P, t+\Delta\rangle | \langle P, t\rangle)\ p(\langle P, t\rangle).$$

In the particular examples discussed in (Dean and Kanazawa, 1988), $p(\langle P, t+\Delta\rangle | \langle P, t\rangle) = e^{-\lambda\Delta}$. Thus, unless otherwise known, fluents (or states, in our theory) *decay* as time passes by—the same point made by (McDermott, 1982). Corresponding to our typical durations, the exponent λ is determined by the rate of decay (and, hence, the life-time) of the given fluent/state.

Although not introduced to present inherent temporal constraints on the persistence of states, Shanahan's notion of *trajectories* (Shanahan, 1997, ch. 13) may be used in that venue. In Shanahan's system, a formula "Trajectory(f1,t1,f2,d)" roughly means that if fluent f1 is initiated at time t1, then fluent f2 holds at time t1+d.[1] If the logic allows fluents of the form "¬f1" to represent the fluent

[1]Trajectories are used in (Shanahan, 1997) in order to represent continuous change. Shanahan's semantics involve a qualification of f1 and f2 that does not concern us here: f1 is a "discrete" fluent and f2 is a "continuous" fluent. Basically, a discrete fluent holds over intervals of non-zero durations, while a continuous fluent may hold instantaneously. The reader may notice that this is the same distinction that (Galton, 1990) makes between "states of motion" and "states of position" (see Section 2.1.2). However, unlike Galton's notions, a fluent's being discrete or continuous is not an intrinsic

that holds whenever f1 does not (i.e., the *complement* of the fluent à la Chapter 3), then the formula "Trajectory(f1,t1,¬f1,d)" may be used to state that d is the life-time/typical duration of the fluent f1.

In all of the above systems, knowledge of how much time has passed is crucial. Where would this knowledge come from? It should be obvious that our model readily provides the answer: the pacemaker. Using Cassie's perception of time as provided by the pacemaker, she may defeasibly reason about the persistence of states. In particular, just as with reference intervals, a state interval associated with a state for which there is a known typical duration may be excluded from a new MTF if the amount of time that has elapsed since the state was determined to be holding is longer than its typical duration. Alternately, a state, that is not known to have ceased, may be assumed to extend into a new MTF so long as it has not exceeded its typical duration. For an agent acting on-line, I believe that, in most cases, this is how it determines the persistence of states, since it seldom comes to know of events causing them to cease (unless these are actions of the agent itself).

Consider the following example. Suppose Cassie is in a room with two sources of light, red and green. Cassie looks at the red light, notices that it is on, and then turns toward the green light. Is the red light still on? The answer would reasonably be "yes" if the amount of time elapsed is within the same HOM as the typical duration of the red light. Otherwise, a suitable answer is "no". But now consider the following situation. Cassie turns back to the red light and finds that it is turned off. Now there are two possibilities. If the amount of time elapsed since she had last observed the light on is longer than its typical duration, then no problem; Cassie would have already assumed that the light is off. On the other hand, if not that much time has passed, then what is observed (the red light is off) contradicts what would be otherwise assumed (the red light is still on). Needless to say, knowledge induced through perception, proprioception, or direct assertion (our model for communication) is more credible than assumptions based on intuitions about typical durations of states. Thus, making use of typical durations to extend a state interval into the new MTF should be a final resort, should only be turned to in case Cassie cannot infer that the state interval has already moved into the past.

More formally, this can be achieved as per the third (and final) revision of algorithm `setup_new_MTF` in Figure 8.1. This version of the algorithm contains many features not in that

property thereof, but an extrinsic one. To take an example from (Shanahan, 1997, p. 260), while a ball is falling, its height is a continuous fluent, but, once it lands, the height becomes a discrete fluent.

Algorithm `setup_new_MTF`$(S^{\downarrow} \subseteq \Psi(\text{TEMP}), t_{tr})$

1. `move_NOW`
2. For all $\mu \in \mathcal{M}_{prop}$
    3. If there are s and t such that $*\mu = \mathsf{MHolds}(s,t)$, then $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t\}$.
4. For all $\mu \in \mathcal{M}_{per}$
    5. For all s and t such that $\mathsf{MHolds}(s,t) \in *\mu$
        6. $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t\}$.
7. For all $t \in \lceil \mathsf{*NOW}_{i-1} \rceil \setminus \{\mathsf{*NOW}_{i-1}\}$
    8. If $\eta(\pi(t)) \leq \mathcal{A}(\delta(t))$, then $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t\}$.
    9. Else $\beta \longleftarrow \beta \cup \{t \prec \mathsf{*NOW}\}$.
10. $\beta \longleftarrow \beta \cup \{t_{tr} \prec \mathsf{*NOW}\}$.
11. For all $s \in S^{\downarrow}$ $\beta \longleftarrow \beta \cup \{\neg\mathsf{Holds}(s, \mathsf{*NOW})\}$.
12. For every state interval $t \in \Phi(\mathsf{*NOW}_{i-1})$,
    13. If $\beta \vdash \mathsf{*NOW} \sqsubset t$, then $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t\}$.
    14. Else, if $\beta \vdash \mathsf{*NOW} \not\sqsubset t$, then `event_cease`$(s,t,t_{tr})$, where s is the state with which t is associated.
    15. Else, if $\eta(\pi(t)) \leq \mathcal{A}(q)$, where $\mathsf{SDur}(s,q)$ and $\mathsf{MHolds}(s,t)$, then

        $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t\}$.
    16. Else `event_cease`(s,t,t_{tr}),
        where s is the state with which t is associated.

Figure 8.1: Algorithm `setup_new_MTF`: Version 3.

of Figure 7.16 (in particular, steps 10 through 16). First, however, note that this version of the algorithm takes two arguments. The second argument, t_{tr}, is the transition interval introduced by whichever algorithm initiates `setup_new_MTF` to represent the time of the transition inducing the movement of NOW. Step 10 adds the assertion that t_{tr} precedes the newly-generated *NOW. The primary use of t_{tr} is in step 14 discussed below. The first argument, $S^{\downarrow}$, represents the set of states that have just been directly asserted to cease. That is, $S^{\downarrow}$ is itself the argument S of algorithm `assert_cease` of Figure 5.22. Step 11 simply makes sure that members of $S^{\downarrow}$ are excluded from the new MTF. I will explain below why $S^{\downarrow}$ is included as an argument to `setup_new_MTF`, but let us, first, inspect the rest of the algorithm.

First, consider the loop starting in step 12. Step 12 restricts the application of the following steps to the set of state intervals in the previous MTF. Step 13 includes into the new MTF all those states that hold in the previous MTF and that may be inferred to hold in the new one. For example, these include states that Cassie has just ceased to perceive but that she, nonetheless, has reason to believe that they still hold. Step 14 takes a complementary action; it introduces a cessation event for all those state intervals that may be inferred to *not* be members of the new MTF. These may include bodily states that have just ceased to hold, perceivable states that were just perceived to cease, and members of $S^{\downarrow}$ (which are guaranteed to satisfy the conditional in step 14 given the assertions made in step 11). If a state cannot be inferred to hold, or to not hold, in the new MTF, then it is included, or excluded, by steps 15 and 16 depending on whether it has exceeded its typical duration. Note that this is exactly the defeasible assumption made by (McDermott, 1982).

The reason we need to include $S^{\downarrow}$ and step 11 in algorithm `setup_new_MTF` is purely technical. Inspecting Figure 5.22, algorithm `assert_cease` is responsible for two main tasks: (i) moving NOW (with the construction of the new MTF) and (ii) updating β so that Cassie believes that states in the argument S no longer hold. The problem, however, is that updating β (as per Figure 5.22) takes place after executing algorithm `setup_new_MTF`, and, indeed, it must, since the new value of NOW needed for the propositions updating β is only introduced then. Why is this a problem? It is a problem, given the current version of `setup_new_MTF`, since step 15 may incorporate some of the states in S into the new MTF. We might, thus, be introducing inconsistencies into the system when this is, clearly, unwarranted. The only way that updating β be performed *after* introducing the new value of NOW but *before* executing step 15 of `setup_new_MTF` is if updating β is itself part

PRE-CONDITIONS:

1. *For every* $s \in S$, $\beta \vdash \mathsf{Holds}(s, {}^*\mathsf{NOW})$.

Algorithm `assert_cease`($S \subseteq$ TEMP)

1. Pick some $t_{\mathrm{tr}} \in \mathcal{T}$, such that $t_{\mathrm{tr}} \notin \Psi(\mathcal{T})$.
2. `setup_new_MTF`(S, t_{tr}).
3. $P_{\mathrm{new}} \longleftarrow \{\}$.
4. For all $s_i \in S$, $P_{\mathrm{new}} \longleftarrow P_{\mathrm{new}} \cup \{\neg\mathsf{Holds}(s_i, {}^*\mathsf{NOW})\}$.
5. `Forward`(P_{new}).

Figure 8.2: Revised version of algorithm `assert_cease`.

of algorithm `setup_new_MTF`. This is exactly what step 11 is responsible for. Note that, this way, algorithm `assert_cease` becomes much simpler than before (see Figure 8.2).[2]

For ease of reference, Appendix D contains a compilation of the final versions of all the temporal progression algorithms presented.

8.2 On Soundness and Completeness

At this point, something should be said about the correctness of the system presented here. As the title of this section shows, the plan is to discuss issues of soundness and completeness. A word of caution though. I am *not* going to provide proofs of soundness and completeness that usually accompany the presentation of a logical system (hence the "on" in the title). Rather, I shall present criteria for soundness and completeness that are slightly different from the traditional ones, and more appropriate for our purposes, and outline, in terms as precise as possible, proofs that our system observes these criteria.

First of all, recall that, as far as our logic is concerned, we do not need to show any proofs of (traditional) soundness and completeness; our logic is a standard first-order logic: no modal operators, no default rules, nothing exotic. What is the problem then? The problem is that traditional soundness and completeness worry about rules of inference, whether one can infer all and only

[2]It should be noted that, despite this revision of `assert_cease`, Theorem 5.10 still holds.

those things that are true in every interpretation in which an initial set of premises is true. Where the premises are coming from is irrelevant to traditional soundness and completeness, and rightfully so. But in our system we need to worry about where some of the premises are coming from. In particular, according to Sections 5.2.2 and 8.1, the temporal projection of states is sometimes not the result of logical inference, but of subconscious temporal progression processes. Beliefs generated thus are, technically, premises since their presence in β is not the result of inference. On the other hand, they are not exactly normal premises pushed into β by fiat; they are the result of the interaction of the algorithms, the meta-theoretical axioms, the contents of β, and the rules of inference. Since these constitute our system, Cassie's temporal machinery, we need to show that it does the right job, that Cassie believes (or can infer) all and only those propositions that she is justified to believe based on what she knows, and more importantly, what she feels (which makes our task different from that of autoepistemic logicians (Moore, 1984; Moore, 1985)).

In particular, given the focus of our investigation, we are interested in two types of beliefs:

1. A state s's holding in the present.

2. The event of some state s's holding being in the past.

In the terminology developed above, propositions of type 1 are of the form $*\mathsf{NOW} \sqsubset t$, where t is a state interval associated with s. Similarly, type 2 propositions are of the form $t \prec *\mathsf{NOW}$. We need to show that our system results in Cassie's believing propositions of type 1 or 2 if and only if it is reasonable for her to do so. The question now is what "reasonable" means. This depends primarily on the type of state in question and how Cassie comes to believe propositions about it.

8.2.1 Feeling is Believing: A "Completeness" Result for Bodily States

For a bodily state, s, characterizing when it is reasonable for Cassie to believe that s holds in the present is clear-cut—when, and only when, the state occupies some proprioceptual modality. Similarly, Cassie should believe that an event of s holding has moved into the past when, and only when, the state stops occupying proprioceptual modalities. Note that nowhere here have I required s to be actually holding, in the first case, or have actually ceased to hold in the second. All that concerns us is whether Cassie *feels* that it does. For example, Cassie might be holding an object but feel that she is empty-handed due to some glitch at the SAL. Yet, believing that she is empty-handed is justified by what Cassie feels.[3] Thus, I do not care whether what Cassie believes is in fact knowledge—justified true belief; I only care about its being justified. That is, I only require Cassie to be *rational*.

The main result of this section is Theorem 8.1 below. The theorem simply states that, by the system presented above (that is, the logic, the axioms, and the PML algorithms), whenever a bodily state occupies some proprioceptual modality, Cassie has a belief that that state holds. Note what this means. It means that bodily sensations, or feelings, purely PML phenomena, are aligned with conscious beliefs at the KL. Thus, in a sense, the theorem states a completeness result. In order to prove the theorem, we need to prove a lemma first. Informally, the lemma states that the values of modality variables are faithful to what Cassie feels.

Lemma 8.1 *For every* $i \in \mathbb{N}$ *(*$i > 0$*),* $s \in$ TEMP *such that* $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$*, and* $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$*, at* $[\![*\mathsf{NOW}_i]\!]$*,* $[\![s]\!]$ *occupies the modality corresponding to* μ *if and only if there is some* $t \in \Psi(\mathcal{T})$ *such that* $*\mu = \mathsf{MHolds}(s,t)$ *and* $\beta \vdash \mathsf{MHolds}(s,t)$*.*

Proof. Let s be a bodily state and let μ be in $\mathrm{Mod}_{\mathrm{prop}}(s)$. I use induction on i to prove the lemma.

Basis. Let $i = 1$. $\Phi(*\mathsf{NOW}_1)$ is established by initiating algorithm `initialize`. Suppose that $[\![s]\!]$ occupies the modality corresponding to μ. Therefore, s is a member of S, the argument of `initialize`. By executing algorithm `state_start` in step 6, the proposition $\mathsf{MHolds}(s,t)$ gets added to β (step 1) and algorithm `state_ceive` gets executed with arguments s and t (step 3), where t is the interval introduced by step 4 of `initialize`. Since $[\![s]\!]$ occupies the modality corresponding to μ, then, by initiating algorithm `state_proprioceive` in step 1 of `state_ceive`, $*\mu = \mathsf{MHolds}(s,t)$. In addition, trivially, $\beta \vdash \mathsf{MHolds}(s,t)$ since $\mathsf{MHolds}(s,t) \in \beta$ by step 1 of `state_start`. Conversely, suppose that $[\![s]\!]$ does not occupy the modality corresponding to μ. By Axiom 5.6, some other state, $[\![s']\!]$ occupies that modality. Following the same argument outlined above, algorithm `initialize` results in assigning μ a proposition $\mathsf{MHolds}(s',t')$, for some $t' \in \Psi(\mathcal{T})$. Therefore, $*\mu \neq \mathsf{MHolds}(s,t)$ for any $t \in \mathcal{T}$.

Induction Hypothesis. Assume that, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ if and only if there is some $t \in \Psi(\mathcal{T})$ such that $*\mu = \mathsf{MHolds}(s,t)$ and $\beta \vdash \mathsf{MHolds}(s,t)$.

Induction Step. Consider the situation at $[\![*\mathsf{NOW}_{i+1}]\!]$. I will break the proof of the induction

[3]Unless Cassie believes that there is some SAL problem.

step into four sub-proofs: sub-proofs (a) and (b) cover two complementary cases of the *if* part of the theorem, sub-proofs (c) and (d) cover two similar cases of the *only-if* part.

(a). Suppose that $[\![s]\!]$ occupies the modality corresponding to μ at both $[\![{}^*\mathsf{NOW}_{i+1}]\!]$ and $[\![{}^*\mathsf{NOW}_{i}]\!]$. By the induction hypothesis, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, μ is set to $\mathsf{MHolds}(s,t)$, for some $t \in \Psi(\mathcal{T})$. In addition, $\beta \vdash \mathsf{MHolds}(s,t)$. Thus, we only need to show that μ retains its value over $[\![{}^*\mathsf{NOW}_{i+1}]\!]$. The value of NOW can change from ${}^*\mathsf{NOW}_i$ to ${}^*\mathsf{NOW}_{i+1}$ only if one of algorithms `state_change`, `assert_start`, or `assert_cease` is executed. Since `assert_start` and `assert_cease` do not involve any steps that set proprioception modality variables, then if NOW moves due to the execution of either algorithm, μ would retain its value. Now, suppose that NOW moves due to the execution of algorithm `state_change`. The only places in `state_change` where proprioception modality variables are set are steps 5 and 6b (through the execution of algorithm `state_proprioceive`). But since $[\![s]\!]$ occupies the modality corresponding to μ at both $[\![{}^*\mathsf{NOW}_{i+1}]\!]$ and $[\![{}^*\mathsf{NOW}_{i}]\!]$, then by Axiom 5.6 and pre-condition 4, there is no $s' \in S^{\uparrow}$ such that $[\![s']\!]$ occupies the modality corresponding to μ. Thus, μ never gets changed by algorithm `state_proprioceive`. Therefore, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, following the execution of `state_change`, $*\mu = \mathsf{MHolds}(s,t)$.

(b). Suppose that, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, but not at $[\![{}^*\mathsf{NOW}_{i}]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ. By Axiom 5.7, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, there is no $\mu' \in \mathcal{M}_{\text{prop}}$ such that $[\![s]\!]$ occupies the modality corresponding to μ'. Further, by Axiom 5.10, there is no $\mu' \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ'. By the induction hypothesis, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, there is no $t \in \mathcal{T}$ and $\mu' \in \mathcal{M}$ such that $*\mu' = \mathsf{MHolds}(s,t)$. Therefore, as $[\![s]\!]$ undergoes the transition from not occupying the modality corresponding to μ to occupying it, s satisfies pre-conditions 3 and 4 of algorithm `state_change`, and the algorithm gets initiated with $s \in S^{\uparrow}$. By the proof of Theorem 5.5, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, for some $t \in \mathcal{T}$, $*\mu = \mathsf{MHolds}(s,t)$ and $\beta \vdash \mathsf{MHolds}(s,t)$.

(c). Suppose that, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, but not at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ. By Axiom 5.6, there is some s' ($s' \neq s$) such that, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, but not at $[\![{}^*\mathsf{NOW}_{i}]\!]$, $[\![s']\!]$ occupies the modality corresponding to μ. Thus, following the proof of part (b) above (switching s' and s), at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, $*\mu = \mathsf{MHolds}(s',t')$, for some $t' \in \mathcal{T}$. Therefore, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, $*\mu \neq \mathsf{MHolds}(s,t)$, for any $t \in \mathcal{T}$.

(d). Suppose that, at neither $[\![{}^*\mathsf{NOW}_{i}]\!]$ nor $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, does $[\![s]\!]$ occupy the modality corresponding to μ. By Axiom 5.6, there are s' and s'', different from s, such that, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s']\!]$ occupies the modality corresponding to μ and, at $[\![*\mathsf{NOW}_{i+1}]\!]$, $[\![s'']\!]$ occupies the modality corresponding to μ. If $s'' = s'$, then, following the proof of part (a) above, at $[\![*\mathsf{NOW}_{i+1}]\!]$, $*\mu = \mathsf{MHolds}(s'',t'')$, for some $t'' \in \mathcal{T}$. If, on the other hand, $s'' \neq s'$, then following the proof of part (b) above, at $[\![*\mathsf{NOW}_{i+1}]\!]$, $*\mu = \mathsf{MHolds}(s'',t'')$, for some $t'' \in \mathcal{T}$. Since $s \neq s''$, then in either case, at $[\![*\mathsf{NOW}_{i+1}]\!]$, $*\mu \neq \mathsf{MHolds}(s,t)$, for any $t \in \mathcal{T}$.

From (a), (b), (c), and (d) the induction step follows. Since s and μ are arbitrary, the lemma follows. $\square$

Given the above result, we are now ready to prove the main theorem for this section.

Theorem 8.1 *For every $i \in \mathbb{N}$ ($i > 0$), $s \in$ TEMP such that* $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$*, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$, at $[\![*\mathsf{NOW}_i]\!]$, if $[\![s]\!]$ occupies the modality corresponding to μ, then* $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$.

Proof. Pick some $i \in \mathbb{N}$ ($i > 0$), $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. Suppose that, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ. By Lemma 8.1, at $[\![*\mathsf{NOW}_i]\!]$, $*\mu = \mathsf{MHolds}(s,t)$ and $\beta \vdash \mathsf{MHolds}(s,t)$ for some $t \in \Psi(\mathcal{T})$. Step 3 of algorithm `setup_new_MTF` results in adding the proposition $p = *\mathsf{NOW}_i \sqsubset t$ to β. Given p, **AS2**, and **AS3**, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$. Since i, s, and μ are arbitrary, then the result applies to all $i \in \mathbb{N}$, $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. $\square$

8.2.2 Is Believing Feeling?

The above theorem is essentially a completeness-like result. Similarly, we may prove the following soundness-like Theorem.

Theorem 8.2 *For every $i \in \mathbb{N}$ ($i > 0$), $s \in$ TEMP such that* $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$*, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$, at $[\![*\mathsf{NOW}_i]\!]$, if $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$, then $[\![s]\!]$ occupies the modality corresponding to μ.*

Proof. We prove the theorem using contraposition. Pick some $i \in \mathbb{N}$ ($i > 0$), $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. Suppose that, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ does not occupy the modality corresponding to μ. By Axioms 5.6 (p. 152) and 5.7 (p. 153), there is some $s' \in$ TEMP such that $s' \neq s$, $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s')$, and, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s']\!]$ occupies the modality corresponding to μ. By Theorem 8.1, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s',*\mathsf{NOW}_i)$. Given Axiom 5.8 and **TS1**, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_i)$. Since i, s, and μ are arbitrary, then the result applies to all $i \in \mathbb{N}$, $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. $\square$

Having proved the above theorem, I feel obliged to raise a number of issues that may render the reader a little skeptical about Theorem 8.2 (and Theorem 8.1, for that matter). This is due to a number of reasons that have to do with gullibility, astuteness, and the need for future research.

First of all, recall what I mean by "soundness" here. I do not merely mean the property of a *logical system* that allows only valid *inferences*, I mean that feature of a *cognitive agent* that allows it to have only *justified beliefs*. As per Theorem 8.1, bodily sensations justify beliefs about bodily states, but there could be other justifications for such beliefs too. For example, a human (or otherwise) agent may *tell* Cassie about states of her own body, for example that she is holding a fruit (recall our discussion in Section 5.2.5 in relation to the sixth principle of change). If Cassie feels that she is holding a fruit, then no problem. But what if she does not? Should she *believe* that she is holding a fruit? It depends. If Cassie is a highly-gullible agent (which most AI systems are), then she should believe whatever she is told, and that would be a counter example for the *believing-is-feeling* hypothesis for bodily states. But this is only considering one extreme of the scale of agent-gullibility. On the other extreme, Cassie always ignores, and never believes, anything she is *told* about her bodily states; if she feels them, she will believe in them, according to Theorem 8.1. But this too is an extreme position. In fact, in order for Cassie to initially learn what different bodily sensations mean, some outside help is inevitable. A moderate position is, therefore, required. A possibility would be for Cassie to believe those assertions about states of her body that do not contradict beliefs invoked by what she feels. For example, feeling and believing that she is holding an apple, she should reject an outside statement indicating that she is empty-handed. However, given the way the system has been developed (in particular, Axioms 5.6 and 5.8), unless Cassie is told exactly what she feels, the asserted information would have to be contradictory to Cassie's beliefs. Thus, in this case too, Cassie would reject anything that she is told about bodily states she does not feel.

The most lenient position, the one I tacitly adopt, is to allow Cassie to hold beliefs contradictory to what she feels. There are justifications for this. For suppose that the holding-an-apple bodily sensations are mere hallucinations due to some SAL problem. In that case, Cassie may learn about such a problem from the contradictory external input as long as it is coming from some credible source (for example, a human supervisor, not a fellow robot). What this boils down to, then, is the issue of belief revision in a multi-source environment. Progress on this front is currently underway but still not complete enough to be integrated into our theory (see (Johnson and Shapiro, 2000a) for a preliminary report). Were Cassie to choose to believe that she is holding an apple (because the source of the "empty-handed"-assertion is less credible than proprioception) then no problem; Theorems 8.1 and 8.2 are upheld. On the other hand, if Cassie were to believe that she is empty-handed, then both theorems would be violated. Thus, Theorems 8.1 and 8.2 are valid as long as Cassie does not choose to believe, as mandated by belief revision principles, that a bodily state holds when (Cassie feels) that it does not occupy any modalities.

8.2.3 Seeing is Believing: A "Completeness" Result for Perceived States

Similar to what we did in Section 8.2.1, I shall show that the system guarantees that, whenever Cassie is perceiving a state holding, then she may syntactically infer that it does. We first prove a lemma similar to Lemma 8.1

Lemma 8.2 *For every $i \in \mathbb{N}$ ($i > 0$) and $s \in$ TEMP such that $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, at $[\![{}^*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ if and only if there is some $t \in \Psi(\mathcal{T})$ such that $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s,t)$.*

Proof. Pick some $s \in$ TEMP such that $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined. I use induction on i to prove the lemma.

Basis. Let $i = 1$. $\Phi({}^*\mathsf{NOW}_1)$ is established by initiating algorithm `initialize`. Suppose that $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. Therefore, s is a member of S, the argument of `initialize`. By executing algorithm `state_start` in step 6, the proposition $\mathsf{MHolds}(s,t)$ gets added to β (step 1) and algorithm `state_ceive` gets executed with arguments s and t (step 3), where t is the interval introduced by step 4 of `initialize`. Since $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$, then, by algorithm `state_perceive` (step 2 of `state_ceive`), $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$. Conversely, suppose that $[\![s]\!]$ is not perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By Axiom 5.9 and pre-condition 1, $s \notin S$. Thus, the execution of algorithm `state_perceive` would not result in adding any propositions concerning s to $*\mathrm{Mod}_{\mathrm{per}}(s)$. Since, for $[\![{}^*\mathsf{NOW}_1]\!]$, algorithm `initialize` is the only place where modality variables are set, then, at $[\![{}^*\mathsf{NOW}_1]\!]$, $*\mathrm{Mod}_{\mathrm{per}}(s) \neq \mathsf{MHolds}(s,t)$ for any $t \in \mathcal{T}$.

Induction Hypothesis. Assume that, at $[\![{}^*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ if and only if there is some $t \in \Psi(\mathcal{T})$ such that $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and

$\beta \vdash \mathsf{MHolds}(s,t)$.

Induction Step. Consider the situation at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$. Similar to the proof of Lemma 8.1, I break down the proof into four parts.

(a). Suppose that $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ at both $[\![{}^*\mathsf{NOW}_{i+1}]\!]$ and $[\![{}^*\mathsf{NOW}_{i}]\!]$. By the induction hypothesis, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s,t)$, for some $t \in \Psi(\mathcal{T})$. Thus, we only need to show that the proposition $\mathsf{MHolds}(s,t)$ remains a member of $*\mathrm{Mod}_{\mathrm{per}}(s)$ over $[\![{}^*\mathsf{NOW}_{i+1}]\!]$. The value of NOW can change from ${}^*\mathsf{NOW}_i$ to ${}^*\mathsf{NOW}_{i+1}$ only if one of algorithms `state_change`, `assert_start`, or `assert_cease` is executed. Since `assert_start` and `assert_cease` do not involve any steps that set perception modality variables, then if NOW moves due to the execution of either algorithm, $\mathrm{Mod}_{\mathrm{per}}(s)$ would not be changed. Now, suppose that NOW moves due to the execution of algorithm `state_change`. The only place in `state_change` where propositions get removed from perception modality variables is step 5, through the execution of algorithm `cease_perceive`. But since $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ at both $[\![{}^*\mathsf{NOW}_{i+1}]\!]$ and $[\![{}^*\mathsf{NOW}_{i}]\!]$, then, by pre-condition 2 of `state_change`, $s \notin S^{\downarrow}$. Thus, algorithm `cease_perceive` does not get applied to s, and the proposition $\mathsf{MHolds}(s,t)$ never gets removed from $\mathrm{Mod}_{\mathrm{per}}(s)$ by algorithm `state_change`. Therefore, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, following the execution of `state_change`, $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s,t)$.

(b). Suppose that, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, but not at $[\![{}^*\mathsf{NOW}_{i}]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By the induction hypothesis, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, there is no $t \in \mathcal{T}$ such that $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$. Thus, by Axiom 5.9, s satisfies pre-conditions 3 and 4 of algorithm `state_change`, and the algorithm gets initiated with $s \in S^{\uparrow}$. By the proof of Theorem 5.5, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$, for some $t \in \Psi(\mathcal{T})$.

(c). Suppose that, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, but not at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By Axiom 5.9, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, there is no $\mu \in \mathcal{M}_{\mathrm{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ, and, by Axiom 5.10, $[\![s]\!]$ does not occupy any proprioception modality. Given the induction hypothesis, at $[\![{}^*\mathsf{NOW}_{i}]\!]$, $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$, for some $t \in \Psi(\mathcal{T})$. Therefore, s satisfies pre-conditions 1 and 2 of algorithm `state_change`, and the algorithm gets initiated with $s \in S^{\downarrow}$. By Theorem 5.6, at $[\![{}^*\mathsf{NOW}_{i+1}]\!]$, after executing the algorithm, $\mathsf{MHolds}(s,t) \notin *\mathrm{Mod}_{\mathrm{per}}(s)$, for any $t \in \mathcal{T}$.

(d). Suppose that, at neither $[\![\text{*NOW}_i]\!]$ nor $[\![\text{*NOW}_{i+1}]\!]$, is $[\![s]\!]$ perceived via the modality corresponding to $\text{Mod}_{\text{per}}(s)$. By the induction hypothesis, at $[\![\text{*NOW}_i]\!]$, $\mathsf{MHolds}(s,t) \notin *\text{Mod}_{\text{per}}(s)$, for any $t \in \mathcal{T}$. Therefore, we only need to show that $*\text{Mod}_{\text{per}}(s)$ continues to not include any propositions concerning s over $[\![\text{*NOW}_{i+1}]\!]$. The value of NOW can change from *NOW_i to *NOW_{i+1} only if one of algorithms `state_change`, `assert_start`, or `assert_cease` is executed. Since `assert_start` and `assert_cease` do not involve any steps that set perception modality variables, then if NOW moves due to the execution of either algorithm, $\text{Mod}_{\text{per}}(s)$ would not be changed. Now, suppose that NOW moves due to the execution of algorithm `state_change`. Propositions may only be added to perception modality variables through the execution of algorithm `state_perceive`. In algorithm `state_change`, this may happen in steps 5, 6b, and 6c. However, since s violates pre-condition 4 of the algorithm, $s \notin S^{\uparrow}$, and `state_perceive` never gets applied to it. Therefore, at $[\![\text{*NOW}_{i+1}]\!]$, following the execution of `state_change`, $\mathsf{MHolds}(s,t) \notin *\text{Mod}_{\text{per}}(s)$, for any $t \in \mathcal{T}$.

From (a), (b), (c), and (d) the induction step follows. Since s is arbitrary, the lemma follows. □

Given the above result, we can now prove a theorem corresponding to Theorem 8.1.

Theorem 8.3 *For every $i \in \mathbb{N}$ ($i > 0$) and $s \in$ TEMP such that* $\text{Mod}_{\text{per}}(s)$ *is defined, if, at* $[\![\text{*NOW}_i]\!]$, $[\![s]\!]$ *is perceived via the modality corresponding to* $\text{Mod}_{\text{per}}(s)$*, then* $\beta \vdash \mathsf{Holds}(s,\text{*NOW}_i)$.

Proof. Pick some $i \in \mathbb{N}$ ($i > 0$) and $s \in$TEMP such that $\text{Mod}_{\text{per}}(s)$ is defined. Suppose that, at $[\![\text{*NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\text{Mod}_{\text{per}}(s)$. By Lemma 8.2, at $[\![\text{*NOW}_i]\!]$, $\mathsf{MHolds}(s,t) \in *\text{Mod}_{\text{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s,t)$, for some $t \in \Psi(\mathcal{T})$. Steps 4 through 6 of algorithm `setup_new_MTF` result in adding the proposition $p = \text{*NOW}_i \sqsubset t$ to β. Given p, **AS2**, and **AS3**, $\beta \vdash \mathsf{Holds}(s,\text{*NOW}_i)$. Since i and s are arbitrary, then the result applies to all $i \in \mathbb{N}$ ($i > 0$) and $s \in$TEMP such that $\text{Mod}_{\text{per}}(s)$ is defined. □

Similar to Theorem 8.1, the above theorem is a completeness-like results for perceivable states. However, unlike with Theorem 8.2, we cannot prove a corresponding soundness-like result. Note that a soundness result would state that, whenever Cassie believes that a perceivable state holds, then she perceives it. Evidently, such a result cannot be justified on any empirical grounds; Cassie may be told that a perceivable state holds even though she does not perceive it. Unlike the case with bodily states, one cannot argue that Cassie should not believe any such assertions if she does not actually perceive the state. For example, there is nothing wrong in Cassie's believing Stu's assertion that the walk-light is on when she is not looking towards the walk-light.

8.2.4 Persistence through Time-Perception

Theorems 8.1 and 8.3 have illustrated that Cassie may syntactically infer that a state holds based on the non-syntactic, PML phenomena of perception and proprioception. As pointed out in Section 8.1, Cassie may also believe that a state continues to hold based on her knowledge of the typical duration of the state and her sense of how much time has passed. The latter is another PML phenomenon that falls outside the bounds of the logical theory. In this section, we prove that, in the appropriate circumstances, Cassie's feel for how much time has passed results in the appropriate beliefs about the persistence of states.

Theorem 8.4 *For every $t \in \Psi(\mathcal{T})$, $s \in \Psi(\text{TEMP})$, $q \in \Psi(Q)$, and $i, n \in \mathbb{N}$ ($i > 0$), if*

1. $\beta \vdash \mathsf{MHolds}(s,t)$,
2. $\beta \vdash \mathsf{SDur}(s,q)$,
3. $\mathcal{A}(q)$ *is defined,*
4. *for every* $\Phi \in \text{Span}(t)$, Φ *is a* NOW*-MTF or a transition-MTF,*
5. $\Phi(*\mathsf{NOW}_i)$ *is the smallest element of the poset* $\langle \text{Span}(t), \text{precedes} \rangle$,
6. $\eta(d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_{i+n}))) \leq \mathcal{A}(q)$, *and*
7. *for every* $m \in \mathbb{N}$, $0 < m \leq n$, *Cassie does not determine that s ceases to hold at* $*\mathsf{NOW}_{i+m}$, *then*
8. $\Phi(*\mathsf{NOW}_{i+n}) \in \text{Span}(t)$.

The reader should note that the statement of the theorem is similar to the statement of Lemma 7.1; the proof closely follows that of the lemma.

Proof. I use induction on n to prove the lemma.

Basis. Let $n = 0$. Given statement 5, trivially, $\Phi(*\mathsf{NOW}_i) \in \text{Span}(t)$.

Induction Hypothesis. Assume that, for every $t \in \Psi(\mathcal{T})$, $s \in \Psi(\textsc{Temp})$, $q \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ ($i > 0$), the conjunction of statements 1 through 7 implies statement 8, for some $n \in \mathbb{N}$.

Induction Step. We need to show, for $n+1$, that the conjunction of statements 1 through 7 implies statement 8. Given statement 6,

$$\eta(d_t(\Phi(^*\mathsf{NOW}_i), \Phi(^*\mathsf{NOW}_{i+n+1}))) \leq \mathcal{A}(q).$$

By Corollary 7.1,

$$\eta(d_t(\Phi(^*\mathsf{NOW}_i), \Phi(^*\mathsf{NOW}_{i+n}))) \leq \mathcal{A}(q),$$

which, by the induction hypothesis, implies that $\Phi(^*\mathsf{NOW}_{i+n}) \in \text{Span}(t)$. By Definition 5.8 (Span), $t \in \Phi(^*\mathsf{NOW}_{i+n})$. As NOW moves from $^*\mathsf{NOW}_{i+n}$ to $^*\mathsf{NOW}_{i+n+1}$, algorithm `setup_new_MTF` gets executed. Since t is a state interval in $\Phi(^*\mathsf{NOW}_{i+n})$, then we are only concerned with steps 12 through 16 of the algorithm.

If $\beta \vdash {}^*\mathsf{NOW}_{i+n+1} \sqsubset t$, then the conditional of step 13 is satisfied, and the proposition $^*\mathsf{NOW}_{i+n+1} \sqsubset t$ is added to β. It follows by Definitions 5.4 (MTFs) and 5.8 that $\Phi(^*\mathsf{NOW}_{i+n+1}) \in \text{Span}(t)$.

On the other hand, if $\beta \nvdash {}^*\mathsf{NOW}_{i+n+1} \sqsubset t$, then control flows to step 14 of `setup_new_MTF`. Since $t \in \Phi(^*\mathsf{NOW}_{i+n})$, then $\beta \vdash {}^*\mathsf{NOW}_{i+n} \sqsubset t$. Thus, by **AS2** and **AS3**, $\beta \vdash \mathsf{Holds}(s, {}^*\mathsf{NOW}_{i+n})$. Given statement 7, Cassie does not determine that s ceases to hold at $^*\mathsf{NOW}_{i+n+1}$. Therefore, by Definition 5.17, $\beta \nvdash {}^*\mathsf{NOW}_{i+n+1} \not\sqsubseteq t$. Thus, the conditional in step 14 of algorithm `setup_new_MTF` fails and control flows to step 15. Following the same reasoning in the proof of Lemma 7.1, the conditional in 14 is true with respect to t, and the proposition $^*\mathsf{NOW}_{i+n+1} \sqsubset t$ gets added to β. By Definitions 5.4 and 5.8, $\Phi(^*\mathsf{NOW}_{i+n+1}) \in \text{Span}(t)$. $\square$

8.2.5 Past States

We have shown that, at any time, if it is reasonable for Cassie to believe that some state holds (based on PML phenomena), then she indeed believes (at least implicitly) that it holds. In this section, we show that, at any time, Cassie's beliefs also reasonably reflect that the event of a state holding has moved into the past.

First, we prove the following post-condition of algorithm `setup_new_MTF`.

Lemma 8.3 *For every $t \in \Psi(\mathcal{T})$, $s \in \Psi(\textsc{Temp})$, $q \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ ($i > 1$), if*

1. $\beta \vdash \mathsf{MHolds}(s,t)$,

2. $^*\mathsf{NOW}_{i-1} \sqsubset t \in \beta$,

3. $\beta \vdash \mathsf{SDur}(s,q)$,

4. $\mathcal{A}(q)$ *is defined, and*

5. *for every* $\Phi \in \mathrm{Span}(t)$, Φ *is a* NOW*-MTF or a transition-MTF, then*

6. *following the* $i-1^{st}$ *execution of algorithm* `setup_new_MTF`, $\beta \vdash {}^*\mathsf{NOW}_i \sqsubset t$ *or* $\beta \vdash {}^*\mathsf{NOW}_i \not\sqsubset t$.

Proof. Consider the $i-1^{st}$ execution of algorithm `setup_new_MTF`. Note that this is the execution responsible for the transition from $^*\mathsf{NOW}_{i-1}$ to $^*\mathsf{NOW}_i$. By statement 2 and Definition 5.4 (MTFs), $t \in \Phi(^*\mathsf{NOW}_{i-1})$ and, by statement 1, t is a state interval. Thus, we only need to consider steps 12 through 16 of the algorithm. If, at step 13, $\beta \vdash {}^*\mathsf{NOW}_i \sqsubset t$, then the algorithm ends, and, trivially, $\beta \vdash {}^*\mathsf{NOW}_i \sqsubset t$ following its execution. On the other hand, if the conditional of step 13 is not true, then step 14 is executed. If, at step 14, $\beta \vdash {}^*\mathsf{NOW}_i \not\sqsubset t$, then, trivially, $\beta \vdash {}^*\mathsf{NOW}_i \not\sqsubset t$ following the execution of the algorithm. Otherwise, step 15 gets executed. By statement 5, $\pi(t)$ is defined and, given statement 4, the conditional in 14 may be evaluated. If $\eta(\pi(t)) \leq \mathcal{A}(q)$, then step 15 guarantees that, following the execution of the algorithm, $\beta \vdash {}^*\mathsf{NOW}_i \sqsubset t$. Otherwise, algorithm `event_cease` gets executed with arguments s, t, and t_{tr} (the second argument to `setup_new_MTF`). By step 2 of `event_cease`, $\beta \vdash t \supset\subset t_{\mathrm{tr}}$. Therefore, by **AT7**, $\beta \vdash t \prec t_{\mathrm{tr}}$. By step 10 of `setup_new_MTF`, $\beta \vdash t_{\mathrm{tr}} \prec {}^*\mathsf{NOW}_i$. Therefore, by the transitivity of $\prec$ (**AT2**), $\beta \vdash t \prec {}^*\mathsf{NOW}_i$. By **TT1**, it follows that, following the execution of the algorithm, $\beta \vdash {}^*\mathsf{NOW}_i \not\sqsubset t$. $\square$

We now prove the main result.

Theorem 8.5 *For every* $t \in \Psi(\mathcal{T})$, $s \in \Psi(\textsc{Temp})$, $q \in \Psi(\mathcal{Q})$, *and* $i \in \mathbb{N}$ $(i > 1)$, *if*

1. $\beta \vdash \mathsf{MHolds}(s,t)$,

2. $^*\mathsf{NOW}_{i-1} \sqsubset t \in \beta$,

3. $\beta \vdash \mathsf{SDur}(s,q)$,

4. $\mathcal{A}(q)$ *is defined,*

5. *for every* $\Phi \in \text{Span}(t)$*,* Φ *is a* NOW*-MTF or a transition-MTF, and*

6. *Cassie determines that s ceases to hold at* $^*\text{NOW}_i$*, then*

7. *for every* $n \in \mathbb{N}$*,* $\beta \vdash t \prec {}^*\text{NOW}_{i+n}$*.*

Proof. By Lemma 8.3, at $[\![{}^*\text{NOW}_i]\!]$ (following the $i-1^{st}$ execution of algorithm `setup_new_MTF`), either $\beta \vdash {}^*\text{NOW}_i \sqsubset t$ or $\beta \vdash {}^*\text{NOW}_i \not\sqsubset t$. Given statement 6 and Definition 5.17, the former cannot be the case.[4] Therefore, following the $i-1^{st}$ execution of algorithm `setup_new_MTF`, $\beta \vdash {}^*\text{NOW}_i \not\sqsubset t$. Following the proof of Lemma 8.3, this may happen either through step 14 or step 16. In either case, the proposition $t \prec {}^*\text{NOW}_i$ gets added to β. By Theorem 5.4 and the transitivity of $\prec$ (**AT2**), for every $n \in \mathbb{N}$, $\beta \vdash t \prec {}^*\text{NOW}_{i+n}$. □

8.3 Conclusions

In this chapter, we proved a number of results about the persistence of states. We can sum up those results as follows.

1. Whenever a bodily state occupies some proprioceptual modality, Cassie believes that it holds (and vice versa).

2. Whenever Cassie continuously perceives a state, she believes that it persists.

3. In the absence of knowledge to the contrary, Cassie believes that a state persists if it has not exceeded its typical duration.

4. Whenever a state is believed to have ceased to persist, the particular event of that state holding is believed to be in the past.

[4]Recall that I am dismissing the possibility of contradictory information about current states (see Section 5.2.1). If this were not the case, then statement 6 may still be accommodated if belief revision algorithms choose to retract the proposition $^*\text{NOW}_i \sqsubset t$ (especially if introduced as a default assumption by step 15 of `setup_new_MTF`) in favor of a newly derived/asserted contradictory proposition.

Chapter 9

Cascaded Acts

In the previous chapters, a detailed account of reasoning about time and situations, temporal progression, and temporal perception has been presented. In this chapter, I turn to a slightly different matter: action execution. The reader, however, should note that this is not totally unrelated to what has been done in the previous chapters. In particular, recall that a primary concern of this work is accounting for what a cognitive agent has in mind as it executes actions on-line. The main utility of the theory of time and situations developed so far is to establish the formal and conceptual grounds required for providing the agent with awareness of what it has done and is doing. A complete investigation of everything that is involved in such awareness is not possible within the scope of this work. Nevertheless, a basic, and hence crucial, issue is investigated in this chapter. The logic of events presented in Chapter 4 will prove particularly useful in addressing the issues raised here.[1]

9.1 The Problem

Before getting into a detailed discussion of what I am trying to do in this chapter, let me first briefly point out what it is that I am *not* doing. First, I do not address planning problems. I assume that Cassie has a prestored library of plans (or recipes, à la (De Eugenio, 1998)) or a collection of high-level programs (Levesque et al., 1997) that it uses to achieve its goals. Second, I do not address problems that may arise from various concurrent processes communicating with each other to affect the behavior of the agent. Finally, in this chapter, I do not consider errors and interrupts that may

[1] The work presented in this chapter is based on (Ismail and Shapiro, 1999).

happen during acting (but see Chapter 10). Nevertheless, the approach has been designed with these considerations firmly in mind. The problem being addressed here is, I believe, more fundamental.

Consider an agent that is performing a sequence of acts $\langle \alpha_1 \ldots \alpha_n \rangle$. The agent may be involved in such an activity either because it was instructed to do so, or because it was asked to achieve some goal for which it needs to perform that sequence of acts as indicated by some recipe. What does performing a *sequence* of acts mean? Intuitively, to perform α_1 and then α_2 and then α_3, etc. The simple concept of performing one act and *then* another seems to be a very fundamental one. However, when considering what it requires to actually behave in such a manner, it turns out that this concept is indeed fundamental but far from simple, or at least not as simple as it may seem. For example, consider the following instructions:

(9.1) Pick up the block and then walk to the table and then put the block on the table.

(9.2) Run to the store and then buy a bottle of milk and then run back here.

(9.3) Stick a stamp on the envelope and then bring the secretary here and then give her the envelope.

In the three cases, the English word *then* stands for an important constraint on how to behave according to the given instruction. In particular, in (9.1) for instance, Cassie should first *start* picking up the block and it should start walking to the table when, and only when, it is holding the block. Similarly, it should start putting the block on the table when and only when it is near the table. That is, the above instructions could be more explicitly represented as follows:

(9.1′) Start picking up the block; and when you are holding the block, start walking to the table; and when you are near the table, start putting the block on the table.

(9.2′) Start running to the store; and when you are in the store, start buying a bottle of milk; and when you have bought a bottle of milk, start running back here.

(9.3′) Start sticking a stamp on the envelope; and when there is a stamp on the envelope, start bringing the secretary here; and when the secretary is here, start giving her the envelope.

As should be obvious, the above versions of (9.1), (9.2), and (9.3) are linguistically awkward. They, however, represent a more detailed account of what an agent would do in order to behave correctly.

There are two main differences between the two versions of the instructions. First, note the use of *start*. Indeed, the only thing that an agent is guaranteed to do is to start performing some act.[2] Whether it will actually perform the act or not depends on many factors including the occurrence of errors and various environmental conditions. Second, the agent moves on to the next step in a sequence of acts when and only when it has completed the previous one. The important point here is that the agent does not start the first act and then start the second and then the third etc.; it must start the i^{th} act only when the $(i-1)^{st}$ act has been successfully completed.

In this chapter, I propose a theory for the performance of sequences of acts. Given instructions as those in (9.1)-(9.3) (either in a formal or a natural language), the theory ought to ensure that the agent would actually behave along the lines of (9.1′)-(9.3′). One important feature of the theory (and, for that matter, any theory that would address the same issue) is that perception, proprioception, and reasoning should all be integrated into the acting system in such a way that they actually direct the execution of acts. In particular, knowing when an act is complete is a reasoning process that in many cases is initiated by perception and/or proprioception (see the above examples). It is this conscious awareness of what it has done that directs the agent in its execution of a sequence of acts.

Now, the reader might (rightfully) wonder why there is a problem in the first place. After all, there have been numerous AI applications where robots adequately performed the tasks required of them without the need for a sophisticated theory of sequential acting. How were these robots at all successful? To answer this question, it might be sufficient to review the relevant literature and discuss how existing acting systems approach the issue of sequencing, which is done in Section 9.2. However, part of the answer to the question could be provided without delving into the details of particular systems. Granted, years of research have led to the development of successful robots that do not seem to have any problems performing a sequence of acts. But that can only be said based on the observable behaviors of those robots. Correct behavior could be achieved in a number of ways, but is the underlying methodology robust? Is it theoretically well-worked-out so that it may be generalizable to novel, more ambitious domains? The problem that we are addressing here is not just how to achieve appropriate behavior, but how to explain it and flesh out the fine details that contribute, as generally as possible, to the correct execution of a sequence of acts. At the core of our problem is determining what constitutes the completion of an act. Part of Chapter 4 was dedicated

[2]Where *start* could be interpreted as *intend*.

to developing an ontology of events, of which acts are a special case, categorized according to their completion conditions. This is an issue that has not been addressed before in any depth and that, I believe, should be carefully examined to develop better action executors.

As shall be illustrated in Section 9.2, there have been two main attitudes toward the issue of action execution within the AI community. The first demotes action execution to the level of implementation details that are typically not accounted for in the formal theory of action. This essentially means that there is no *principled* method by which the correct sequencing of acts is achieved. The second attitude grants execution more attention but only considers acts for which no reasoning is required to determine when they are complete. This facilitates building robots where the reasoning system does not have to worry about the sequential execution of acts; it is notified by a lower-level component when it should issue the next act in the sequence. Nevertheless, such systems are obviously not robust since they only work under the assumption that no reasoning is needed to determine completion. As we shall argue below, this is far from reasonable.

9.2 Related Work

As far as I know, nobody else has *explicitly* addressed the problem of sequential acting before. To start with, research oriented towards developing languages for representing and reasoning about actions and plans (see Chapter 2) does not say anything about the correct execution of sequences of acts in real time in the real world.[3] For example, see (Traverso and Spalazzi, 1995), (Artale and Franconi, 1998), or (Chen and De Giacomo, 1999) for recent proposals with various concerns. In these systems, there is some construct in the formal language denoting sequences of acts.[4] The semantics of these constructs either specify temporal constraints on the component acts or implicitly use variants of the English *then*. No mention is made of acts being complete and agents being aware of it. I believe that such issues should be accounted for in the semantics of the action language.

As pointed out in Chapter 2, research within the GOLOG family (including standard GOLOG (Levesque et al., 1997), its concurrent version: CONGOLOG (De Giacomo et al., 2000), and its temporal version (Reiter, 1998)) may be divided into those versions of GOLOG with an off-line interpreter and those with an on-line interpreter. With an off-line interpreter, the output of the

[3]In the sense of $(9.1')$-$(9.3')$.

[4]Indirectly in the case of (Traverso and Spalazzi, 1995).

system is a sequence of actions that would lead to the goal sought should they be executed in order. However, execution is assumed to be carried out by a different component of the system and there is no mention of how sequences of acts would actually be performed in real time in the world.

On-line interpreters of GOLOG (De Giacomo et al., 1998; De Giacomo and Levesque, 1998) account for the actual execution of actions and recovery from errors that may affect their outcome:

> The robot is executing a program on-line. By this, we mean that it is physically performing the actions in sequence, as these are specified by the program. *After* each execution of a primitive action or of a program test action, the execution monitor observes whether an exogenous action has occurred. (De Giacomo et al., 1998, pp. 453-454; my emphasis)

Nevertheless, nothing is mentioned about how the monitor knows that the action has been successfully completed. The implicit assumption is that when the agent is acting, it is acting, and control returns to the main execute-monitor loop when the action is finished (see the sample Prolog implementations presented in (De Giacomo and Levesque, 1998) and (De Giacomo et al., 1998)). This would work so long as the action is merely simulated.[5] However, if the action initiates certain activities in a hardware robot, the monitor should *wait* for these activities to terminate not merely for their initiation to be over. The mail-delivery GOLOG agent described in (Lespérance et al., 1998), which is an actual hardware robot, seems to satisfy this requirement. The following outlines the gist of the recipe presented therein for trying to serve a customer.

1. `Start going to the customer.`
2. `Wait until you are not moving.`
3. `If you have reached the customer, then serve the customer.`
4. `Else if you are stuck, then handle failure.`

This looks very much like the sequences in $(9.1')$-$(9.3')$ and indeed results in the appropriate behavior. Nevertheless, it does not provide any insights into the general issue of sequential acting; the correct way to act is explicitly represented in the recipe rather than being implemented in the

[5] This is not to say that *all* GOLOG applications use software simulations. For example, see (Tam et al., 1997) where a discussion of using GOLOG in the control of two hardware robotic platforms is presented.

semantics of the GOLOG sequencing operator. New sequences would also have to be explicitly expanded in the same way, thereby missing important generalizations that could be made regarding the structure of sequential acting. In addition, such representations are also far removed from natural linguistic instructions (note the awkwardness of (9.1$'$)-(9.3$'$)), a consequence that I would like to avoid in building cognitive agents.

The only place where the completion of a component act is explicitly stated in the semantics of sequences is (Davis, 1992, p. 41). Informally, he states that a sequence of two acts is active if the first is active and not yet completed or the second is active and the first is completed. This suggests (but does not guarantee) that the second act may only be activated when it is known that the first one has been completed. (Davis, 1992) also mentions that for some acts (what he calls *finite acts*), the reasoning system is notified that they are complete:

> The command to execute [finite acts] is shipped off to a black box: an effector control unit, or a separate module, or even a lower-level planner. It is assumed that the black box knows what it means to "begin", "continue", "interrupt", "resume", and "complete" such an act, and can report to the plan interpreter when the act has been completed. (Davis, 1992, p. 41)

We may interpret this as an indication of some sort of bodily feedback (where the body is Davis's black box). Nevertheless, this issue is not explicitly related to the execution of sequences. In addition, given Davis's discussion, it appears that the completion of a finite act is somehow determined by the cessation of bodily activity. As will be pointed out below, this is not sufficient for all types of acts, since some *reasoning* might be needed to determine completion.

Research on interleaving sensing, planning, and execution (Georgeff and Lansky, 1987; Ambros-Ingerson and Steel, 1988; Shanahan, 1998, for example) addresses issues that are related to our problem here (see Chapter 2 for a detailed discussion of these systems). In these systems, planning and execution are interleaved (or performed concurrently) in a producer-consumer kind of model. When the planner reaches a primitive act, the executive performs it. The sensory system updates some knowledge base with changes in the environment that are taken into consideration by the planner. It is not clear, however, how the planner knows when to send a new primitive act to the executive, essentially the problem of when to move to the next step in a sequence. For example, Shanahan discusses the *sense-plan-act cycle*:

> The robot's sensors are read, then a bounded amount of time is allotted to processing sensor data before the robot moves on to planning. Then a bounded amount of time is given over to planning before the robot proceeds to act. *Having acted*, the robot consults its sensors again, and the cycle is repeated. (Shanahan, 1998, p. 8; our emphasis)

How does the robot know that it *has acted*? It is not clear how to answer this question in light of Shanahan's discussion. It is, however, clear that the sense-plan-act cycle is not consciously controlled by the agent, i.e., it is not encoded in Shanahan's logic. Apparently, some *subconscious* mechanism controls the cycle. Why is that a problem? It is a problem because, as will be argued below, it is sometimes *necessary* to reason about whether an act is complete.

Reviewing previous work, we obviously needed to *read between the lines*, in an attempt to come up with answers to the question of how the discussed systems address the problem of sequential acting outlined in Section 9.1. This shows that the problem has not been explicitly addressed. Even though the discussed systems certainly have their means of overcoming, or overlooking, the problem, there is yet no generalized theory to be found.

One might think that there is nothing fundamentally deep about this problem. In particular, it may be suggested that concurrent processing can totally eliminate any problems with sequential acting. For example, consider the following system. There are two concurrent processes, p_1 and p_2. p_1 carries out reasoning and possibly interacting with another agent (typically, a human user). p_2 controls the execution of sequences of acts and has direct access to the status of the body (this is, more or less, the organization assumed in (Davis, 1992)). Suppose that the agent is instructed to perform a sequence of acts. Having received the instruction, p_1 sends it to p_2, which starts executing the sequence. p_2 initiates one act, sleeps until the body finishes execution, and then initiates the next act. This keeps on repeating until the whole sequence has been performed. In the meantime, p_1 is active, available for interacting with the user and can, at any time, interrupt p_2, causing it to stop the execution of the sequence.

Although this seems to solve the problem, it actually does not. The main problem is how p_2 operates. Note that p_2 initiates an act when the current activity of the body terminates. However, just termination is no guarantee that the act has actually been performed successfully. Something might have gone wrong during execution and it may not be appropriate to assume that the act has been completed. Suppose for the sake of the argument, however, that p_2 can somehow tell whether the act actually succeeded. That is, p_2 initiates an act only when the body has actually carried out the previous one. In this case, such a system might be sufficient for *some* sequences. For example, it may correctly perform the sequence in (9.1). However, consider (9.2). The act “buy a bottle of milk” does not seem to be primitive. That is, it is not one continuous bodily activity; the agent would need to reason and follow a plan in order to buy a bottle of milk. Such a plan may be another sequence of acts including walking to the dairy section, picking up a bottle of milk, walking to the cashier, paying the money, and possibly interacting with other agents. Merely monitoring the state of the body is obviously not sufficient for p_2 to tell when to initiate the act following that of buying the milk. The agent has to *know*, using both bodily feedback and reasoning, that the *goal* of buying a bottle of milk has been achieved: according to (9.2′), that it has bought a bottle of milk. Only then could the agent move on to the next step.

Consider another example. An interactive tutoring system is given the following high-level instruction: “Explain the problem to the student and then give a test”. Determining when the test should be given is based on whether the student has understood the problem. But determining the latter cannot be done without reasoning; the agent needs to reason, and even act, in order to figure out whether the first act in the two-act sequence has been completed. No acting system that we are aware of offers a principled, domain-independent way to do this.

Even more interesting is the sequence in (9.3). The act “bring the secretary here” may be performed by the agent calling the secretary, for instance. However, once it does that, its body is no longer actively doing anything. This, by no means, should make the agent move on to the next step; the secretary has to actually arrive, an event that can happen at any time and that is not under the control of the agent. In general, the same applies to any $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ act (see Section 4.6)

Waiting until bodily activities terminate is obviously not sufficient to initiate the next step in a sequence of acts. Nor is it sufficient to depend on the sub-conscious body to notify the planner/reasoning system that an act is complete; determining whether an act is complete often requires reasoning. We need a theory of an agent that is conscious of what it is doing, aware of the outcome of its activities, and whether they actually achieve their intended goals. In the rest of the Chapter, such a theory is proposed.

9.3 Primitive and Composite Acts

At this point, an important distinction must be made—that between *primitive* and *composite* acts. But, first, a word on what exactly is meant by "act" here. Technically, by "act" I mean a particular event category. Intuitively, however, not any event category qualifies as an act: an act is an event category involving an agent carrying out some process. In this section, I'm concerned with a particular classification of acts, one that should be familiar but that, nonetheless, requires explicit discussion at this point. An act is either primitive or composite. Two notions are complementary; in order to define one, it suffices to define the other. In what follows, I shall attempt to give an informal characterization of primitive acts; acts that do not conform with such a characterization shall be considered composite.

Consider two ways by which we may characterize primitive acts. We may call these the *epistemic* characterization and the *kinesthetic* characterization. For the former, an act is primitive if the agent does not need to (or even *cannot*) be "told" how to perform it. For instance, it is hard to linguistically explain to an agent how to ride a bicycle (or tie a shoe lace): maybe show it, but not tell it. Accordingly, such an act may be considered epistemically primitive. On the other hand, one may explain to the agent how to cross the street, for instance: press the button, wait for the walk light, and then walk to the other side. Crossing the street would therefore be considered epistemically composite. In general, an act is epistemically primitive if the agent knows *how to* do it but cannot (easily) reason about how it does it. This is, more or less, the distinction between procedural and declarative knowledge.

Kinesthetic characterization of primitive acts is based on the relation between an agent's intentions and its body. Here, an act is primitive if the agent has no control over its performance. For instance, I may intend to move my arm, and I can do that by contracting my muscles in a certain way. I have some control over the degree and speed of these contractions, and I can interrupt the motion of my arm by simply deciding to do so. Nevertheless, I do not have full control over the whole process. The movement of the arm is made up of bits of *events* that are predetermined by our neuro-muscular make-up. I do not have control over, or awareness of, the quality, duration, or speed of such events; neither can I interrupt them once they start. Actions that directly reduce to such uncontrollable events are what we may call kinesthetically primitive. Anything else is kinesthetically composite. In this case, the epistemically primitive acts of riding the bicycle or tying the shoe lace would be kinesthetically composite. However, if an act is kinesthetically primitive then it is also epistemically primitive.

The above characterizations are just two examples of what a primitive act may be. The philosophy of action literature contains various possible accounts for what may be a *basic* act; a notion that is similar (if not identical) to that of primitive acts (Goldman, 1970; McCann, 1998, for instance). This research, however, is primarily concerned with *human* action. What is reasonable to assume about humans need not be suitable for other agents, robots for instance. Humans are extremely complex agents; they are provided with a set of primitive acts that could be combined in various ways to yield a large set of composite acts. This is required because of the complex environment in which humans exist. Other agents might exist in less demanding environments and therefore need not be as complex. In particular, computational agents are usually designed to operate in relatively simple environments. In such environments, due to the limited number of behaviors expected from the agent, primitiveness may be very coarsely defined. For example, finding a red robot, making spaghetti, and giving coffee to a person are considered primitive acts in the systems described by Shapiro (1998), Artale and Franconi (1998), and Reiter (1998), respectively. Such acts are arguably not primitive for humans (not even epistemically primitive).

The main point is that an act's being primitive or composite depends on the very nature of the agent. Following (Shapiro et al., 1989), I assume that "[a]ny behaving entity has a repertoire of *primitive actions* it is capable of performing" (Shapiro et al., 1989, original emphasis). In designing an artificial agent, one has to make decisions regarding which acts are primitive and which are not. The notion of primitiveness that I adopt in this work has mainly an epistemic nature but also has some kinesthetic features. In particular, I make the following assumptions (**P** stands for *primitive*).

P1. Cassie can perform any of her primitive acts; she cannot reason about how she performs them.

P2. When performing a primitive act, Cassie is aware that she is performing the act. Nevertheless, she has no conscious awareness of its progression, nor of its different stages if it has any. Note that this is not a very unreasonable assumption; people, with enough skills, can perform certain acts while their attention is totally directed to something else. I view Cassie to be skillful enough to carry out her primitive acts without any intellectual interference.

P3. In principle, Cassie may interrupt her performance of a primitive act at any point. Of course, "point" here is very coarse-grained; Cassie cannot interrupt kinesthetically primitive acts (the switching of a relay, for instance).

For example, we may consider running to the store a primitive act. That is, running, reaching the store, and stopping are not consciously planned and serialized by Cassie; she just knows *how* to "run-to-the-store". In particular, there is a PML routine that implements the process of running to the store. The innards of this routine are not accessible at the KL, and, thus, Cassie has no beliefs about how running to the store takes place. On a different account, running to the store may be a composite act involving Cassie's continuing to run based on explicit beliefs about her position with respect to the store. The SNePS acting system allows us to adopt such an approach and to reduce all (physical) primitive acts to basic *movements* as suggested in (Israel, 1995).

In general, an act is represented in the logic by some term in $\overset{\bullet-\bullet}{\mathcal{EC}}$—a durative event category. For a primitive act, $\overset{\bullet-\bullet}{ec}$, there is a PML procedure associated with it. More precisely, $\mathcal{A}(\overset{\bullet-\bullet}{ec})$ is defined and evaluates to a particular PML routine that directs Cassie's body to doing whatever a performance of $[\![\overset{\bullet-\bullet}{ec}]\!]$ requires (see Section 7.4.2). On the other hand, if $\overset{\bullet-\bullet}{ec}$ is a composite act, then $\mathcal{A}(\overset{\bullet-\bullet}{ec})$ is undefined. Now, if $\overset{\bullet-\bullet}{ec}$ is an act that Cassie is supposed to be able to perform (rather than just reason and talk about it), then Cassie must have a plan for how to perform it. The plan should eventually reduce to the performance of primitive acts. For more on the representation of plans, see (Kumar and Shapiro, 1994a; Kumar and Shapiro, 1994b; Kumar, 1994).

9.4 Telicity Revisited

At the core of the problem of sequential acting presented in Section 9.1 is the issue of act completion. In Section 4.6, the issue was discussed in detail. To recapitulate, whether an act token (and, in general, an event token) is complete depends on how it is categorized. In particular, it depends on whether it is categorized as telic, $\overleftarrow{\mathsf{atelic}}$, $\overrightarrow{\mathsf{atelic}}$, or $\overleftrightarrow{\mathsf{atelic}}$.[6] Recall that an act is telic if there is a state whose onset signals its completion. It is $\overleftarrow{\mathsf{atelic}}$ if a state puts an upper temporal bound on its completion. Similarly, it is $\overrightarrow{\mathsf{atelic}}$ if a state puts a lower temporal bound on its completion.

[6] The distinction between $\overset{\rightarrow\bullet}{\mathsf{telic}}$ and $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ acts does not play a significant role at this level of the analysis, but see Section 9.8.

Otherwise, the act is $\overleftrightarrow{\mathsf{atelic}}$.

In this section, I will illustrate, in some detail, how the system presented in the previous chapters may allow Cassie to believe that one of her acts is complete. I shall do this rather informally since a formal *proof* would be tedious and hard to follow. Below, I will only consider the case of a primitive durative act; a more comprehensive analysis is provided in Section 9.8.

1. Cassie decides to perform an act $\overset{\bullet-\bullet}{ec}$.

2. The acting system introduces a new event token, $\overset{\bullet-\bullet}{e}$, representing the particular performance of $\overset{\bullet-\bullet}{ec}$ that Cassie is about to undertake (see Section 9.8 for a more precise characterization of how the acting system does this).

3. The PML process $\mathcal{A}(\overset{\bullet-\bullet}{ec})$ is initiated.

4. As $\mathcal{A}(\overset{\bullet-\bullet}{ec})$ starts execution, it invokes algorithm `state_change` with the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ in $S^{\uparrow}$ (see Section 4.6.2 for Int). Note that $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ is a bodily state. That is, $\mathrm{Mod}_{\mathrm{prop}}(\mathsf{Int}(\overset{\bullet-\bullet}{e})) \neq \{\}$.

5. Following the execution of `state_change`, by Theorem 5.5, $\beta \vdash \mathsf{MHolds}(\mathsf{Int}(\overset{\bullet-\bullet}{e}), t)$, for some newly-introduced $t \in \Psi(\mathcal{T})$.

6. By executing algorithm `state_start` in step 6b of `state_change`, a new event token, $\overset{\bullet}{e}_1$, of category $\uparrow \mathsf{Int}(\overset{\bullet-\bullet}{e})$ is introduced. Note that, by step 2 of algorithm `event_start`, $t_1 \supset\subset t \in \beta$, where t_1 is the transition interval associated with $\overset{\bullet}{e}_1$ (this is the local variable t_{tr} generated by `state_change`).

7. As $\mathcal{A}(\overset{\bullet-\bullet}{ec})$ finishes, algorithm `state_change` is invoked with the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ in $S^{\downarrow}$. Note that, by the third principle of change (see Section 5.2.5), there must be some state $s \neq \mathsf{Int}(\overset{\bullet-\bullet}{e})$ in $S^{\uparrow}$, such that $\mathrm{Mod}_{\mathrm{prop}}(\mathsf{Int}(\overset{\bullet-\bullet}{e})) \cap \mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$. By Axiom 5.10, $\beta \vdash \mathsf{Holds}(s, {*}\mathsf{NOW})$ implies $\beta \vdash \neg\mathsf{Holds}(\mathsf{Int}(\overset{\bullet-\bullet}{e}), {*}\mathsf{NOW})$.

8. By executing algorithm `setup_new_MTF` in step 8 of `state_change`, a new event token, $\overset{\bullet}{e}_2$, of category $\downarrow \mathsf{Int}(\overset{\bullet-\bullet}{e})$ is introduced (step 14 of `setup_new_MTF`). Note that, by step 2 of algorithm `event_cease`, $t \supset\subset t_2 \in \beta$, where t_2 is the transition interval associated with $\overset{\bullet}{e}_2$ (again, t_2 is the local variable t_{tr} generated by `state_change`).

9. Since $\beta \vdash \mathsf{MHolds}(\mathsf{Int}(\overset{\bullet-\bullet}{e}),t)$, then, by **AS2** and **AS3**, $\beta \nvdash \neg\mathsf{Holds}(\mathsf{Int}(\overset{\bullet-\bullet}{e}),t')$, for any t' between t_1 and t_2.

10. By **TOC4**, $\beta \vdash \mathsf{NoOccurs}(\downarrow \mathsf{Int}(\overset{\bullet-\bullet}{e}),t_1,t_2)$. It follows that $\beta \vdash \mathsf{OCPair}(\overset{\bullet}{e}_1,\overset{\bullet}{e}_2,\mathsf{Int}(\overset{\bullet-\bullet}{e}))$.

    Digression

    At this point, the type of the act $\overset{\bullet-\bullet}{ec}$ comes into play. In particular, given the type of the act ($\overset{\rightarrow\bullet}{\mathsf{telic}}$, $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$, $\overset{\longleftarrow}{\mathsf{atelic}}$, $\overset{\longrightarrow}{\mathsf{atelic}}$, or $\overset{\longleftrightarrow}{\mathsf{atelic}}$), Axioms **AE14–AE18** would determine whether the durative event $\mathsf{Clos}(\overset{\bullet}{e}_1,\overset{\bullet}{e}_2)$ is of category $\overset{\bullet-\bullet}{ec}$. If $\overset{\bullet-\bullet}{ec}$ is $\overset{\longleftrightarrow}{\mathsf{atelic}}$, then **AE18** readily implies $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{e}_1,\overset{\bullet}{e}_2),\overset{\bullet-\bullet}{ec})$. Otherwise, there is some state, s, such that whether $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{e}_1,\overset{\bullet}{e}_2),\overset{\bullet-\bullet}{ec})$ follows is based on the temporal position of its onset with respect to t_2. In particular, note that if $\overset{\rightarrow\bullet}{\mathsf{telic}}(\overset{\bullet-\bullet}{ec},s)$, then, for $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{e}_1,\overset{\bullet}{e}_2),\overset{\bullet-\bullet}{ec})$ to be the case, s would typically be a member of $S^{\uparrow}$ in the invocation of `state_change` indicated in 7 above.[7]

    But, recalling a previous discussion in Section 4.6.4, it should be noted that *inferring* that the onset of s is appropriately positioned may not be possible. In particular, note that, for Axioms **AE14–AE18** to be usable at this point, Cassie should infer that no onset of s occurs within the entire interval, or an appropriate subinterval in case of $\overset{\longrightarrow}{\mathsf{atelic}}$ acts, bounded by $\overset{\bullet}{e}_1$ and $\overset{\bullet}{e}_2$. Now, Cassie may indeed by unaware of any such occurrence, but that does not allow her to infer that it does not exist. To get around this problem, there are two routes we can take (as pointed out in Section 4.6.4): (i) drop the problematic NoOcc clause from **AE14–AE18** or (ii) introduce non-monotonicity. The latter is certainly the ultimate solution to the problem and, in what follows, I will be assuming the system includes non-monotonic reasoning capabilities. Note that this is done in order to illustrate how the axioms may theoretically be used by Cassie to infer act completion; what is actually done is slightly different as shall be shown below. The former route would result in Cassie correctly behaving if the primitive actions are implemented appropriately. In particular, the PML procedure associated with a primitive act ($\mathcal{A}(\overset{\bullet-\bullet}{ec})$ above) should terminate whenever the appropriate state, s, starts to hold. Thus, Cassie would stop running to/toward the store once she is at the store (more on this in Section 9.5). This way, we may drop the NoOcc clauses while ensuring that Cassie would not hold beliefs

[7] At this point it should be clear why a unique transition interval is shared by all onset and cessation events introduced through `state_change`. In addition to the assumption that these events are simultaneous, this also serves the purpose of ensuring that the onset of s takes place at the same time as $\overset{\bullet}{e}_2$, which is needed to satisfy the antecedent of **AE14**.

about her actions that contradict what she actually did. In a sense, the non-monotonicity required is implicitly built into the system. The problem, of course, is that changing the axioms thus would render them incorrect and might lead to wrong conclusions in case the events involved are not Cassie's own actions. In any case, I will argue below that, for time-critical action, the reasoning steps outlined here are too lengthy to be practical; short-cuts that would allow us to avoid the NoOcc problem altogether will be presented.

With the assumption of non-monotonicity, I will be assuming that the onset of s is appropriately positioned so that $\beta \vdash \mathsf{Cat}(\mathsf{Clos}(\dot{e}_1, \dot{e}_2), \overset{\bullet-\bullet}{ec})$. Note that, if it is not, then Cassie would, rightfully, not be able to conclude that $\overset{\bullet-\bullet}{e}$ is complete.

End of Digression

11. Given an appropriate instance of schema (4.26) (repeated below for convenience) and assuming that $\overset{\bullet-\bullet}{e}$ was introduced with $\overset{\bullet-\bullet}{ec}$ as its only category, $\beta \vdash \mathsf{Equiv}(\overset{\bullet-\bullet}{e}, \mathsf{Clos}(\dot{e}_1, \dot{e}_2))$.

$$(4.26)\ [\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{e})) \wedge \bigwedge_{\overset{\bullet-\bullet}{ec} \in \mathcal{C}(\overset{\bullet-\bullet}{e})} \mathsf{Cat}(\mathsf{Clos}(\dot{E}v_1, \dot{E}v_2), \overset{\bullet-\bullet}{ec})] \Rightarrow \mathsf{Equiv}(\overset{\bullet-\bullet}{e}, \mathsf{Clos}(\dot{E}v_1, \dot{E}v_2))]$$

There are two things to note here.

(a) Assuming that $\overset{\bullet-\bullet}{ec}$ is the only category of $\overset{\bullet-\bullet}{e}$ at the time $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ starts to hold is reasonable and typical of acts. In particular, note that Cassie is instructed (or, for some reason, decides) to perform $\overset{\bullet-\bullet}{ec}$, and $\overset{\bullet-\bullet}{e}$ is only introduced as a particular instance thereof.

(b) As pointed out in Section 4.6.2, "[f]or a given event, $\overset{\bullet-\bullet}{e}$, Cassie comes to believe an appropriate instance of (4.26) once she determines that the state $\mathsf{Int}(\overset{\bullet-\bullet}{e})$ holds". Coming to believe such an instance of (4.26) may reasonably be the responsibility of algorithm `Forward`. A revision of the algorithm is shown in Figure 9.1.

12. By **AE13**, it follows that $\beta \vdash \exists Tv'[\mathsf{Holds}(\mathsf{Complete}(\overset{\bullet-\bullet}{Ev}), Tv')]$.

13. If $\overset{\bullet-\bullet}{ec}$ is not $\overset{\to\ldots\bullet}{\mathsf{telic}}$, then by (4.24) (repeated below for convenience), the permanence of Complete, and the fact that the occurrence of $\overset{\bullet-\bullet}{e}$ is in the past (step 10 of `setup_new_MTF`), it follows that $\beta \vdash \mathsf{Holds}(\mathsf{Complete}(\overset{\bullet-\bullet}{e}), {}^{*}\mathsf{NOW})$.

$$(4.24)\ [\mathsf{Occurs}(Ev, Tv) \wedge Tv \prec Tv'] \Rightarrow \mathsf{Holds}(\mathsf{Complete}(Ev), Tv')$$

Algorithm `Forward`$(P \subseteq \Psi(\mathcal{P}))$

1. $P_{\text{inf}} \longleftarrow$ `Forward_old`(P).
2. For every $p \in P \cup P_{\text{inf}}$
   3. If $p = \mathsf{Holds}(s,{}^*\mathsf{NOW})$, for some $s \in \Psi(\text{T}{\small EMP})$, then
      4. If there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s,t)$, then `state_present`$(s,t)$.
      5. Else
         5a. Pick some $t \in \mathcal{T}$ such that $t \notin \Psi(\mathcal{T})$.
         5b. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t), {}^*\mathsf{NOW} \sqsubset t\}$.
         5c. `state_present`(s,t).
      6. If there is some $\overset{\bullet-\bullet}{e} \in \Psi(\overset{\bullet-\bullet}{\mathcal{E}})$ such that $s = \mathsf{Int}(\overset{\bullet-\bullet}{e})$, then

$$\beta \longleftarrow \beta \cup \{[\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{e})) \wedge \bigwedge_{\overset{\bullet-\bullet}{ec} \in \mathcal{C}(\overset{\bullet-\bullet}{e})} \mathsf{Cat}(\mathsf{Clos}(\dot{E}v_1, \dot{E}v_2), \overset{\bullet-\bullet}{ec})] \Rightarrow \mathsf{Equiv}(\overset{\bullet-\bullet}{e}, \mathsf{Clos}(\dot{E}v_1, \dot{E}v_2))]\}.$$

Figure 9.1: Algorithm `Forward`: Version 3.

14. On the other hand, if $\overset{\bullet-\bullet}{ec}$ is $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$, then:

    (a) It follows from **AE8** that, prior to the occurrence, of $\overset{\bullet-\bullet}{e}$, it is not complete.

    (b) By **AOC5**, there is some transition interval, t_3, at which the state $\mathsf{Complete}(\overset{\bullet-\bullet}{e})$ starts to hold.

    (c) By the definition of $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}(\overset{\bullet-\bullet}{ec}, s)$, $\beta \vdash \mathsf{Equiv}(t_2, t_3)$.

    (d) Since t_2 is in the past, and given the permanence of $\mathsf{Complete}$, it follows that $\beta \vdash \mathsf{Holds}(\mathsf{Complete}(\overset{\bullet-\bullet}{e}), {}^*\mathsf{NOW})$.

Evidently, whenever appropriate, Cassie would be able to conclude that an act of hers is complete. If such an act is a step in a sequence, then inferring that it is complete (as per the above steps) would be the sign for her to move on to the next step in the sequence. Nevertheless, the reader should note the complicated and lengthy reasoning steps required to conclude that an act is complete. This might prove impractical in a situation where swift action is required. Fortunately, however, by making some reasonable assumptions, we may side-step most of the above steps. In particular, in many cases, Cassie does not need to explicitly infer that an act is complete in order for her to move to the next step in a sequence of acts. In others, there may be *short-cuts* to reaching the same conclusion. In particular, note that if Cassie is performing some act token $\overset{\bullet-\bullet}{e}$, of category $\overset{\bullet-\bullet}{ec}$, such that $\mathsf{telic}(\overset{\bullet-\bullet}{ec}, s)$, then, assuming that $\overset{\bullet-\bullet}{ec}$ is the only category of $\overset{\bullet-\bullet}{e}$, once s starts to hold, $\overset{\bullet-\bullet}{e}$ should be considered complete and many of the reasoning steps outlined above would not be needed. Informally, I will assume that every act has an associated state which I will call its *goal*. Once the goal of an act starts to hold, then it should be considered complete. In some cases, the goal of an act is simply the state of that act being complete, in others it is not. In the sections that follow, I will investigate this notion of goals in detail. First, we need to consider, in some depth, how telicity interacts with primitiveness.

9.5 The Anatomy of Acts

An agent that knows how to perform its primitive acts should be able to perform any sequence of them. An arbitrary sequence of primitive acts is not primitive, though. Faced with a novel sequence, the agent will have to consciously control its execution. Unlike primitive acts, sequences of them have stages that the agent is aware of; it should proceed to one stage when, and only when, the previous one has been completed. In this section we put together the ideas developed in sections 4.6 and 9.3 and take a close look at the resulting system. In doing so, we develop a deeper understanding of the structure of acts; in particular, the structure of their execution. Such an understanding is essential to the study of the properties different kinds of acts exhibit when they are embedded in a sequence. In section 4.6, we analyzed telicity in terms of two binary features: **R** and **L**. Adding a feature, **P**, for primitiveness, we end up with a three-dimensional feature space: the RLP cube (see Figure 9.2).

Are there acts corresponding to the eight vertices of the cube? And if so, what is the structure of these acts? First, consider the bottom plane, that of composite acts. Consider an agent with two primitive acts: pouring coffee from a pot into a cup and taking one sip of coffee from the cup. In such a case, one can readily come up with four acts to fill the bottom plane (I am using the imperative mood so that the reader can think of these as examples of instructions to act).

(9.4) Drink coffee. $(\overleftrightarrow{\mathsf{atelic}})$

(9.5) Drink no less than three cups of coffee. $(\overrightarrow{\mathsf{atelic}})$

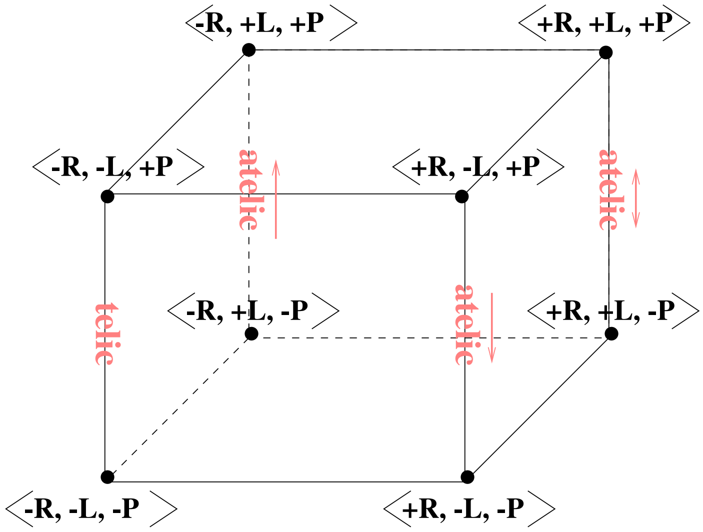

Figure 9.2: The RLP Cube.

(9.6) Drink no more than three cups of coffee. ($\overleftarrow{\text{atelic}}$)

(9.7) Drink three cups of coffee. (telic)

Drinking one cup of coffee is the composite act of sipping from a cup of coffee until it is empty. A precondition for such an act is for the cup to contain coffee, a state that may be achieved by pouring coffee from the pot into the cup. The above acts are therefore composite for such an agent. To act according to (9.4), the agent will start the process of drinking one cup of coffee. Note that since this process is composite, it requires the agent's conscious monitoring and control. The process may then be repeated for an indefinite number of times. At any point, the agent may stop drinking coffee by either finishing one cup and not starting to drink another or by just stopping sipping coffee from a nonempty cup. Such a decision to stop drinking is certainly a conscious one and may be caused by various events including the agent's finding out that no more coffee is left in the pot. Before starting to perform the act, the agent does not have a definite scenario of how it will end, the only thing it knows is that *at some time the act will be complete*.

On the other hand, performing (9.7) requires the agent to not only monitor what it is doing, but also keep track of how many cups of coffee it has drunk to stop drinking when and only when it finishes the third cup. In this case, the agent a-priori knows what completes the act. Performing

(9.5) is a sequence of the performance of (9.7) and (9.4). First the agent drinks three cups of coffee to reach the lower bound the instruction indicates and then continues to drink coffee indefinitely.[8] In this case, like in (9.4), the agent only knows that the act will eventually be complete. However, unlike (9.4), the agent knows that the act will reach completion in a state in which it has drunk three cups of coffee.

Whereas $\overrightarrow{\mathsf{atelic}}$ acts like (9.5) are decomposable into two clearly distinguished components: one telic and another $\overleftrightarrow{\mathsf{atelic}}$, $\overleftarrow{\mathsf{atelic}}$ acts like (9.6) are not as simply structured. The agent starts to drink coffee while consciously monitoring the three-cups upper limit. In that respect, (9.6) is essentially similar to (9.7). However, in the case of (9.6), the agent has one more degree of freedom; it does not have to actually reach the three-cups limit for the act to be complete. Similar to (9.4), prior to finishing three cups, the agent may decide to stop drinking at any point for reasons that may have nothing to do with drinking coffee. It is as if the agent executes both (9.4) and (9.7) in parallel and the act completes whenever one of them does.

Thus, while the execution of $\langle +\mathbf{R}, -\mathbf{L}, -\mathbf{P}\rangle$ acts is structured by the sequencing of telic and $\overleftrightarrow{\mathsf{atelic}}$ acts, the execution of $\langle -\mathbf{R}, +\mathbf{L}, -\mathbf{P}\rangle$ acts is made up of the interleaving, or parallel execution, of telic and $\overleftrightarrow{\mathsf{atelic}}$ acts. This reveals the more complex nature of $\overrightarrow{\mathsf{atelic}}$ and $\overleftarrow{\mathsf{atelic}}$ acts over that of telic and $\overleftrightarrow{\mathsf{atelic}}$ acts. This complexity is formally manifest in the different signs of the **R** and **L** features which reflect the heterogeneous nature of these acts.

Things get more elusive as we move up to the plane of primitive acts. As pointed out in section 9.3, designating an act as primitive or composite by and large depends on the nature of the agent. One might, therefore, argue that we can fill the upper plane by just thinking of an agent for which (9.4)-(9.7) designate primitive acts. However, as I shall argue, such a move is, at least, not obviously valid. Consider our running agent from section 4.6.[9] One might consider (9.8) to be primitive, and hence $\langle +\mathbf{R}, +\mathbf{L}, +\mathbf{P}\rangle$.

(9.8) Run.

[8] Actually, the action is a non-deterministic choice between the above-mentioned sequence and the simple (9.7). This is important since drinking no less than three cups of coffee is compatible with drinking exactly three cups. On how non-deterministic choice among actions is accounted for, see Section 9.6.1.

[9] As mentioned above, we might consider an agent for which (9.4)-(9.7) are primitive. However, since those examples are not even epistemically primitive for humans, I choose to discuss other examples that could be more appreciated by the reader. Such a choice of examples has no effect on the claims to follow.

To perform (9.8), the agent starts running and then may cease to run at any point it decides to. Note that this is compatible with our assumptions **P1**-**P3**: (i) the agent may know how to run but not how it runs, (ii) the agent may be unaware of the different stages of the running process, and (iii) it may interrupt the running at any point thereby putting an end to its activity. Thus, the main difference between primitive and composite $\overleftrightarrow{\textsf{atelic}}$ acts is that the agent has no control over or awareness of the structure of the former but consciously controls and monitors that of the latter. In both cases, however, the agent has the same epistemic status regarding their completion: it only knows that at some state they would be complete, without any further qualification of such a state.

For the same, or a slightly different, agent, (9.9) may be considered primitive.

(9.9) Run to the store.

How would (9.9)-as-primitive be performed? The only way to conceive of (9.9) as a primitive act, is to assume that the agent is designed to reactively (not deliberately) stop when it reaches the store. That is, the agent starts to run and as a reaction to reaching the store it stops running. This involves no cognitive processing. In this case, it is easy to see that (9.9) is compatible with **P1**-**P3**.

Restricting ourselves to the same domain, can (9.10) and (9.11) be considered primitive (and therefore examples of $\langle -\mathbf{R}, +\mathbf{L}, +\mathbf{P}\rangle$ and $\langle +\mathbf{R}, -\mathbf{L}, +\mathbf{P}\rangle$ acts, respectively)?

(9.10) Run toward the store.

(9.11) Run past the store.

First, consider (9.10). Although, as noted above, $\overleftarrow{\textsf{atelic}}$ acts have a complex structure, there are ways to conceive of (9.10) as primitive. A performance of (9.10) may be just a performance of (9.9)-as-primitive. The only difference is that interrupting the first completes the act while interrupting the second does not.[10] It could be shown that such an account is compatible with **P1**-**P3**. Regarding **P1**, the agent need not know how it runs toward the store, in particular it need not know that it does that by simply doing what it would do if it were running *to* the store. That is, the agent knows that it can perform two acts, (9.9) and (9.10), what it is not aware of is how they are both performed and that

[10]Another interpretation is one where (9.10) merely indicates a direction for the running agent. In that case, the agent does not really have to stop if it reaches the store. Although this is fine as far as the interpretation of the particular English sentence goes, it falls outside the category of $\overleftarrow{\textsf{atelic}}$ acts which we are now investigating.

they are both actually performed in the same manner. Regarding **P2**, the agent has no awareness (and need not have any) of the internal structure of the act. Finally, the agent may interrupt the act at any point and like $\overleftrightarrow{\mathsf{atelic}}$ acts, and unlike telic acts, the act would be considered complete.

I am aware that there might be some *pragmatic* difficulties with the above discussion. First, it is not very clear how instructions like (9.10) may be useful in practice. Running toward the store is *usually* something that one realizes *has* happened not something that one *intends* to do. That is, the agent may intend to run *to* the store but end up only running toward the store. In addition, since the existence of a $\overleftarrow{\mathsf{atelic}}$ primitive act would be accompanied by that of a telic act, it does not seem reasonable to make the former primitive. For example, one might consider only (9.9) to be primitive, and performing (9.10) would simply be an intentional performance of (9.9) that may be interrupted. Nevertheless, it should be noted that such concerns, although valid, do not provide any argument against the *logical* possibility of $\langle -\mathbf{R}, +\mathbf{L}, +\mathbf{P} \rangle$ acts. Other examples, may also stand more strongly in the face of the above objections.

(9.12) Pour some coffee into the cup.

(9.13) Lift the block above the table.

Now, let us consider (9.11). One might argue that since (9.8) and (9.9) may be thought of as primitive, then so is (9.11). Similar to our analysis of (9.5), (9.11) may be a sequence of (9.9) and (9.8). However, for (9.11) to be primitive, the sequencing of (9.9) and (9.8) must be hard-wired into the agent; instead of stopping when reaching the store as in (9.9)-as-primitive, the agent *reactively* starts performing (9.8). Such an apparently plausible account is not totally sound though. According to **P3**, the agent should be able to interrupt any of its primitive acts. Now, suppose that, for some reason, the agent's performance of (9.11) is interrupted. A reasonable requirement of an intelligent agent is to remember what it has done. For our agent to know whether it has run past the store, it needs to know whether it had reached the store at some point during the performance of (9.11). This simply means that the agent is aware of the internal structure of (9.11). Obviously this contradicts the assumption of (9.11)-as-primitive since it runs counter to **P2**. Note that this also means that the agent knows (whether explicitly or implicitly, fully or partially) how it runs past the store, which is incompatible with **P1**.

In general, starting from the assumption that an intelligent agent should know the outcome of its activities, if a $\overrightarrow{\mathsf{atelic}}$ act is interrupted, the agent needs to know during which of its two stages the interruption occurred. Such necessary awareness of the structure of an act is contradictory to the notion of primitiveness. Note that this issue does not arise with telic and $\overleftrightarrow{\mathsf{atelic}}$ (also, $\overleftarrow{\mathsf{atelic}}$) acts since for the first, an interruption directly means that the act has not completed; and for the second, an interruption signals the completion of the act. I shall therefore assume that no acts correspond to the $\langle +\mathbf{R}, -\mathbf{L}, +\mathbf{P}\rangle$ vertex of the RLP cube.

To summarize the above discussion we highlight the main points:

1. For atelic acts (those with $+\mathbf{R}$ and/or $+\mathbf{L}$ features), the agent does not know *how* they will end, only that at some point, they will. In particular, atelic acts reach completion when they are interrupted.

2. For telic acts, the agent knows exactly in what state they will be complete.

3. $\langle +\mathbf{R}, -\mathbf{L}, +\mathbf{P}\rangle$ acts are not logically possible.

9.6 Acts and Their Goals

9.6.1 Types of Acts

The RLP categorization of acts does not cover all types of acts that the SNePS acting system, SNeRE, can accomodate. The following presents an overarching classification.

1. **RLP Acts:** These are acts that may be characterized by the RLP features. They are acts that affect the state of Cassie's external environment.

2. **Mental Acts:** These are primitive acts affecting Cassie's mental state; adding or retracting beliefs. Two mental acts are of particular relevance here. The following is their syntax and informal operational semantics; a more precise semantics is introduced in Section 9.8.

   - $\mathsf{believe}: \mathcal{S} \longrightarrow \dot{\mathcal{E}}\mathcal{C}$. For any $s \in \mathcal{S}$, if $\neg\mathsf{Holds}(s,{*}\mathsf{NOW}) \in \beta$, then $\mathcal{A}(\mathsf{believe}(s))$ removes it, adds $\mathsf{Holds}(s,{*}\mathsf{NOW})$ to β, and initiates forward inference.[11]

   - $\mathsf{disbelieve}: \mathcal{S} \longrightarrow \dot{\mathcal{E}}\mathcal{C}$. For any $s \in \mathcal{S}$, $\mathcal{A}(\mathsf{disbelieve}(s))$ removes $\mathsf{Holds}(s,{*}\mathsf{NOW})$ from β.

[11] I am assuming that if $s \in$ ETERNAL, then s is equivalent to $\mathsf{Holds}(s,{*}\mathsf{NOW})$ (see Section 3.6).

Note that believe and disbelieve are punctual. Hence, notions of telicity cannot be applied to them.

3. **Control Acts:** These are primitive acts that control various ways in which a set (possibly a singleton) of acts are to be performed. They are typically used to represent plans for performing composite acts. The following is the syntax and informal operational semantics of the three SNeRE control acts that are relevant at this point (for a complete specification of the control acts provided by SNeRE, see (Shapiro and the SNePS Implementation Group, 1999)).

   - $\mathsf{snsequence} : \mathcal{EC}^n \longrightarrow \mathcal{EC}$ $(n \in \mathbb{N})$. $\mathcal{A}(\mathsf{snsequence}(ec_1, ec_2, \ldots, ec_n))$ executes the acts ec_1, ec_2, ..., ec_n in order.[12]

   - $\mathsf{snif} : 2^{\mathcal{S}\times\mathcal{EC}} \times \mathcal{EC} \longrightarrow \mathcal{EC}$. $\mathcal{A}(\mathsf{snif}(\{\langle s_i, ec_i\rangle\}_{i=1}^{n}, ec_{n+1}))$ executes ec_i for some randomly-chosen i $(1 \leq i \leq n)$ such that $\texttt{Backward}(s)$.[13] If there is no such i, then ec_{n+1} is executed.

   - $\mathsf{sniterate} : 2^{\mathcal{S}\times\mathcal{EC}} \times \mathcal{EC} \longrightarrow \mathcal{EC}$. $\mathcal{A}(\mathsf{sniterate}(\{\langle s_i, ec_i\rangle\}_{i=1}^{n}, ec_{n+1}))$ executes ec_i for some randomly-chosen i $(1 \leq i \leq n)$ such that $\texttt{Backward}(s)$. This process is repeated until there is no such i, where, in that case, ec_{n+1} is executed.

   - $\mathsf{DoOne} : 2^{\mathcal{EC}} \longrightarrow \mathcal{EC}$. $\mathcal{A}(\mathsf{DoOne}(\{ec_i\}_{i=1}^{n})$ executes ec_i for some randomly-chosen i $(1 \leq i \leq n)$.

Note that the domain of a control-act-forming function essentially involves the set $\mathcal{EC}$. This makes control acts conceptually-complex and cannot be expressed in natural language by single-clause constructions. RLP acts, on the other hand, are characterized by being conceptually-simple; $\mathcal{FOCS}$ terms denoting them do not involve elements of $\mathcal{EC}$ as arguments. Linguistically, RLP acts may be expressed by (loosely-speaking) simple clauses and, hence, can exhibit telic features.

[12] Note that, strictly speaking, there is a family of functions $\{\mathsf{snsequence}_n\}_{n\in\mathbb{N}}$, one for each possible arity. However, I choose to be a little sloppy (since sloppiness here does not pose any dangers) and assume that there is one function that may be applied to arbitrary n-tuples.

[13] For simplicity, I assume that `Backward` evaluates to "true" or "false" depending on whether the deduction succeeds. The actual SNePS system has a more sophisticated inference engine that implements a system of message-passing among processes; see (Shapiro and McKay, 1980; McKay and Shapiro, 1981).

9.6.2 Type-Token Associations

Towards the end of Section 9.4, the notion of a goal was introduced. In general, the goal of an act is a state that signals its completion. The notion of goal alluded to here is a very localized one and is tightly associated with the particular categorization of the corresponding act. In particular, it is not the *reason why* the agent performs an act. The goal of an act, α, as far as this work is concerned, is not a pragmatic one; it is that state that signals successful completion of α thereby allowing the agent to start executing other acts contingent upon its completion.[14] For a telic act, things are fairly straightforward; there is a unique state that always holds at the end of a successful performance (and, hence, completion) of a telic act. Thus, the goal of the telic "run to the store", for instance, is the state of being at the store.

But now a question arises: what is a goal associated with, an act category or an act token? For telic acts, one can see that the goal is *token-independent*; all instances of a telic act category have the same goal. More precisely, the goal of a telic act category, $\overset{\bullet-\bullet}{ec}$, is the state s such that $\mathsf{telic}(\overset{\bullet-\bullet}{ec}, s)$. However, things are not always that straightforward; as will be argued below, the goal of some acts can be nothing but the very state of the acts's being complete. But recall that such a state is one of act tokens, not act categories. I will, therefore, define goals for what I call a *type-token association*. This is a pair, $ec[e]$, where ec is an act category and e is a particular instance of it.[15] Thus, a goal is defined for a token under a certain categorization or, alternately, for a tokenized act category. Informally, I will continue to refer to "the goal of an act" with the understanding that goals are formally associated with type-token associations (see Definition 9.1 below).

We now take a closer look at the goals of non-telic acts: atelic, mental, and control acts.

9.6.3 Goals

Consider the goal of an atelic act. The completion of an atelic act (whether $\overset{\longleftrightarrow}{\mathsf{atelic}}$, $\overset{\longleftarrow}{\mathsf{atelic}}$, or $\overset{\longrightarrow}{\mathsf{atelic}}$) is not necessarily simultaneous with the onset of some unique state. Let us first get $\overset{\longrightarrow}{\mathsf{atelic}}$ acts out of the way. As pointed out in Section 9.5, a $\overset{\longrightarrow}{\mathsf{atelic}}$ act, ec, is necessarily composite; it is made up

[14]In McCann's terms (McCann, 1998), this is the *result* of the act. I choose not to adopt this terminology, however, since it only works with $\overset{\rightarrow\bullet}{\mathsf{telic}}$ acts. Goals of $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ acts are *consequences*, according to McCann. Also the term *goal* has been used by some linguists to refer to the same notion (see (Declerck, 1979, p. 762) and the quotes therein).

[15]This will be slightly revised in Section 9.8, where e will not be required to be an instance of ec.

of a sequence $\langle ec_1, ec_2 \rangle$ where ec_1 is telic and ec_2 is $\overleftrightarrow{\mathsf{atelic}}$.[16] Thus, the goal of ec is whatever the goal of the sequence may be—an issue that is discussed below. But the reader may wonder how the state, s (where $\beta \vdash \overrightarrow{\mathsf{atelic}}\,(ec, s)$), required to precede the completion of ec, is accounted for. This state would be accounted for by the choice of ec_1, the telic component of ec. More precisely, ec_1 should be such that there is some s', where $\beta \vdash \mathsf{telic}(ec_1, s')$ and believing that s' holds is sufficient for Cassie to infer that s holds.

According to Section 9.5, the only thing the agent knows about the completion of a $\overleftrightarrow{\mathsf{atelic}}$ act is that it will eventually happen. In particular, this can only take place if the act is interrupted, since the act itself does not have an inherent end-point. Cassie may decide to interrupt an atelic act for different reasons. First, she might interrupt it in order to allocate resources used by the act to other, more important acts. Second, she can interrupt it if she realizes that the act can no longer continue. This could be the result of various kinds of failures that Cassie is capable of perceiving or proprioceiving. Evidently, states in which such interruption events may take place are totally unpredictable. Because of this peculiar feature, the only state that an agent can a-priori rely on to mark the end of an $\overleftrightarrow{\mathsf{atelic}}$ act is the very state of the act's being complete. Thus, for a type-token association, $ec[e]$, such that $\beta \vdash \overleftrightarrow{\mathsf{atelic}}\,(ec)$, the goal is simply the state $\mathsf{Complete}(e)$. Coming to believe that such a state holds may be achieved through the lengthy reasoning procedure outlined in Section 9.4. However, a more efficient mechanism may also be built in the interrupt-handling mechanism.[17]

Given the discussion in Section 9.5, $\overleftarrow{\mathsf{atelic}}$ acts have a dual nature. In some respects they are like telic acts, in other respects they are more like $\overleftrightarrow{\mathsf{atelic}}$ acts. Consider a type-token association, $ec[e]$, such that $\beta \vdash \overleftarrow{\mathsf{atelic}}\,(ec, s)$, for some states s. The performance of e may complete in two ways, reflecting the telic and $\overleftrightarrow{\mathsf{atelic}}$ facets of ec: e completes if the state s starts to hold or if it is simply interrupted. Because of the second possibility, s cannot be considered the goal of $ec[e]$, and, similar to the discussion of $\overleftrightarrow{\mathsf{atelic}}$ acts above, the goal may be nothing but the state $\mathsf{Complete}(e)$. Again, there may be different ways to get around the long reasoning procedure of Section 9.4. One that will be incorporated within the system is for Cassie to hold a belief, while performing e, that the

[16]But see fn. 8.

[17]Of course, one may interrupt an act and then resume it; the interrupt is not a sign of completion in such a case. Chapter 10 presents a cursory account of interruption, and until that is fully developed, I make the assumption that interrupting an act is equivalent to ending it.

fact that the state s holds implies that e is complete (see Section 9.8).

The above seems to provide fairly precise descriptions of what the goals of atelic acts are. What about control acts? Control acts are not the kind of acts that may be described as telic or atelic. As long as the act does not have a *cover term*; i.e., an atomic conceptual structure expressible by a simple sentence; it cannot be said to be telic or atelic (see Section 9.6.1). For example, consider the following sequence describing a recipe for a useful exercise.

$$\mathsf{snsequence(run\text{-}a\text{-}mile, swim, take\text{-}a\text{-}shower)}$$

It is not clear whether one can ascribe any telic properties to such an act without making a category error. Control acts have structures that are far more complex than those of the acts discussed so far. In particular, several states with complex temporal and causal (in case of snif) relations make up the internal structure of control acts. What then is the goal of a control act, for example an arbitrary snsequence? One might argue that the goal of a sequence of acts is the goal of the last act in the sequence. That might sound appealing, but there are reasons not to adopt such a position. For example, consider the following sequence:

$$\mathsf{snesequence(pick\text{-}up\text{-}a\text{-}block, mark\text{-}the\text{-}block, put\text{-}down\text{-}the\text{-}block)}$$

Suppose that the three acts in the sequence are primitive. What is the goal of the put-down-the-block act? Given the localized sense of "goal" that we are assuming, it would be something like *being empty-handed*; this is the state that signals the completion of putting down the block. Intuitively, however, it does not seem right to say that this is the goal of the whole sequence. Being empty-handed may be achieved if the agent drops the block before marking it. This, by no means, signals successful completion of the sequence. In addition, a precondition of the above act is the very state of being empty-handed (to allow for picking up a block); if the goal of performing an act already holds, there is no need to perform it. In fact, in many cases, if the goal of an act already holds, then there is no *way* to perform it (think of running to the store when one is already at the store or putting down a block when one is empty-handed, for instance).

One may object by arguing that an agent might still need to perform an act even if its goal already holds, the reasoning being that an act may be performed to achieve one of its effects (the goal being only a distinguished effect). To support such an objection, one has to come up with a situation in which the goal of some act holds and an agent, nevertheless, needs to perform it for one of its effects. To start with, it is extremely difficult to come up with such a situation. Given the notion of "goal" that we are adopting (i.e., the state that signals the successful completion of an act), it is very hard to describe a realistic situation in which it is appropriate (or even possible) to perform an act while its goal holds. However, let us try to construct such a situation.

> John is a bachelor who leads a pretty dull social life. Due to the nature of his job, he spends the whole day driving his car and visiting customers. To get out of the car, John unlocks his door, opens it, and steps out. John can unlock his door either manually or electronically by pressing a button. Because John is always alone in the car, for him the goal of pressing that button is to unlock *his* door. Nevertheless, pressing the button causes the four doors of the car to be unlocked as a side effect. One day John offers his colleague, Pat, a ride to work. When they arrive, John starts to perform the usual three-step sequence to get out of the car. John's door happens to be unlocked because, for some reason, he has manually unlocked it on their way. That is, the *goal* of unlocking the door is already achieved; John can directly open his door and step out of the car. However, he still needs to perform the act of electronically unlocking *his* door in order to unlock Pat's door, a side effect of the act.

The above example seems to be a suitable one for an act whose goal holds and still needs to be performed for one of its effects. The problem though is that this assumes that the relation between an act and its goal is accidental; the goal is just an arbitrary effect of the act that happens to get some special status due to pragmatic reasons (John's life-style in the example). Such a position confuses two concepts: an act category, *ec*, and the PML procedure, $\mathcal{A}(ec)$, associated to it. The act category is a particular conceptualization of the motor program ($\mathcal{A}(ec)$). The same PML routine may be associated with different act categories. The act categories themselves have their goals as inherent parts of their characterization. In the above example, one would have two act categories: $\mathsf{Unlock\text{-}my\text{-}door}$ and, for example, $\mathsf{Unlock\text{-}passenger\text{-}door}$ (or even more specifically, $\mathsf{Unlock\text{-}Pat's\text{-}door}$). The goal of the first is my (i.e., John's) door's being unlocked and that of the second is the passenger's door's being unlocked. Even though they may both be associated with the same PML procedure (pushing the same button), the two act categories are distinct mental entities, and part of that distinction is their distinct goals. More generally, for every effect of a PML procedure, one may have a distinct act category with that effect as its goal.

This indeed suggests that an act should not be performed if its goal already holds. The simple assumption that the goal of a sequence is that of its last act is therefore not quite satisfactory.[18] Intuitively, the goal of a sequence of acts is to achieve the goals of its individual acts in order. However, this is a more complicated way of saying "to correctly complete the act". The problem is that there is no one single state that, when achieved, would signal the completion of a control act; the act being complete is something that the agent concludes based on its conscious monitoring of its progression. Therefore, like atelic acts, the goal of a type-token association, $ec[e]$, where ec is a control act category, is again the state $\mathsf{Complete(e)}$, such a state is asserted to be achieved not merely based on sensory input or interruptions but on the very process of executing the control act. Of course, this is not an adequate description of what exactly establishes $\mathsf{Complete(e)}$ for a control act. This point will be discussed in detail in Section 9.8.

The following is a semi-formal definition of goals.

Definition 9.1 *Let Γ be a partial function from type-token associations to states. More precisely, $\Gamma: \mathcal{EC} \times \mathcal{E} \rightharpoonup \textsc{Temp} \cup \textsc{Perm}$, such that $\Gamma(ec[e])$ is the goal of $ec[e]$.*[19] *$\Gamma(ec[e])$ is defined as follows.*

1. *If $\beta \vdash \mathsf{telic}(ec, s)$, then $\Gamma(ec[e]) = s$,*

2. *if ec is an atelic or a control act category, then $\Gamma(ec[e]) = \mathsf{Complete(e)}$, and*

3. *If ec is a mental act category, then $\Gamma(ec[e])$ is undefined.*

Note the following:

1. Γ is a function. That is, I am assuming that each act has a *unique* goal (if it has any).

2. Cassie has explicit beliefs about the goals of telic acts. These are represented using propositions of the form $\mathsf{telic}(ec, s)$ (see Section 4.6.3), where s is the goal associated with type-token associations involving the type ec. Ideally, however, knowledge about these associations should be structural rather than assertional (Woods, 1975) since many linguists (in reference to telic situations) argue that "reference to the goal is an essential part of the description of the situation" (Declerck, 1989, p. 277; also see Depraetere, 1995).

[18]Things are even more complicated with other control acts like snif.

[19]*Note that the goal must be a state with the $+$onset feature (see Section 3.5.4)—hence, the range of Γ.*

9.7 Cascades

Given a sequence of acts, Cassie should start performing the first and form the belief that, when its goal is achieved, she can perform the rest of the sequence. That is, achieving the goal of the first act will have a *cascade* effect resulting in the rest of the sequence being performed in a similar fashion. Such an effect requires some way to *transform* the belief of a goal having been achieved into a performance of an act (in this case a sequence). (Kumar, 1994) formalizes the required notion of transformers (also see (Kumar and Shapiro, 1994a) and (Kumar and Shapiro, 1994b)). The following is a slight modification of a *proposition-act* transformer suggested by Kumar and informally discussed in Chapter 6 (also see (Shapiro and the SNePS Implementation Group, 1999)).

- $\mathsf{whendo} : \mathcal{S} \times \mathcal{EC} \longrightarrow$ ETERNAL. For any $s \in \mathcal{S}$ and $ec \in \mathcal{EC}$, if forward inference causes both $\mathsf{whendo}(s, ec)$ and $\mathsf{Holds}(s, \text{*}\mathsf{NOW})$ to be asserted, then ec is performed and $\mathsf{whendo}(s, ec)$ is disbelieved.[20]

We are now ready to give a precise account of the performance of sequences. This is achieved by the following control act.

- $\mathsf{cascade} : \mathcal{EC}^n \longrightarrow \mathcal{E}$. For some instance, e, of ec_1, $\mathcal{A}(\mathsf{cascade}(ec_1, ec_2, \ldots, ec_n))$ executes as follows:

  1. If $\Gamma(ec_1[e])$ is not defined, then $\mathcal{A}(\mathsf{snsequence}(ec_1, \mathsf{cascade}(ec_2, \ldots, ec_n)))$ is executed.
  2. If $\mathtt{Backward}(\Gamma(ec_1[e]))$, then $\mathcal{A}(\mathsf{cascade}(ec_2, \ldots, ec_n))$ is executed.
  3. Otherwise,

  $$\mathcal{A}(\mathsf{snsequnce}(\mathsf{believe}(\mathsf{whendo}(\Gamma(ec_1[e])), \mathsf{cascade}(ec_2, \ldots, ec_n)), ec_1).$$

  is executed.

As shown above, what Cassie will exactly do when performing a cascade depends on two factors: whether the first act has a goal and whether its goal is already achieved. If the first act does not have a goal, then the cascade reduces to the sequential initiation of this act directly followed by

[20]It has been brought to my attention, by Sam Steel and Antony Galton (personal communication), that there might be categorial problems with considering whendo constructs to be propositions. The exact ontological status of these constructs is a topic for future research.

the execution of the rest of the cascade. According to Definition 9.1, this will happen in case the first act is a mental act. It should be noted, however, that this is only the way *I* am proposing to use cascade; other users may have other theories of acts and their goals, and this, by no means, would affect the way cascaded acts are performed. In other words, the procedural semantics of cascade is not dependent on Definition 9.1. If the goal of the first act is defined, the second clause makes sure that the first act is not performed if its goal is already achieved. Note that I am assuming that the first act in the cascade is tokenized. This is important to give a precise semantics for the execution of cascades. Section 9.8 provides a more precise account of how, and when, acts are tokenized. If the goal of the first act does not already hold, then Cassie starts performing the first act only after forming the belief that, when she is done, she will perform the rest of the cascade. Note that this is guaranteed to work, since coming to believe that the goal of an act has been achieved is done with forward inference that would *activate* the believed whendo. Breaking up the sequence in this way, allows for simple solutions to deal with errors and interrupts. For more on this, see Chapter 10.

9.8 Semantics of Act Execution and Completion

As pointed out in Sections 9.6.3, the goal of a type-token association, $ec[e]$, where ec is a control act, is the state $\mathsf{Complete}(e)$. Cassie comes to believe in such a state by the very process of performing the control act. To precisely explain what this last sentence means, I shall give formal operational semantics for the different SNeRE constructs discussed. Recall from Section 1.5.1 that SNeRE makes use of a network activation mechanism to control the performance of acts. The exact details of how this works may be found in (Kumar, 1994). However, most of the details are irrelevant; the only relevant notion is that of activation. Basically, an act is activated whenever Cassie decides to perform it. Constructs like whendo and wheneverdo are activated when their conditions are asserted with forward inference. The operational semantics of the different SNeRE constructs will be given by inductively defining a single operation: *Activate*. Different definitions of the operation correspond to different types of constructs it operates upon. One can think of these defintions as *reduction axioms* used in, for example, (Mitchell, 1996) to provide operational semantics for programming languages. Essentially, such axioms are high-level renderings of λ-reductions. The syntax of a single definition for *Activate* will follow the following general schema.

$$Activate(\langle\text{pattern}\rangle) \rightarrow \langle\text{body}\rangle.$$

The arrow represents a single reduction: the left-hand side reduces in one step to the right-hand side. The body of the definition would look pretty much like the body of an algorithm employing control structures such as sequencing (represented by ";") and selection (represented by the familiar *if* ... *then* ... *else* ...). These structures may be thought of as macros whose meanings can be formally fleshed out using either λ-abstractions or some abstract machine model in the style of VDL (Vienna Description Language), for example (see (Pagan, 1981)). I shall, however, rely on the reader's intuitive understanding of algorithmic notation.

The "pattern" appearing on the left-hand side is one of five *basic* patterns of SNeRE constructs:

1. $ec[e]$; an act type-token association as presented in Section 9.6.2

2. $ec[]$; a SNePS term representing an act category that is not (yet) associated with a particular token.

3. ec; a SNePS term representing an act category which may, or may not, be associated with a token. That is, the pattern ec subsumes the above two patterns.

4. $\mathsf{whendo}(s, ec)$, where s represents a state and ec is a pattern as in (3) above.

5. $\mathsf{wheneverdo}(s, ec)$, where s represents a proposition and ec is a pattern as in (3) above.

Any of the above basic patterns (except the last two) may be annotated with a superscript that further restricts the types of constructs that it can match. Triples of the form $\langle\pm\mathbf{R}, \pm\mathbf{L}, \pm\mathbf{P}\rangle$ (with the appropriate setting of $+$ and $-$) are used to refer to RLP act categories. For example, the pattern $ec^{\langle-\mathbf{R},+\mathbf{L},\pm\mathbf{P}\rangle}$ would match any $\overleftarrow{\mathsf{atelic}}$ act category. In addition, I use the superscripts "$+M$", "$-M$" and "Ctrl" to designate categories of mental, non-mental, and control acts, respectively.

Figure 9.3 shows the definitions of *Activate* for mental acts and whendo and wheneverdo constructs. The wheneverdo construct is similar to whendo (see Section 9.7) but the rule does not get discarded when the act is activated; it is used for general acting rules, rather than occasional ones (an example of its use is given in Section 11.1). Activating a mental act reduces, in one step, to initiating it. Initiating a mental act is the simple (theoretically instantaneous and error-free) addition (believe) or retraction (disbelieve) of a proposition from the knowledge base (note that initiation

1. $Activate(ec^{+M}) \rightarrow Initiate(\mathcal{A}(ec))$.
2. $Activate(\mathsf{whendo}(s, ec)) \rightarrow Activate(ec); Activate(\mathsf{disbelieve}(\mathsf{whendo}(s, ec)))$.
3. $Activate(\mathsf{wheneverdo}(s, ec)) \rightarrow Activate(ec)$.

Figure 9.3: Reduction axioms for the activation of mental acts, whendo, and wheneverdo.

4. $Activate(ec^{\langle -\mathbf{R}, +\mathbf{L}, \pm\mathbf{P}\rangle}[]) \rightarrow Activate(\mathsf{believe}(\mathsf{Cat}(e, ec)));$
$Activate(\mathsf{believe}(\mathsf{whendo}(s, \mathsf{believe}(\mathsf{Complete}(e)))));$
$Activate(ec[e])$.
Where e is a newly introduced term in $\mathcal{E}$ and $\beta \vdash \overleftarrow{\mathsf{atelic}}(ec, s)$.
5. $Activate(ec^{\langle -\mathbf{R}, -\mathbf{L}, \pm\mathbf{P}\rangle}[]) \rightarrow Activate(\mathsf{believe}(\mathsf{Cat}(e, ec))); Activate(ec[e])$.
Where e is a newly introduced term in $\mathcal{E}$.
6. $Activate(ec^{\langle +\mathbf{R}, \pm\mathbf{L}, \pm\mathbf{P}\rangle}[]) \rightarrow Activate(\mathsf{believe}(\mathsf{Cat}(e, ec))); Activate(ec[e])$.
Where e is a newly introduced term in $\mathcal{E}$.
7. $Activate(ec^{\langle \pm\mathbf{R}, \pm\mathbf{L}, +\mathbf{P}\rangle}[e]) \rightarrow Initiate(\mathcal{A}(ec[e]))$.
8. $Activate(ec^{\langle -\mathbf{R}, -\mathbf{L}, -\mathbf{P}\rangle}[e]) \rightarrow Activate(\mathsf{DoOne}(\{ec_i : \beta \vdash \mathsf{ActPlan}(ec, ec_i)\})[])$.
9. $Activate(ec^{\langle -\mathbf{R}, +\mathbf{L}, -\mathbf{P}\rangle}[e]) \rightarrow Activate(\mathsf{DoOne}(\{ec_i : \beta \vdash \mathsf{ActPlan}(ec, ec_i)\})[e])$.
10. $Activate(ec^{\langle +\mathbf{R}, \pm\mathbf{L}, -\mathbf{P}\rangle}[e]) \rightarrow Activate(\mathsf{DoOne}(\{ec_i : \beta \vdash \mathsf{ActPlan}(ec, ec_i)\})[e])$.

Figure 9.4: Reduction axioms for the activation of RLP acts.

may only be applied to primitive acts). As stated in Section 9.6.1, addition and retraction are done in such a way to maintain the consistency of the knowledge base (Martins and Shapiro, 1988). It should be noted that, once a mental act is activated, it is assumed that its effects apply immediately. The only thing worth pointing out regarding definitions (2) and (3) is that a whendo proposition is immediately retracted once it is activated, whereas a wheneverdo persists.

Figure 9.4 outlines the definitions of *Activate* for RLP acts. Note that these exclude mental and control acts to which the notion of telicity does not apply. Before discussing definition (4), let us first consider (5) and (6). The patterns on the left-hand sides of these definitions match any non-$\overleftarrow{\mathsf{atelic}}$ act category that is not associated with a particular token. This is typically the situation when an act is first activated. It is at this time that a token for the act is conceived—when Cassie starts entertaining the performance of the act. Thus, the right-hand side of the definition simply introduces a new token and then activates the resulting type-token association. Definition (4) is similar, but it only applies to $\overleftarrow{\mathsf{atelic}}$ acts. The difference is that for $\overleftarrow{\mathsf{atelic}}$ acts, a particular performance may reach completion by either being interrupted or by achieving the state, s, beyond which the act cannot continue. The second activation on the right-hand side of (4) guarantees that, once s is achieved, the agent would believe that the act is complete (see the discusssion in Section 9.6.3). Definition (7) states that activating a primitive RLP act reduces to initiating it. Definitions (8), (9), and (10) state that activating a composite act reduces to activating one of the plans that it decomposes to, where $\mathsf{ActPlan}(ec_1, ec_2)$ means that performing ec_2 is a plan for performing the composite ec_1 (the activation of DoOne is defined below).

Why is more than one definition needed for composite acts? As (8), (9), and (10) show, we are treating telic composite acts differently from atelic composite acts. In particular, for atelic acts, the token associated with the composite act is *passed down* to its DoOne decomposition, whereas, for telic acts, it is not. The basic idea behind passing the token to the decompostion of a composite act is for Cassie to come to believe that the act is complete when its plan is (this is precisely illustrated below). Now, this is reasonable if the composite act is atelic, and hence (9) and (10). For telic acts, however, there are two reasons why passing the token is not a good idea. First, we do not need it; a telic act is complete when some token-independent state starts to hold. Second, and more important, it is outright wrong. Consider a composite $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ act. Such an act would reduce to some plan. However, the performance of the plan's being complete is no reason for Cassie to believe that the act is. For example, although calling the secretary may be complete, the act of bringing her is not necessarily complete at the same time. To avoid such complications, we choose not to pass the token to the decomposition of any telic act.

Figure 9.5 represents a summary of the main outcome of the chapter. It outlines definitions of $\mathsf{cascade}$ activations and is, thus, a precise statement of how sequential acting works according to our theory. Definition (11) simply states that activating an empty $\mathsf{cascade}$ with an associated token, e, reduces (in two steps) to believing that e is complete. Definitions (12) through (16) outline how activation works for different cases of the non-tokenized head of a cascade. (12), (13) and (14) correspond to definitions (4), (5), and (6), respectively, of Figure 9.4. Their main purpose is to introduce a token for the head of the cascade. The reason this is needed is that performing a cascade essentially involves reasoning about the goal of its first act. As argued in Section 9.6.2, goals are only defined for type-token associations of acts, not for non-tokenized act categories. In addition, Cassie's performing a cascade requires her to reason about the particular future performance of

11. $\textit{Activate}(\mathsf{cascade}()[e]) \rightarrow \textit{Activate}(\mathsf{believe}(\mathsf{Complete}(e)))$.

12. $\textit{Activate}(\mathsf{cascade}(ec_1^{\langle -\mathbf{R}, +\mathbf{L}, \pm\mathbf{P}\rangle}[], ec_2, \ldots, ec_n)[e_1]) \rightarrow$

if $\texttt{Backward}(s)$ *then* $\textit{Activate}(\mathsf{cascade}(ec_2, \ldots, ec_n)[e_1])$

else

$\textit{Activate}(\mathsf{believe}(\mathsf{Cat}(e_2, ec_1)))$;

$\textit{Activate}(\mathsf{believe}(\mathsf{whendo}(s, \mathsf{believe}(\mathsf{Complete}(e_2)))))$;

$\textit{Activate}(\mathsf{cascade}(ec_1[e_2], ec_2, \ldots, ec_n)[e_1])$.

Where e_2 is a newly introduced $\mathcal{E}$-term and $\beta \vdash \mathsf{telic}(ec, s)$.

13. $\textit{Activate}(\mathsf{cascade}(ec_1^{\langle -\mathbf{R}, -\mathbf{L}, \pm\mathbf{P}\rangle}[], ec_2, \ldots, ec_n)[e_1]) \rightarrow$

$\textit{Activate}(\mathsf{believe}(\mathsf{Cat}(e_2, ec_1)))$; $\textit{Activate}(\mathsf{cascade}(ec_1[e_2], ec_2, \ldots, ec_n)[e_1])$.

Where e_2 is a newly introduced $\mathcal{E}$-term.

14. $\textit{Activate}(\mathsf{cascade}(ec_1^{\langle +\mathbf{R}, \pm\mathbf{L}, \pm\mathbf{P}\rangle}[], ec_2, \ldots, ec_n)[e_1]) \rightarrow$

$\textit{Activate}(\mathsf{believe}(\mathsf{Cat}(e_2, ec_1)))$; $\textit{Activate}(\mathsf{cascade}(ec_1[e_2], ec_2, \ldots, ec_n)[e_1])$.

Where e_2 is a newly introduced $\mathcal{E}$-term.

15. $\textit{Activate}(\mathsf{cascade}(ec_1^{\mathrm{Ctrl}}[], ec_2, \ldots, ec_n)[e_1]) \rightarrow$

$\textit{Activate}(\mathsf{believe}(\mathsf{Cat}(e_2, ec_1)))$; $\textit{Activate}(\mathsf{cascade}(ec_1[e_2], ec_2, \ldots, ec_n)[e_1])$.

Where e_2 is a newly introduced $\mathcal{E}$-term.

16. $\textit{Activate}(\mathsf{cascade}(ec_1^{+M}, ec_2, \ldots, ec_n)[e]) \rightarrow \textit{Activate}(ec_1)$; $\textit{Activate}(\mathsf{cascade}(ec_2, \ldots, ec_n)[e])$

. **17.** $\textit{Activate}(\mathsf{cascade}(ec_1^{-M}, ec_2, \ldots, ec_n)[e]) \rightarrow$

if $\texttt{Backward}(\Gamma(ec_1))$ *then* $\textit{Activate}(\mathsf{cascade}(ec_2, \ldots, ec_n)[e])$

else $\textit{Activate}(\mathsf{believe}(\mathsf{whendo}(\Gamma(ec_1), \mathsf{cascade}(ec_2, \ldots, ec_n)[e])))$; $\textit{Activate}(ec_1)$

Figure 9.5: Reduction axioms for the activation of cascade.

its first act (before actually starting to perform it) and thus conceiving of it (see Section 4.6.2). Definition (15) is similar to (13) and (14); it introduces a new token for the control act that heads the cascade. Since mental acts are assumed to be complete once activated, (16) activates the rest of the cascade right after activating the mental head-act. Finally, (17) represents the general case for any type of tokenized non-mental head-act. In this case, the function Γ is defined for the head of the cascade and the definition follows the semi-formal outline presented in Section 9.7. In general, note how the token associated with the activated cascade is passed from the activation of a non-empty cascade to the activation of its tail. When activations recurse deep enough to reach an empty cascade, definition (11) applies and Cassie comes to believe that the whole cascade is complete.

To further illustrate how Cassie comes to believe that a control act is complete (the case of cascade in Figure 9.5 is already an example), Figure 9.6 shows definitions of *Activate* for the SNeRE control acts mentioned in Section 9.6.1. Just like with RLP acts, (18) introduces a new token for a newly-activated control act (this is where the cascade tokens in Figure 9.5 come from). Definitions (19) and (20) for snsequence are self-explanatory. The main reason they are included here is to make clear the difference between cascade and snsequence. In particular, an snsequence is complete once all of its elements are activated in order. Note that, as per (20), snsequence does not *wait* for the completion of any of its acts, which is the intended semantics. The tail of a non-empty cascade, on the other hand, *blocks* following the activation of the head. It is reactivated (indirectly through the activation of whendo) only when the head-act is complete.

Definition (21) for DoOne provides another example of control act activation and is also needed since it is used on the right-hand sides of definitions (8), (9), and (10) of Figure 9.4. Informally, activating a DoOne reduces to activating an arbitrary act from its set-of-acts argument. There are two points to note. First, a cascade is wrapped around the act chosen for activation. This is done to ensure that the whole DoOne act is believed to be complete when and only when the chosen act is. Second, the definition does not apply when the argument of DoOne is the empty set. This means that the act would simply *fail* if this were the case. Thus, given Defintions (8), (9) and (10), the activation of a composite act fails if no plan could be deduced for it.

The definition for snif should be self-explanatory. Again note that a cascade is wrapped around the chosen act to ensure that the whole snif act is believed to be complete when and only when the chosen act is. Definition (23) for sniterate follows a similar pattern. The repetitive behavior of

18. $\mathit{Activate}(ec^{\mathrm{Ctrl}}[]) \rightarrow \mathit{Activate}(\mathsf{believe}(\mathsf{Cat}(e, ec))); \mathit{Activate}(ec[e])$.
Where e is a newly introduced $\mathcal{E}$-term.
19. $\mathit{Activate}(\mathsf{snsequence}()[e]) \rightarrow \mathit{Activate}(\mathsf{believe}(\mathsf{Complete}(e)))$.
20. $\mathit{Activate}(\mathsf{snsequence}(ec_1, ec_2, \ldots, ec_n)[e]) \rightarrow \mathit{Activate}(ec_1); \mathit{Activate}(\mathsf{snsequence}(ec_2, \ldots, ec_n)[e])$.
21. $\mathit{Activate}(\mathsf{DoOne}(\{ec_i\}_{i=1}^{n})[e]) \rightarrow \mathit{Activate}(\mathsf{cascade}(ec_i)[e])$
for some arbitrary $1 \leq i \leq n$.
22. $\mathit{Activate}(\mathsf{snif}(\{\langle s_i, ec_i\rangle\}_{i=1}^{n}, ec_{n+1})[e]) \rightarrow$
if $\exists j, 1 \leq j \leq n$, such that $\mathtt{Backward}(s_j)$ *then* $\mathit{Activate}(\mathsf{cascade}(ec_j)[e])$
else $\mathit{Activate}(\mathsf{cascade}(ec_{n+1})[e])$.
23. $\mathit{Activate}(\mathsf{sniterate}(\{\langle s_i, ec_i\rangle\}_{i=1}^{n}, ec_{n+1})[e]) \rightarrow$
if $\exists j, 1 \leq j \leq n$, such that $\mathtt{Backward}(s_j)$ *then*
$\mathit{Activate}(\mathsf{cascade}(ec_j, \mathsf{sniterate}(\{\langle s_i, ec_i\rangle\}_{i=1}^{n}, ec_{n+1})[e]))$
else $\mathit{Activate}(\mathsf{cascade}(ec_{n+1})[e])$.

Figure 9.6: Reduction axioms for the activation of a collection of control acts.

$\mathsf{sniterate}$ is effected by cascading the chosen act with the entire $\mathsf{sniterate}$.

Now, given the above system of axioms, it should be obvious that the *normal form* of any activation (i.e., the form beyond which it cannot be reduced) is a sequence of initiations. Note that this conforms with our discussion in Section 9.1, where it was stated that the only thing that an agent is guaranteed to do is to *start* performing some act. To illustrate how the reduction axioms work, I shall work through a simple example. Consider the following example which is a formalization of instruction (9.3) from Section 9.1. Here, "Gloria" is the name of the secretary.

$$\mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope}, \mathsf{stamp}), \mathsf{bring}(\mathsf{Gloria}), \mathsf{give}(\mathsf{envelope}, \mathsf{Gloria}))$$

Cassie's deciding to perform this act activates it. Since it is a control act, it matches the pattern in Definition (18), introducing a token for the new performance:

$$\begin{aligned}&\mathit{Activate}(\mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope}, \mathsf{stamp}), \mathsf{bring}(\mathsf{Gloria}), \mathsf{give}(\mathsf{envelope}, \mathsf{Gloria}))) \rightarrow\\ &\quad \mathit{Activate}(\mathsf{believe}(\mathsf{Cat}(e_1, \mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope}, \mathsf{stamp}), \mathsf{bring}(\mathsf{Gloria}), \mathsf{give}(\mathsf{envelope}, \mathsf{Gloria}))))); \\ &\quad \mathit{Activate}(\mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope}, \mathsf{stamp}), \mathsf{bring}(\mathsf{Gloria}), \mathsf{give}(\mathsf{envelope}, \mathsf{Gloria}))[e_1]).\end{aligned}$$

By Definition (1), the first element in the resulting sequence reduces to an initiation:

$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{Cat}(e_1, \mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope}, \mathsf{stamp}), \mathsf{bring}(\mathsf{Gloria}), \mathsf{give}(\mathsf{envelope}, \mathsf{Gloria})))))).$$

The second element is an activation of a tokenized cascade whose head is a telic act. This matches Definition (13):

$$\mathit{Activate}(\mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp}),\mathsf{bring}(\mathsf{Gloria}),\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))[e_1]) \rightarrow$$
$$\mathit{Activate}(\mathsf{believe}(\mathsf{Cat}(e_2,\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp}))));$$
$$\mathit{Activate}(\mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp})[\mathsf{e}_2],\mathsf{bring}(\mathsf{Gloria}),\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))[e_1]).$$

This results, again, in a two-element sequence. The first reduces to an initiation as per Definition (1), and the second matches the left-hand side of Definition (17). Assuming that there is no stamp on the envelope (i.e., $\Gamma(\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp})$ is not deducible), then

$$\mathit{Activate}(\mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp})[\mathsf{e}_2],\mathsf{bring}(\mathsf{Gloria}),\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))[e_1]) \rightarrow$$
$$\mathit{Activate}(\mathsf{believe}(\mathsf{whendo}(\mathsf{on}(\mathsf{envelop},\mathsf{stamp}),\mathsf{cascade}(\mathsf{bring}(\mathsf{Gloria}),\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))[e_1])));$$
$$\mathit{Activate}(\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp})[e_2]).$$

Through two further reductions using Definitions (1) and (7), the initially-activated cascade is reduced to the following normal form:

$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{Cat}(e_1,\mathsf{cascade}(\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp}),\mathsf{bring}(\mathsf{Gloria}),\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{Cat}(e_2,\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp})))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{whendo}(\mathsf{on}(\mathsf{envelop},\mathsf{stamp}),\mathsf{cascade}(\mathsf{bring}(\mathsf{Gloria}),\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))[e_1]))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{stickon}(\mathsf{envelope},\mathsf{stamp})[e_2])).$$

It should be noted that, in practice, SNeRE does not fully reduce an activation to some normal form before starting the actual initiations. Rather, once an *Initiate* form is reached, the corresponding actual initiation immediately takes place.

Now, suppose that Cassie succeeds in sticking the stamp on the envelope. Perceiving this state results in activating the believed whendo (third element in the above normal form), which, following Definition (2), activates the pending cascade. Without getting into much detail, the activated cascade reduces to yet another normal form ($\rightarrow^*$ denotes reduction in zero or more steps):

$$\mathit{Activate}(\mathsf{cascade}(\mathsf{bring}(\mathsf{Gloria}),\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))[e_1]) \rightarrow^*$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{Cat}(e_3,\mathsf{bring}(\mathsf{Gloria})))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{whendo}(\mathsf{here}(\mathsf{Gloria}),\mathsf{cascade}(\mathsf{give}(\mathsf{envelope},\mathsf{Gloria}))[e_1]))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{Cat}(e_4,\mathsf{DoOne}(\{\mathsf{say}(\text{“Come here”},\mathsf{Gloria})\})))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{Cat}(e_5,\mathsf{say}(\text{“Come here”},\mathsf{Gloria})))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{whendo}(\mathsf{said}(\text{“Come here”},\mathsf{Gloria}),\mathsf{cascade}()[e_4]))));$$
$$\mathit{Initiate}(\mathcal{A}(\mathsf{say}(\text{“Come here”},\mathsf{Gloria})[e_5])).$$

The above normal form is a sequence of six initiations. The first two should be obvious enough. The next four are what the activation of the composite act $\mathsf{bring(Gloria)}$ reduces to. There are two points to note here. First, note that the token associated with $\mathsf{bring(Gloria)}$ (e_3) is not the same one associated with the DoOne to which it decomposes (e_4). The reason is that $\mathsf{bring(Gloria)}$ is a $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$ act and is not complete by simply performing the plan of calling Gloria (see Definition (8)). Second, I assume said(“Come here”, x) to be the goal of the primitive act say(“Come here”, x). This might seem rather dubious. However, for some primitive telic acts like saying something, it is not exactly clear if there is an easily-indentifiable state that signals their completion.[21] It could be argued that for such acts, the goal is simply $\mathsf{Complete}(e)$, for some token e. The difference between these and atelic acts is that $\mathsf{Complete}(e)$ should not be asserted if the act is interrupted, but only when the lower PML process runs normally to completion. For the sake of the above example, however, I choose not to complicate matters by carving out a further partition in the ontology of acts.

Now suppose that Gloria arrives. This activates the single-element pending cascade, resulting in the following reduction to normal form:

$$
\begin{aligned}
&\mathit{Activate}(\mathsf{give(envelope,Gloria)}\,[e_1]) \rightarrow^* \\
&\quad \mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{Cat}(e_6, \mathsf{give(envelope,Gloria)})))); \\
&\quad \mathit{Initiate}(\mathcal{A}(\mathsf{believe}(\mathsf{whendo}(\mathsf{has(Gloria,envelope)}, \mathsf{cascade()}\,[e_1])))); \\
&\quad \mathit{Initiate}(\mathcal{A}(\mathsf{give(envelope,Gloria)}\,[e_6])).
\end{aligned}
$$

Once Gloria has the envelope, the pending empty cascade is activated and, following Definition (11), Cassie comes to believe that the original act that she set out to perform is finally complete.

9.9 Conclusions

Performing sequences of acts is not as simple as it may seem. An embodied agent that properly executes a sequence of acts reasons, making use of its awareness of the environment and its own body, about when to move to the next step in the sequence. Such an apparently simple concept does not seem to have been explicitly addressed in general, precise, and abstract enough terms.

[21] (McDermott, 1982, p. 109) discusses a similar class of acts that “are done for their own sake”.

When an agent is acting, it has a certain conceptualization of what it is doing. Such conceptualizations vary along the dimension of telicity. Different types of telicity correspond to different criteria regarding what signals successful completion of an act. Successful completion of an act is signaled by some state —the goal— starting to hold. Telic acts have a built-in goal that is essentially part of their structure. Atelic acts, on the other hand, do not have such a built-in goal, and complete by being consciously terminated by the agent at some arbitrary point. Control acts, which are neither telic nor atelic in the classical senses of the terms, are consciously performed. The agent is aware of their progression and the goal of such acts is simply the state of their being complete, something that the agent comes to believe by virtue of its very performance of the act.

A mechanism was introduced for performing sequences of acts, *cascades*, that gives the agent conscious control of their execution. To cascade a sequence of acts, the agent starts performing the first and forms the belief that when its goal is achieved it shall (recursively) perform the rest of the cascade. Although the presentation of cascades has been primarily within the confines of the SNePS system, the proposal is a theoretical one and, thus, system-independent. The cascading mechanism itself is conceptually simple, but it is founded on a thorough investigation of the non-trivial notion of act completion.

Chapter 10

Towards a Theory of Interrupts

In the real world, any of the actions of a cognitive agent may fail to achieve its goals. A theory of acting agents should therefore be constructed with failure deeply in mind. At each step, the agent should be aware of the outcome of its previous actions and should behave appropriately to recover from errors. In addition, such behavior should be the result of the agent's reasoning, not of hardwired reactive mechanisms. In particular, the agent should be able to reason and discuss its actions and failures.

This chapter sketches the beginnings of a theory of interrupts and interrupt handling. The work reported here is based on (Ismail and Shapiro, 2000a) and is still very cursory at this stage.

10.1 Related Work

The action literature within symbolic artificial intelligence contains various, though essentially similar, proposals to deal with the problem of interrupt handling. The basic recurring theme is that interrupt handling involves the definition of priorities among acts (or goals). Reactive planners (Georgeff and Lansky, 1987; Shanahan, 1998, for instance) (see Chapter 2) typically interleave planning and action execution; once the planner generates a primitive act, it starts to execute while planning is still going on. Interrupts in this setting may be handled by simply generating the appropriate reaction. This is feasible since the system never commits to a particular sequence of actions, only one act at a time.

Other systems, where plans are specified (or generated off-line) in some action specification language, need to provide appropriate means for handling interrupts. (Davis, 1992) presents precise semantics for an action language equipped with control structures for handling interrupts. For example, the expression "interrupt_for($T1, T2$)" corresponds to the execution of task $T2$, interrupting it when necessary for $T1$. Essentially, this means that $T1$ has higher priority over $T2$. Given the formal semantics (Davis, 1992, p. 42), it is not clear how the system can represent priorities that may change over time. Such an issue, I believe, is crucial and, as it turns out, is overlooked by many of the existing systems.

Within the GOLOG family (Levesque et al., 1997), interrupts are handled in CONGOLOG using special control structures for priorities and reactions (De Giacomo et al., 2000). In a CONGOLOG program, the expression "$(\sigma_1\rangle\rangle\sigma_2)$" denotes the concurrent execution of the actions σ_1 and σ_2 with σ_1 having higher priority over σ_2. Note that this is essentially Davis's "interrupt_for" control structure. Further, "$(\sigma_1\rangle\rangle\sigma_2)$" denotes an act, a step in a CONGOLOG program that the agent should execute in a certain manner as indicated by the semantics. Thus, once the agent starts performing $(\sigma_1\rangle\rangle\sigma_2)$, it is not obvious how it may decide to "change its mind" regarding the priorities of σ_1 and σ_2, and, for example, interrupt σ_1 to perform σ_2.

Interrupt handling obviously involves the notion of priorities. The problem with approaches such as the above (where priorities are represented as actions in plans) is that they do not provide enough flexibility for the agent to reason about what to do next. In the system presented below, priority information is represented as context-dependent domain knowledge that may be communicated on-line (for example, in natural language) while the agent is acting. Interrupt handling is not represented by means of explicit control structures, but is built into the acting executive. Whenever the agent is about to act, it reasons about what it is about to do, and what it is currently doing, to decide what to do next.

10.2 The Concept of Interruption

To develop a theory of recovering from errors, one needs to be precise about what sort of thing an error is. An error, as far as this work is concerned, is a special kind of an interrupt, where an interrupt is an event that causes the agent to stop what it is doing and handle an unexpected situation.

An error is special in that the unexpected situation is the failure of one of the agent's acts. A general theory for interrupt handling would subsume one for error recovery. I will, therefore, discuss general interrupt handling in this chapter, error recovery being a by-product of the theory.

An interrupt is an event that causes the agent to change its intentions and/or actions. It involves three main components: an event $\mathfrak{I}$ (typically the onset or cessation of some state), a reaction $\mathfrak{R}(\mathfrak{I})$, and a *non-empty* set Π of on-going processes. For there to be an interrupt, the three components have to be present. For example, Cassie may be carrying out a number of concurrent processes (Π) when her battery goes low ($\mathfrak{I}$), requiring her to move to the recharging station ($\mathfrak{R}(\mathfrak{I})$). There are a number of points to note:

1. The reaction is a function of the event $\mathfrak{I}$. Thus, there can be no situation in which an act, viewed as a reaction, is present without a corresponding event.

2. The event $\mathfrak{I}$ may be an instruction by another superior agent (possibly a human operator) to perform some action, which, in that case, would be $\mathfrak{R}(\mathfrak{I})$.

3. The set Π is not empty. If the agent is not doing anything, then whatever happens is not an interrupt, just an event that may require some appropriate reaction.[1]

4. A valid reaction is the act of stopping one of the processes in Π.

For the sake of simplicity, let us assume for the moment that Π is a singleton containing only one act, α. α could be either *primitive* or *composite*. When $\mathfrak{I}$ occurs (and is detected by Cassie), Cassie could be performing either a primitive or a composite act. Corresponding to these two situations, there are two types of interrupts: one that happens while performing a composite act and another that happens in the midst of executing a primitive act. To make things more concrete, we can identify these two categories of interrupts as follows:

1. Cassie is executing a composite act α which reduces to the execution of some sequence of acts $\langle\alpha_1 \ldots \alpha_n\rangle$. She has just finished performing α_i, and is about to perform α_{i+1}, when $\mathfrak{I}$ occurs. For example, Cassie may be performing the sequence of acts $\langle\mathsf{PickUp\text{-}Block}, \mathsf{Goto\text{-}Table}, \mathsf{Put\text{-}Down\text{-}Block}\rangle$ and have just picked up the block when she senses that her battery is low.

[1]Note that this subsumes errors. Nothing can qualify as an error if the agent is not doing anything.

2. Cassie is in the midst of executing a primitive act (which could be part of executing a composite act) when $\mathfrak{I}$ occurs. For example, while performing the $\langle$PickUp-Block, Goto-Table, Put-Down-Block$\rangle$ sequence, the battery goes low while Cassie is moving toward the table.

In the first case, Cassie needs to merely change (or, more precisely, *revise*) her intentions regarding what to do next (go to the table or recharge the battery). In the second case, she may need to *stop* what she is doing to handle the interrupt. In any situation, there are two main requirements on any interrupt handling mechanism:

1. Cassie should first perform the act with higher priority. If continuing to perform α is more important than performing $\mathfrak{R}(\mathfrak{I})$, then this is what she should do. Note that priorities are context-sensitive; they change according to the current situation. For instance, if Cassie is at the table, then putting the block down may have higher priority than recharging the battery. If, on the other hand, she is at the recharging station, then recharging the battery should have higher priority.

2. If Cassie chooses to stop α and perform $\mathfrak{R}(\mathfrak{I})$, she should resume α when $\mathfrak{R}(\mathfrak{I})$ is complete. On the other hand, if she chooses to continue performing α, Cassie should somehow *remember* to perform $\mathfrak{R}(\mathfrak{I})$ when she is done.

10.3 On Priorities

Appropriately handling interrupts requires the agent to reason about the priorities of its acts. Therefore, to solve the problem of interrupt handling (i.e., what exactly an agent should do in case of an interrupt), we need to consider the following more general problem: Given a set of acts, A, to be executed, if due to resource contention not all the acts in A may be concurrently executed, which elements of A should the agent execute first? And how can it, then, attend to the rest of the acts? The first step is for the agent to somehow determine which elements of A have top priorities. Let us then closely examine the notion of priorities.

Priorities define a partial order over the set of acts. Two acts that are not prioritized relative to each other are assumed to have the same priority. Specification of priorities among acts may be explicitly represented in β by means of assertions involving the function Prior.

- $\mathsf{Prior}: \mathcal{EC} \times \mathcal{EC} \longrightarrow$ TEMP, where $[\![\mathsf{Prior}(ec_1, ec_2)]\!]$ is the temporary state of $[\![ec_1]\!]$'s having higher priority, for Cassie, over $[\![ec_2]\!]$.

Note that, unless $[\![ec_1]\!]$ and $[\![ec_2]\!]$ are acts that Cassie can perform, the idea of prioritizing them would not make much sense. Also note the crucial fact that one act's having priority over another is a *temporary* state. Thus, priorities are dynamic, they change as time goes by and the overall situation changes.

The following axioms capture the partial-order nature of Prior.

- **APr1.** $\mathsf{Holds}(\mathsf{Prior}(ECv_1, ECv_2), Tv) \Rightarrow \mathsf{Holds}(\neg\mathsf{Prior}(ECv_2, ECv_1), Tv)$

- **APr2.** $[\mathsf{Holds}(\mathsf{Prior}(ECv_1, ECv_2), Tv) \wedge \mathsf{Holds}(\mathsf{Prior}(ECv_2, ECv_3), Tv)] \Rightarrow$
$$\mathsf{Holds}(\mathsf{Prior}(ECv_1, ECv_3), Tv).$$

- **APr3.** $\mathsf{Equiv}(ECv, ECv') \Rightarrow \forall Tv[\mathsf{Holds}(\neg\mathsf{Prior}(ECv, ECv'), Tv)]$.

The following meta-theoretical axiom characterizes the minimum knowledge Cassie should have of priorities among acts.

Axiom 10.1 *For every $ec_1, ec_2 \in \Psi(\mathcal{EC})$, if*

1. $ec_1 \neq ec_2$ *and*

2. $\mathrm{Mod}_{\mathrm{prop}}(\mathsf{Prog}(ec_1)) \cap \mathrm{Mod}_{\mathrm{prop}}(\mathsf{Prog}(ec_2)) \neq \{\}$,

then, for every $i \in \mathbb{N}$ ($i > 0$), at $[\![{}^*\mathsf{NOW}_i]\!]$,

1. $\beta \vdash \mathsf{Holds}(\mathsf{Prior}(ec_1, ec_2), {}^*\mathsf{NOW}_i)$ *or*

2. $\beta \vdash \mathsf{Holds}(\mathsf{Prior}(ec_2, ec_1), {}^*\mathsf{NOW}_i)$

That is, if the states of performing two distinct acts occupy the same proprioception modality, then, at any time, Cassie must have a belief about the relative priorities of these acts. The point is that, since the acts would share some resources, they cannot be concurrently executed, and knowledge of their priorities is needed for Cassie to decide which one to perform first. For example, suppose that

$$\text{Mod}_{\text{prop}}(\mathsf{Prog}(\mathsf{Goto\text{-}Table})) = \{\mathsf{VISION}_{\text{prop}}, \mathsf{LOCOMOTION}_{\text{prop}}\}$$

$$\text{Mod}_{\text{prop}}(\mathsf{Prog}(\mathsf{Recharge\text{-}Battery})) = \{\mathsf{VISION}_{\text{prop}}, \mathsf{LOCOMOTION}_{\text{prop}}\}$$

If Cassie is about to go to the table when she senses that her battery is low, then she must decide whether to go ahead and move to the table or change her plans and recharge the battery. Because both acts make use of the same resources, they cannot be performed simultaneously. Thus, Cassie's decision would have to be based on her knowledge of the priorities of these acts in the current situation.

To describe how the acting executive actually executes acts according to their priorities, some meta-theoretical definitions are needed. The first definition extends the notion of priorities to cover acts that are embedded within cascades. First, three simple pieces of notation that should make the exposition more tractable. In what follows, $\mathcal{B}$ stands for "Body", $\mathcal{F}$ for "First", and $\mathcal{R}$ for "Rest".

Definition 10.1 *For every act ec, if $ec = \mathsf{cascade}(ec_1, \ldots, ec_n)$ (where $\{ec_i\}_{i=1}^{n}$ is a set of acts), then $\mathcal{B}(ec) = \langle ec_1, \ldots, ec_n \rangle$, $\mathcal{F}(ec) = ec_1$, and $\mathcal{R}(ec) = \langle ec_2, \ldots, ec_n \rangle$.*

Now, consider the following deeper notion of priority.

Definition 10.2 *For every two distinct acts $ec_1, ec_2 \in \Psi(\mathcal{EC})$ and $i \in \mathbb{N}$ ($i > 0$), at $[\![*\mathsf{NOW}]\!]$, $ec_1 >_{pr} ec_2$ (read, ec_1 has priority over ec_2) iff*

1. $\beta \vdash \mathsf{Holds}(\mathsf{Prior}(ec_1, ec_2), *\mathsf{NOW}_i)$,
2. $\beta \not\vdash \mathsf{Holds}(\mathsf{Prior}(ec_2, ec_1), *\mathsf{NOW}_i)$, *$ec_1$ is a cascade, and $\mathcal{F}(ec_1) >_{pr} ec_2$; or*
3. $\beta \not\vdash \mathsf{Holds}(\mathsf{Prior}(ec_2, ec_1), *\mathsf{NOW}_i)$, *$ec_2$ is a cascade, and $ec_1 >_{pr} \mathcal{F}(ec_1)$.*

The meta-theoretical relation $>_{pr}$ extends the notion of priorities from that directly defined by the object language's Prior to one covering acts embedded in cascades. The reason behind this move is that it is often the case that cascades are created on the fly, as Cassie acts, and such arbitrary cascades are not expected to be explicitly prioritized, in terms of Prior, with respect to other acts. Note that, in general, *inferring* such priorities is not unproblematic. In particular, one might think that the following possible pair of axioms may do the job.[2]

[2]Note that, to state these axioms, we need to assume a normal form for cascades. In particular, any cascade is equivalent to one that only has two elements. This is possible since the second element may itself be a cascade.

(10.1) $\mathsf{Prior}(ECv_1, ECv_2) \Rightarrow \forall ECv_3[\mathsf{Prior}(\mathsf{cascade}(ECv_1, ECv_3), ECv_2)]$

(10.2) $\mathsf{Prior}(ECv_1, ECv_2) \Rightarrow \forall ECv_3[\mathsf{Prior}(ECv_1, \mathsf{cascade}(ECv_2, ECv_3))]$

The intuition is that (10.1) captures point 2 of Definition 10.2 and (10.2) captures point 3. Unfortunately, this is not the case since the axioms miss the first conjunct in 2 and 3, the one requiring lack of contradictory information about the priorities of the two given acts. The problem is that the condition for $>_{pr}$ to hold made by points 2 and 3 of Definition 10.2 involves a non-monotonic (auto-epistemic) assumption—Cassie's lack of knowledge to the contrary. Since $\mathcal{FOCS}$ is an ordinary monotonic logic, it cannot represent this condition.

To see why $>_{pr}$ is needed and how its definition works, consider the following example discussed in Section 10.2. Suppose Cassie is about to perform the following act.

$$ec_1 = \mathsf{cascade}(\mathsf{Goto}(\text{room}), \mathsf{PickUp}(\text{block}), \mathsf{Goto}(\text{table}), \mathsf{PutDown}(\text{block}))$$

Also suppose that she has the following belief that whenever the battery goes low she should recharge it.

$$\mathsf{wheneverdo}(\mathsf{Low}(\text{battery}), \mathsf{Recharge}(\text{battery}))$$

Cassie starts performing ec_1. According to the semantics given in Section 9.8, this results in two things: Cassie's starting to perform the act $\mathsf{Goto}(\text{room})$ and her coming to believe that, when she is done, she should perform, ec_2, the rest of ec_1.

$$\mathsf{whendo}(\mathsf{In}(\text{room}), \mathsf{cascade}(\mathsf{PickUp}(\text{block}), \mathsf{Goto}(\text{table}), \mathsf{PutDown}(\text{block})))$$

Now, two things happen: Cassie comes to be in the room and her battery goes low. This activates both the whendo and the $\mathsf{wheneverdo}$ above, resulting in the two acts, ec_2 and $\mathsf{Recharge}(\text{battery})$ being scheduled for execution. What should Cassie do next? Note that since ec_2 is an internally-generated cascade, Cassie would probably have no beliefs about its priority with respect to other actions, including $\mathsf{Recharge}(\text{battery})$. However, Cassie should have a belief about the relative priorities of $\mathsf{Recharge}(\text{battery})$ and $\mathsf{PickUp}(\text{block})$, the first act of ec_2 (this may indeed be required by Axiom 10.1). In the given situation, it is reasonable to assume that recharging the battery has priority over picking the block up. Thus, by Definition 10.2, $\mathsf{Recharge}(\text{battery}) >_{pr} ec_2$. The reader might suspect that recharging the battery should have priority over *any* act, and, therefore, over any cascade, such that $>_{pr}$ is not needed after all. If this is the case, then consider the situation when Cassie has picked up the block and started moving towards the table. At this point, Cassie holds the following belief.

$$\mathsf{whendo}(\mathsf{Near}(\mathsf{table}), \mathsf{cascade}(\mathsf{PutDown}(\mathsf{block})))$$

Now, suppose that the battery goes low as Cassie reaches the table. In this case, Cassie has to decide whether to perform $\mathsf{Recharge}(\mathsf{battery})$ or $\mathsf{cascade}(\mathsf{PutDown}(\mathsf{block}))$. In this case, it is reasonable to assume that, since Cassie is near the table, putting down the block has priority over recharging the battery. By Definition 10.2, $\mathsf{cascade}(\mathsf{PutDown}(\mathsf{block})) >_{pr} \mathsf{Recharge}(\mathsf{battery})$.

The next step now is to make use of the notion of priority defined by $>_{pr}$ in order to determine which acts are to be performed given a set of acts A. Note that, given the definition of $>_{pr}$, we will be able, not only to schedule for execution acts in A, but also acts embedded within cascades in A (see Definition 10.4 below). To do this, however, we need to be able to get our hands on the first *executable* acts within a cascade. This, in general, is *not* the first act in the cascade (defined by $\mathcal{F}$), since the first act of a cascade may itself be a cascade, and the embedding may be indefinitely deep. For example, in executing the following cascade, ec_1 is the first act to be actually executed:

$$\mathsf{cascade}(\mathsf{cascade}(\mathsf{cascade}(ec_1, ec_2), ec_3), ec_4)$$

Such strange-looking cascades are unlikely to be explicitly represented in β as, for example, plans for composite acts. However, it is possible that they be generated by the acting executive in the course of decomposing some complex plan (for example, see the semantics of executing $\mathsf{sniterate}$ in Section 9.8). It should be clear, however, that every cascade is equivalent (with respect to execution) to some *flat* cascade—one that does not have any steps that are themselves cascades. For example, the above cascade is equivalent to the following one.

$$\mathsf{cascade}(ec_1, ec_2, ec_3, ec_4).$$

Let us make this more precise.

Definition 10.3 *For every cascade* $ec = \mathsf{cascade}(ec_1, \ldots, ec_n)$*,* $\{ec_i\}_{i=1}^{n} \subset \mathcal{EC}$*, ec is* **flat** *if, for every i,* $1 \leq i \leq n$*,* ec_i *is not a cascade.*

Algorithm $\texttt{flatten}(ec)$

1. If $ec = \mathsf{cascade}()$, then return ec.
2. If $\mathcal{F}(ec)$ is a cascade, then

   return $\texttt{flatten}(\mathsf{cascade}(\mathcal{B}(\mathcal{F}(ec)) \circ \mathcal{R}(ec)))$.
3. Return $\mathsf{cascade}(\langle \mathcal{F}(ec) \rangle \circ \mathcal{B}(\texttt{flatten}(\mathsf{cascade}(\mathcal{R}(ec)))))$.

Figure 10.1: The algorithm $\texttt{flatten}$.

The algorithm in Figure 10.1 transforms an arbitrary cascade into an equivalent flat one ($\circ$ is a sequence concatenation operator).[3]

Definition 10.4 *For every set of acts, $A \subseteq \Psi(\mathcal{EC})$, and for every $i \in \mathbb{N}$ ($i > 0$), define the set of acts $A_\top(i)$ as follows. For every act $ec \in \Psi(\mathcal{EC})$, $ec \in A_\top(i)$ iff, at $[\![{}^*\mathsf{NOW}_i]\!]$,*

1. *$ec \in A$, ec is not a cascade, and there is no $ec' \in A$ such that $ec' >_{pr} ec$ or*
2. *$ec' \in A$ is a cascade, there is no $ec'' \in A$ such that $ec'' >_{pr} ec'$, and $ec = \mathcal{F}(\texttt{flatten}(ec))$.*

Intuitively, $A_\top(i)$ is the set of acts in A, or embedded within cascades in A, that should be performed first, i.e, those with *top* priorities at $[\![{}^*\mathsf{NOW}_i]\!]$. A complementary set contains whatever remains.

Definition 10.5 *For every set of acts, $A \subseteq \Psi(\mathcal{EC})$, and for every $i \in \mathbb{N}$ ($i > 0$), define the set of acts $A_\bot(i)$ as follows. For every act $ec \in \Psi(\mathcal{EC})$, $ec \in A_\bot(i)$ iff, at $[\![{}^*\mathsf{NOW}_i]\!]$,*

1. *ec is not a cascade and $ec \in A \setminus A_\top(i)$,*
2. *ec is a cascade in A and $\mathcal{F}(\texttt{flatten}(ec)) \notin A_\top(i)$, or*
3. *ec' is a cascade in A, $\mathcal{F}(\texttt{flatten}(ec')) \in A_\top(i)$, and $ec = \mathsf{cascade}(\mathcal{R}(\texttt{flatten}(ec')))$.*

The argument, i, of $A_\top$ and $A_\bot$ will be dropped if NOW_i is assumed to be ${}^*\mathsf{NOW}$. How do the sets $A_\top$ and $A_\bot$ feature in the process of Cassie's deciding on what to do next? And how does this solve the problem of interrupt handling? The next section provides answers to these questions.

[3] I will not attempt to prove this, but the algorithm should be intuitive enough.

10.4 Prioritized Acting and Interrupt Handling

To perform actions according to their priorities, a control act, p-DoAll, is introduced. This is based on the DoAll control act which initiates a set of acts in some arbitrary order (Kumar, 1994; Shapiro and the SNePS Implementation Group, 1999).[4] Before describing how p-DoAll operates, one final definition is needed.

Definition 10.6 *For every $i \in \mathbb{N}$ ($i > 0$), the set of* **active processes** *at* $[\![\text{*NOW}_i]\!]$*, denoted* $\Pi(i)$*, is the set of acts Cassie is performing at* $[\![\text{*NOW}_i]\!]$*. More precisely,*

$$\Pi(i) = \{ec : ec \textit{ is an act and, at } [\![\text{*NOW}_i]\!],\ \beta \vdash \mathsf{Holds}(\mathsf{Prog}(ec), \text{*NOW}_i\}$$

Π *denotes the set of active processes at* $[\![\text{*NOW}]\!]$.

The control act p-DoAll schedules for execution a set of acts according to their priorities taking into account the set of active processes. In the $\mathcal{FOCS}$ language, "p-DoAll" denotes a function from the power set $\mathcal{EC}$ to $\mathcal{EC}$. Its operational semantics is given by the following reduction axiom for the activation of a p-DoAll act (in the spirit of Section 9.8.

- $\textit{Activate}(\mathsf{p\text{-}DoAll}(A)[e]) \rightarrow$

$$\begin{aligned} \textit{Activate}(\mathsf{cascade}(&\mathsf{DoAll}(\{\mathsf{Stop}(ec) | ec \in \Pi \cap A_\bot\}),\\ &\mathsf{DoAll}(A_\top \setminus \Pi),\\ &\mathsf{p\text{-}DoAll}(A_\bot))[e]) \end{aligned}$$

That is, if any of the acts to be prioritized is already being performed, the agent first stops it if it has a low priority (the Stop act). It then performs acts with top priorities unless they are already on-going. Finally, the agent performs a p-DoAll of the acts with low priorities (including those that were stopped). Thus, p-DoAll provides a mechanism for performing acts according to their priorities while taking into account the set of on-going processes.

How does p-DoAll solve the problem of interrupt handling? Cassie's acquisition of new information about the current state of the enviroment or her body is always accompanied by forward inference. Such information may be about an interrupting event, for example, the battery going

[4]Note that it just *initiates* them in arbitrary order. Once initiated, the acts may run in parallel.

low. Forward inference may activate believed "whendo"s or "$\mathsf{wheneverdo}$"s causing a reaction or a pending cascade to be scheduled for execution. All these acts are scheduled on an act stack, Σ (Kumar, 1994), where they are then picked for activation. However, to handle interrupts, one needs more than such a reactive behavior. In particular, we need to Give Cassie a chance to reason about the priorities of these scheduled acts. To model this kind of conscious reaction, a new mode of forward inference is introduced, one that may be called *prioritized forward inference* (PFI). With PFI, all the scheduled acts are replaced by a single $\mathsf{p\text{-}DoAll}$. More precisely, PFI results in the following:

$$\Sigma \longleftarrow \mathsf{p\text{-}DoAll}(\{\Sigma \cup \Pi\})$$ [5]

Thus, not only will Cassie reason about the priorities of the scheduled acts, but she will also take all the on-going processes into account, in case she needs to suspend any of them.

At this point, a detailed example would help. Consider the blocks-world example from Section 10.3. Cassie is about to perform the following act:

$$ec_1 = \mathsf{cascade}(\mathsf{Goto}(\mathsf{room}), \mathsf{PickUp}(\mathsf{block}), \mathsf{Goto}(\mathsf{table}), \mathsf{PutDown}(\mathsf{block}))$$

In the meantime, she is talking to Stu. Again suppose that Cassie has the following belief:

$$\mathsf{wheneverdo}(\mathsf{Low}(\mathsf{battery}), \mathsf{Recharge}(\mathsf{battery}))$$

As she starts performing the first step in ec_1, Cassie comes to believe that

$$\mathsf{whendo}(\mathsf{In}(\mathsf{room}), \mathsf{cascade}(\mathsf{PickUp}(\mathsf{block}), \mathsf{Goto}(\mathsf{table}), \mathsf{PutDown}(\mathsf{block})))$$

Now, suppose that Cassie reaches the room and, at the same time, senses the battery's going low. This activates the $\mathsf{wheneverdo}$ and whendo above resulting in two acts being scheduled onto the act stack:

$$\Sigma = \{\mathsf{Recharge}(\mathsf{battery}), \mathsf{cascade}(\mathsf{PickUp}(\mathsf{block}), \mathsf{Goto}(\mathsf{table}), \mathsf{PutDown}(\mathsf{block}))\}$$

By PFI, the contents of Σ get replaced by a single act:

$$\begin{aligned}\Sigma = \{\mathsf{p\text{-}DoAll}(\{&\mathsf{Recharge}(\mathsf{battery}),\\ &\mathsf{cascade}(\mathsf{PickUp}(\mathsf{block}), \mathsf{Goto}(\mathsf{table}), \mathsf{PutDown}(\mathsf{block})),\\ &\mathsf{TalkTo}(\mathsf{Stu})\})\}\end{aligned}$$

[5] I am assuming that Σ is, formally, a set, albeit an ordered one.

where $\Pi = \{\mathsf{TalkTo(Stu)}\}$. Let A be the argument of the scheduled p-DoAll. To execute this act, the sets $A_\top$ and $A_\perp$ need to be computed. Suppose that Cassie has the following belief about the priorities of recharging the battery and picking up the block.

$$\forall Tv[\mathsf{Holds(Prior(Recharge(battery), PickUp(block))}, Tv)]$$

Since it is reasonable to assume that talking to Stu does not use any of the resources required by recharging the battery, the two acts are not prioritized relative to each other. By Definition 10.2,

$$\mathsf{Recharge(battery)} >_{pr} \mathsf{cascade(PickUp(block), Goto(table), PutDown(block))}$$

Therefore, by Definition 10.4 (the first disjunct), $\mathsf{Recharge(battery)} \in A_\top$ and, by Definition 10.5 (the second disjunct), $\mathsf{cascade(PickUp(block), Goto(table), PutDown(block))} \in A_\perp$. Thus,

$$A_\top = \{\mathsf{Recharge(battery), TalkTo(Stu)}\}$$
$$A_\perp = \{\mathsf{cascade(PickUp(block), Goto(table), PutDown(block))}\}$$

Since $A_\perp$ and Π are disjoint, executing p-DoAll(A) results in first performing the act

$$\mathsf{DoAll(\{Recharge(battery)\})}$$

which amounts to simply recharging the battery, and then performing

$$\mathsf{p\text{-}DoAll(\{cascade(PickUp(block), Goto(table), PutDown(block))\})}$$

which amounts to resuming the performance of ec_1. Note that, all the while, talking to Stu continues uninterrupted.

Now, consider a slightly different scenario: the battery goes low as Cassie is moving to the table. At this point Cassie has the following belief:

$$\mathsf{whendo(Near(table), cascade(PutDown(block)))}$$

In addition, she has following belief about priorities:

$$\forall Tv[\mathsf{Holds(Prior(Recharge(battery), GoTo(table))}, Tv)]$$

At this point, the set of active processes and the act stack are, respectively, in the following state:

$$\Pi = \{\mathsf{GoTo(table), TalkTo(Stu)}\}$$
$$\Sigma = \{\mathsf{Recharge(battery)}\}$$

By PFI, the following p-DoAll gets scheduled onto the stack:

$$\mathsf{p\text{-}DoAll}(\{\mathsf{Recharge}(\mathsf{battery}), \mathsf{GoTo}(\mathsf{table}), \mathsf{TalkTo}(\mathsf{Stu})\})$$

Given Cassie's beliefs about the relative priorities of these acts, the set $A_\top$ and $A_\perp$ are defined as follows (I am, again, assuming that A is the argument of the scheduled p-DoAll):

$$A_\top = \{\mathsf{Recharge}(\mathsf{battery}), \mathsf{TalkTo}(\mathsf{Stu})\}$$
$$A_\perp = \{\mathsf{GoTo}(\mathsf{table})\}$$

Since $A_\perp \cap \Pi = \{\mathsf{GoTo}(\mathsf{table})\}$, then executing the scheduled p-DoAll results in, (i) stopping the act of going to the table, (ii) recharging the battery, and (iii) restarting the act of going to the table.

10.5 Conclusions

A cognitive agent should be capable of reasoning about the priorities of its acts in order to appropriately recover from errors and handle interrupts. The system presented here has a number of advantages over other proposed models in the symbolic AI literature.

1. It provides a general mechanism for prioritized acting using the p-DoAll control act. Interrupt handling comes out smoothly as a special case of prioritized acting.

2. Knowledge about priorities of acts may be given to the agent in natural language during an interaction with an operator. This may happen on-line, while the agent is acting.

3. Priorities are context-sensitive, changing according to various conditions in the environment.

4. Given the definition of the $>_{pr}$ relation and p-DoAll, the agent may interleave the execution of two cascades according to the priorities of acts in each.

Chapter 11

Implementation

This chapter discusses the implementation of the system developed in the previous chapters. Let me start by pointing out that the implementation is, in many respects, partial. The implementation of cascades (Chapter 9) and interrupt handling (Chapter 10), though pretty much complete and tested, still requires more extensive testing with more examples. The current implementation of the theory of time laid out in Chapters 5 through 8 is one of an earlier, less complete version of the theory. Thus, it needs to be (i) upgraded to reflect the current status of the theory and (ii) tested with a well-chosen suite of examples.

The system is implemented in Allegro Common Lisp and runs on University at Buffalo Computer Science and Engineering machines running Sun Microsystems' Solaris 5.8. A large part of the system is an extension/revision of SNePS 2.5 (Shapiro and the SNePS Implementation Group, 1999), the rest is part of the implementation of the agent Cassie, not of SNePS proper.

11.1 Examples of Cascades

In this section three simple examples are used to demonstrate the operation of cascades. The three examples are simulations of an agent carrying out the instructions represented by (9.1)-(9.3) in Section 9.1. More impressive examples require either errors and interrupts, my research on which is still in a premature stage,[1] or an actual robot acting in the world, something that cannot be

[1]They are also needed to demonstrate atelic acts.

presented on paper. The examples are only intended to give a feel of how cascading works. The demonstrations are the output of actual SNePS runs. These outputs are slightly edited for formatting and are broken down into sections to allow for explanation. The ":" is the SNePS prompt and inputs are either assertions, commands, or simulated sensory inputs.[2] Cassie's acts are simulated by generating English sentences describing what she is doing. These are surrounded by asterisks in the output.

First, we provide Cassie with some general rules about the goals of acts.

```
: all(x) (telic(pickup(x), holding(x))).

: all(x) (telic({walkto(x), goto(x), runto(x)}, at(x))).

: all(x,y) (telic(puton(x, y), on(x, y))).

: all(x,y) (telic(give(y, x), has(x, y))).
```

The above respectively assert that the goal of picking something up is to be holding it; the goal of walking, going, or running to some place is to be at that place;[3] the goal of putting some object, *x*, on some object, *y*, is for *x* to be on *y*; and the goal of giving some object, *y*, to some agent, *x*, is for *x* to have *y*.

The first example shows Cassie performing the sequence of acts presented in (9.1), repeated here for convenience.

(9.1) Pick up a block and then walk to the table and then put the block on the table.

In the initial situation Cassie is holding a block.

```
: holding(block).

: perform cascade(pickup(block),

                  walkto(table),

                  puton(block, table))

**Walking to TABLE**
```

[2] I am also eliminating time and the use of temporal progression algorithms to simplify the demonstration.

[3] Note the use of set arguments.

```
: perform believe(at(table))  ;;Sensory input. At the table.

**Putting BLOCK on TABLE**

: perform believe(on(block, table))  ;;Sensory input.

                                     ;;The block is on the table.
```

A couple of points to note. First, since the goal of picking up a block already holds, the act was skipped and the second step in the cascade was performed right away (see Section 9.7). Second, note that Cassie does not start putting the block on the table until she comes to know (via simulated perception) that she is at the table (note the prompt).

The second example shows Cassie acting according to (9.2).

(9.2) Run to the store and then buy a bottle of milk and then come back here.

This example illustrates two main points: (i) cascading control acts and (ii) reasoning and acting while performing a cascade. We first give Cassie recipes for performing some composite acts.

```
: all(x)(ActPlan(greet(x), {say("Hi", x), say("Hello", x)})).

: ActPlan(buy(bottle-of-milk),

          cascade(goto(dairy-section),

                  pickup(bottle-of-milk),

                  goto(cashier),

                  give(money, cashier)))
```

The first of these sentences says that to greet X (presumably a person) either say "Hi X" or "Hello X". The second says that to buy a bottle of milk (assuming that you're already in the store), go to the dairy section, pick up a bottle of milk, go to the cashier, and give money to the cashier. The ActPlan function associates a composite act with a plan to perform it (see Section 9.8).

Next, we ask Cassie to perform (9.2). To simplify matters, the last step of the cascade is for Cassie to run to the `house`. This matches (9.2) if we assume that the instruction was given in the `house`. At the same time, it avoids complications introduced by deictic expressions.[4]

[4]This does not mean that Cassie cannot understand deictic expressions. See various papers in (Duchan et al., 1995).

```
: perform cascade(runto(store),
                  buy(bottle-of-milk),
                  runto(house))
**Running to STORE**
```

Now, Cassie has started running to the store but has not gotten there yet. In the meantime, we can talk to Cassie and she can perform simultaneous acts as long as they do not interfere with her running to the store.

```
: all(x)(wheneverdo(near(x), greet(x))).
: perform believe(near(Stu))  ;;Sensory input. Stu is near.
Hello STU
```

The first of the above two sentences tells Cassie that whenever she's near someone, she should greet them (see Section 9.8 for the exact semantics of wheneverdo). By sensing that Stu is near, forward inference activates the rule, and Cassie greets Stu. The important point here is that Cassie can reason and act while in the midst of performing a cascade.

Having reached the store, Cassie carries out the plan for buying a bottle of milk, all the while observing the greeting rule.

```
: perform believe(at(store))  ;;Sensory input. At the store.
**Going to DAIRY-SECTION**
: perform believe(at(dairy-section))  ;;Sensory input.
                                      ;;Reached the dairy section.
**Picking up BOTTLE-OF-MILK**
: perform believe(holding(bottle-of-milk))  ;;Sensory input.
                                            ;;Holding the milk.
**Going to CASHIER**
: perform believe(near(Bill))   ;;Sensory input. Bill is near.
```

```
Hi BILL

: perform believe(at(cashier))  ;;Sensory input. Reached the cashier.

**Giving MONEY to CASHIER**

: perform believe(has(cashier, money))  ;;Sensory input.
                                        ;;The cashier has the money.

**Running to HOUSE**

: perform believe(near(Sally))  ;;Sensory input. Sally is near.

Hello SALLY

: perform believe(at(house))  ;;Sensory input. At the house.
```

The second step of the top-level cascade (buying the milk) is expanded into another lower-level cascade. It is only after the latter has been completed that Cassie resumed performing the former. Observe that Cassie started running back to the house only after (successfully) giving the money to the cashier. What initiated the act of running to the house? According to Definition 9.1, it is achieving the goal Complete(buy(bottle-of-milk)). Note that Cassie does not come to know that this goal has been achieved merely through sensory input; Cassie knows that by successfully finishing the lower-level cascade. Asserting the completion of the act of buying the milk happens internally, essentially when the last step of the lower-level cascade terminates successfully. The exact mechanism by which is done was outlined in Section 9.8.

The final example demonstrates Cassie's performance of the sequence of acts represented by (9.3).

(9.3) Stick a stamp on the envelope and then bring the secretary here and then give her the envelope.

In what follows, I assume *Gloria* to be the secretary.

```
: all(x) (ActPlan(bring(x), say("Come here", x))).

: telic(stickon(envelope, stamp), on(envelope, stamp)).

: telic(bring(Gloria), here(Gloria)).
```

The above defines the goals of sticking a stamp on the envelope and of bringing Gloria, and asserts that to bring someone just call them.

```
perform cascade(stickon(envelope, stamp),
                bring(Gloria),
                give(envelope, Gloria))
**Sticking ENVELOPE on STAMP**
: perform believe(on(envelope, stamp)) ;;Sensory input.
                                       ;;The stamp is on the envelope.
Come here GLORIA
```

At this point, Cassie is physically not doing anything. Having called Gloria, she can only wait for her to arrive in order to hand her the envelope. In the meantime, we can talk to Cassie and she can engage in other activities.

```
: good-secretary(Gloria).
: perform believe(near(Bill)) ;;Sensory input. Bill is near.
Hi BILL
: late(Gloria).
: perform believe(here(Gloria)) ;;Sensory input. Gloria arrives.
**Giving ENVELOPE to GLORIA**
: perform believe(has(Gloria, envelope)) ;;Sensory input.
                                         ;;Gloria has the envelope.
```

Merely terminating bodily activities would not have been an appropriate signal to start giving the envelope to Gloria. Cassie had to wait till she senses that Gloria has arrived. The performance of a cascade could be indefinitely extended in time since achieving some goals might be contingent upon exogenous events (in this case, Gloria's arrival).[5]

[5]Note that if the goal of bringing Gloria is for her to be *near*, rather than *here*, Cassie should have greeted her.

11.2 Crossing the Street

In this section, a demonstration of the solution to the problem of the fleeting now is presented. This is done using the example of crossing the street referred to throughout the dissertation. Let Cassie be the agent depicted in Figure 3.1. Cassie's task is to cross the street (from $S1$ to $S2$) and then walk into the store. In order to cross the street, Cassie wonders whether the walk-light is on "now". She turns her head toward the walk-light and observes that it is indeed on. Should she cross the street? Should she do it even though the "now" of the wondering is different from the "now" of the observation? Of course, since the two "now"s are indistinguishable at a coarse level of granularity determined by the typical duration of walk-lights.

The following demonstration shows that the solution of the problem of the fleeting now presented in Chapter 7 results in the appropriate behavior. In order to illustrate how exactly the solution works, I have turned on a tracing feature of the system that allows us to inspect the consecutive NOW-MTFs. In the sample run, NOW-MTFs are represented according to the generic schema shown in Figure 11.1. The representation shows the NOW-interval corresponding to the MTF and the sets of reference and state intervals in it. In addition, it shows all the states holding in the MTF with their associated state intervals and the reference intervals they span. Finally, it traces non-NOW reference intervals (those introduced by algorithm `state_query`) and how their actual durations progress toward their projected durations.

For example, the following shows the first NOW-MTF of the demonstration. Here, `B10` is the SNePS term denoting the current time (*NOW), which is the only reference interval in this MTF. There is one state interval, `B9`, associated with the state of Cassie's being on $S1$. Thus, initially, Cassie believes that she is on $S1$. Note that, in this run, state terms are echoed using the natural language generating system. Also, as will be shown below, Cassie reports what she is doing in natural language.

```
*----------------The Current MTF----------------*

 NOW: (B10)
```

However, in such a case, Gloria being near would activate two acts: the pending cascade, and the greeting. Such cases are still under investigation, though. The basic idea is to define a system of dynamic priorities among acts in the spirit of Chapter 10.

----------------The Current MTF----------------
NOW: <the value of *NOW >
Reference Intervals: <the set of reference intervals in this MTF, i.e., $\lceil$*NOW$\rceil$ >
State Intervals: <the set of state intervals in this MTF>
******
States:
< s_1: a state that holds in this MTF>
State interval: <the state interval associated with s_1 >
Reference intervals: <the set of reference intervals spanned by s_1 >
< s_2: another state that holds in this MTF>
State interval: <the state interval associated with s_2 >
Reference intervals: <the set of reference intervals spanned by s_2 >
.
.
.
< s_n: another state that holds in this MTF>
State interval: <the state interval associated with s_n >
Reference intervals: <the set of reference intervals spanned by s_n >
******
Durations of Reference Intervals:
< t_1: a reference interval in this MTF>
Projected HOM: <the HOM represented the projected duration of t_1 >
HOM so far: < the HOM of the amount of time elapsed since t_1 was introduced>
< t_2: a reference interval in this MTF>
Projected HOM: <the HOM represented the projected duration of t_2 >
HOM so far: < the HOM of the amount of time elapsed since t_2 was introduced>
.
.
.
t_n: a reference interval in this MTF>
Projected HOM: <the HOM represented the projected duration of t_n >
HOM so far: < the HOM of the amount of time elapsed since t_n was introduced>
____

Figure 11.1: The representation of NOW-MTFs in sample runs

```
 Reference Intervals: (B10)

 State Intervals: (B9)

******

 States:

    I am on S1.

        State interval: (B9)

        Reference intervals: (B10)

******

 Durations of Reference Intervals:

*________________________________________________*
```

We now tell Cassie to cross the street and walk to the store just in case the walk-light is on (note that this uses the snif control act presented in Chapter 9).[6]

```
* (perform
    (build action snif
           object1
           ((build condition (build object *walk-light property *on)
                   then (build action cascade
                              object1 (build action *cross
                                             object *the-street)
                              object2 (build action *walk
                                             |to| *the-store))))))
```

To perform the above act, Cassie turns to the walk-light in order to check if it is on. Note that this behavior is the result of a belief Cassie has about what she should do to determine whether the walk-light is on. Cassie reports what she is doing in English.

[6]This run makes use of a different SNePS interface language from the one used in the examples of Section 11.1.

```
 I am turning to the walk-light.
```

By starting to turn, NOW moves, giving rise to a new MTF.

```
*----------------The Current MTF----------------*

 NOW: (B17)

 Reference Intervals: (B17 B12)

 State Intervals: (B16 B9)

******

 States:

    I turn to the walk-light.

        State interval: (B16)

        Reference intervals: (B17)

    I am on S1.

        State interval: (B9)

        Reference intervals: (B17 B10)

******

 Durations of Reference Intervals:

  B12

      Projected HOM: 3

      HOM so far: 0

*_______________________________________________*
```

There are a couple of things to note here. First, note that a new NOW-interval is introduced (`B17`). Second, the state of Cassie's turning to the walk-light starts to hold (in fact, it is this transition that is responsible for the movement of NOW). Third, a new non-NOW reference interval is introduced (`B12`). Where does this interval come from? It is introduced by algorithm `state_query` as part of Cassie's wondering whether the walk-light is on (see Figure 7.13). In particular, Cassie's turning to the walk-light is the result of back-chaining on the eternal state of whether the walk-light is on at B12. In HOM-terms, the duration of B12 is projected to be 3. Why? According to algorithm state_query, this value is the result of applying backward projection to the typical duration of walk-lights. In the run presented here, the HOM representation of this duration is 4. This value is, informally, computed as follows. I assume that the typical duration of walk-lights is around 15 seconds, and the wave-length of the signal generated by the pacemaker (in this run) is 0.1 seconds. Thus, 150 cycles of the pacemaker correspond to the typical duration of walk-lights, and the HOM to which 150 belongs is the fourth one, hence the value 4. The projected duration of B12, which falls within the third HOM, is the result of applying backward projection, where the backward projection factor is taken to be $\frac{2}{3}$.

Now, Cassie turns (the "**" mark simulated action) and faces the walk-light. The resulting MTF is shown below.

```
**Turning**
 I turned to the walk-light.
 I am looking at the walk-light.
*-----------------The Current MTF-----------------*
 NOW: (B23)
 Reference Intervals: (B23 B12)
 State Intervals: (B22 B9)
******
 States:
    I look at the walk-light.
        State interval: (B22)
        Reference intervals: (B23)
    I am on S1.
        State interval: (B9)
```

```
        Reference intervals: (B23 B17 B10)

******

 Durations of Reference Intervals:

  B12

      Projected HOM: 3

      HOM so far: 2

*________________________________________________*
```

Note that, in the above MTF, the interval `B12` still persists, though it seems to be running out of rope. Also note that the state of turning toward the walk-light does not persist; this is taken care of without reasoning, thanks to modality variables. It is also the use of modality variables that is responsible for the persistence of the state of Cassie's being on *S*1.

Looking at the walk-light, Cassie determines that it is on.

```
 the walk-light is on.

*----------------The Current MTF----------------*

 NOW: (B27)

 Reference Intervals: (B27 B12)

 State Intervals: (B26 B22 B9)

******

 States:

    I look at the walk-light.

        State interval: (B22)

        Reference intervals: (B27 B23)

    I am on S1.

        State interval: (B9)

        Reference intervals: (B27 B23 B17 B10)
```

```
    the walk-light is on.

        State interval: (B26)

        Reference intervals: (B27 B23 B17 B12)

******

 Durations of Reference Intervals:

  B12

      Projected HOM: 3

      HOM so far: 3

*_______________________________________________*
```

In the above MTF, note that `B12` is within the set of reference intervals spanned by `B26`, the state interval associated with the walk-light's being on. Why? Because, according to algorithm `state_present` (see Figure 7.12), perceiving that the walk-light is on results in incorporating all members of $\lceil$*NOW$\rceil$ within the span of `B26`. Since `B12` is a member of $\lceil$*NOW$\rceil$, it makes it into the span of `B26`.

Thus, the deduction process succeeds (given the homogeneity of states) and Cassie starts crossing the street. Again, the appropriate setting of modality variables removes the states of being on *S*1 and looking toward the walk-light from the new MTF.

```
 I am crossing the street.

*-----------------The Current MTF-----------------*

 NOW: (B31)

 Reference Intervals: (B31 B12)

 State Intervals: (B30 B26)

******

 States:

    I cross the street.
```

```
        State interval: (B30)

        Reference intervals: (B31)

    the walk-light is on.

        State interval: (B26)

        Reference intervals: (B31 B27 B23 B17 B12)

******

 Durations of Reference Intervals:

  B12

      Projected HOM: 3

      HOM so far: 3

*_______________________________________________*


**Crossing the street**

 I crossed the street.

 I am on S2.


*----------------The Current MTF----------------*

 NOW: (B37)

 Reference Intervals: (B37)

 State Intervals: (B36 B26)

******

 States:

    I am on S2.

        State interval: (B36)
```

```
        Reference intervals: (B37)

    the walk-light is on.

        State interval: (B26)

        Reference intervals: (B37 B31 B27 B23 B17 B12)

******

 Durations of Reference Intervals:

*_______________________________________________*
```

At this point, `B12` had run out of rope and does not make it into the above MTF. Note, however, that, so far, Cassie still believes that the walk-light is on. Why? First, it has nothing to do with modality variables; check the above MTF, Cassie is not looking toward the walk-light. The reason is that, given Cassie's knowledge of the typical duration of walk-lights, and her sense of how much time has passed, she can assume that the walk-light is still on (see the discussion in Section 8.1).

Having crossed the street, the cascade effect takes place and Cassie starts walking to the store.

```
 I am walking to the store.


*-----------------The Current MTF-----------------*

 NOW: (B41)

 Reference Intervals: (B41)

 State Intervals: (B40 B36 B26)

******

 States:

    I walk to the store.

        State interval: (B40)

        Reference intervals: (B41)

    I am on S2.
```

```
        State interval: (B36)

        Reference intervals: (B41 B37)

    the walk-light is on.

        State interval: (B26)

        Reference intervals: (B41 B37 B31 B27 B23 B17 B12)

******

 Durations of Reference Intervals:

*_______________________________________________*

**Walking to the store**

 I walked to the store.

 I am in the store.

*----------------The Current MTF----------------*

 NOW: (B47)

 Reference Intervals: (B47)

 State Intervals: (B46)

******

 States:

     I am in the store.

        State interval: (B46)

        Reference intervals: (B47)

******

 Durations of Reference Intervals:

*_______________________________________________*
```

Note that, at this point, Cassie ceases to believe that walk-light is on; it took her long enough to walk into the store that the amount of time elapsed is longer that the typical duration of the walk-light.

11.3 The UXO World

This section demonstrates an implementation of the theory of interrupt-handling presented in Chapter 10. This will be done by showing a sample run of a recent incarnation of Cassie as a robot that is assigned the task of clearing a field of unexploded ordnance (UXO). Again, this Cassie only exists as a software simulation. The UXO-Cassie operates in an area consisting of four main zones: a field that possibly contains UXOs (Z1), a safe zone (Z2), a recharging zone (Z3), and a drop-off zone (Z4). Figure 11.2 shows a screen shot of a graphical simulation of the UXO-world.

The UXO-Cassie is equipped with a battery that continuously discharges as she operates. Should the battery reach a low enough level, Cassie should move to the recharging zone to recharge her battery. Cassie also carries charges that she may use to blow up UXOs. Her main task is to search the field for a UXO, and either blow it up using a charge or carry it to the drop-off zone and leave it there. The UXO-Cassie takes directions from, and reports to, a human operator in a fragment of English.

In the following run, Cassie will be shown while performing the act of clearing the field. This is a composite act for which Cassie has the following plan.

1 Search the field (Z1).

2 While there is a UXO

    3 Pick up the UXO.

    4 Go the the drop-off zone (Z4).

    5 Drop the UXO.

    6 Search the field.

7 Go to the safe zone (Z2)

At any point, if the battery goes low, Cassie should decide whether to interrupt what she is doing in order to recharge the battery, or to continue with the current task. The decision depends on Cassie's beliefs about the relative priorities of acts.

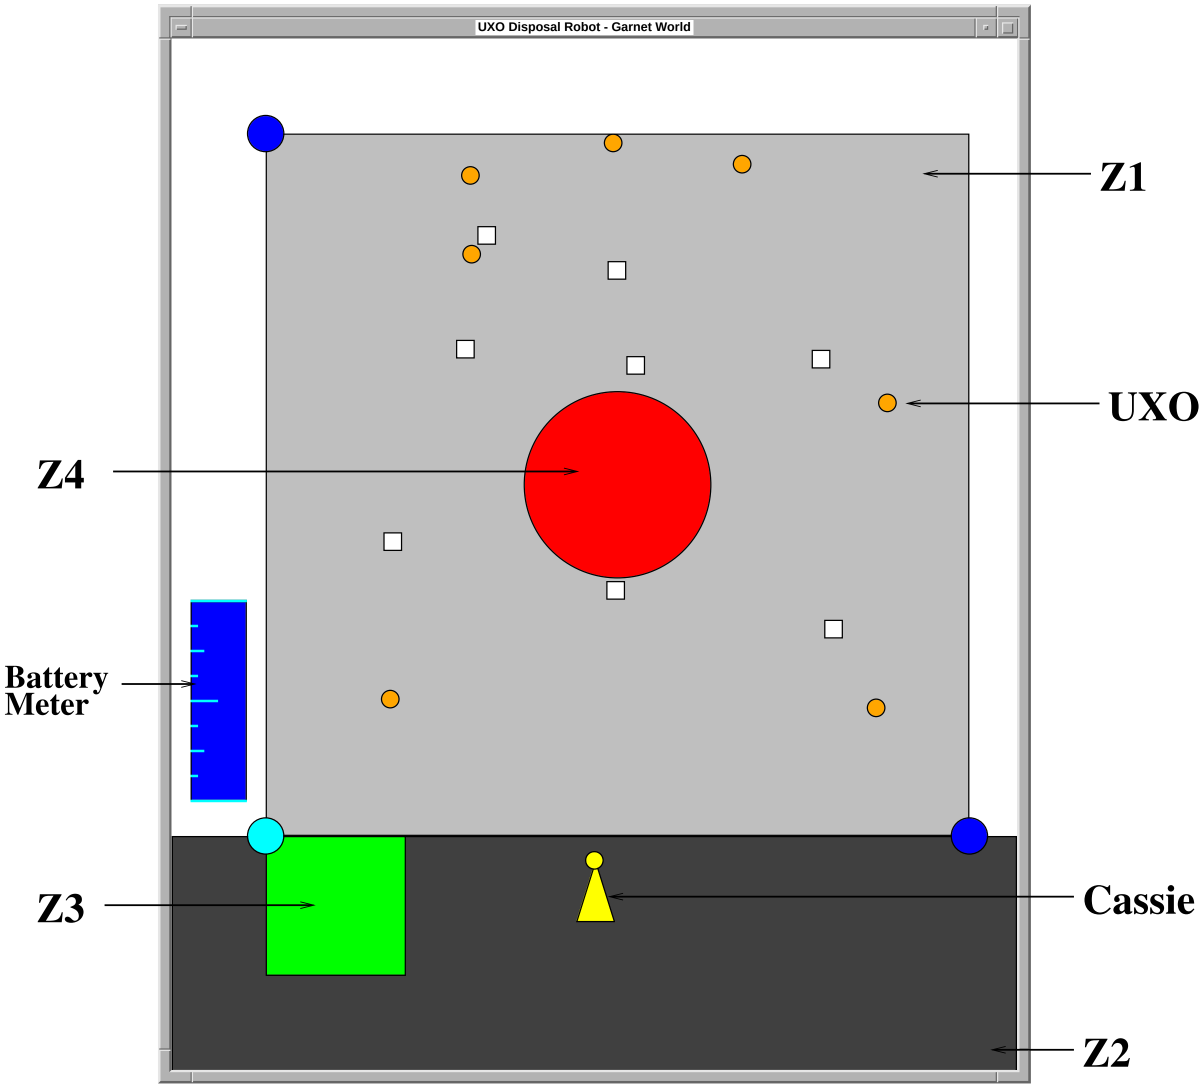

Figure 11.2: The UXO world.

```
 : Clear the field.

 I am going to Z1.

**The robot is going to WORLD:Z1.**

**The robot is in WORLD:Z1 at the point: (150.00, 5.00).**
```

What happens in the world is shown marked by "**". Orange objects are UXOs, white objects are other obstacles.

```
 I went to Z1.

 I am in Z1.

 I am searching.

**The robot is searching for a UXO . . .**

**Object found at: (204.06, 22.04).**

**The robot is going near the object

  (WORLD::ORANGE 204.0638 22.035358) . . .**

**The robot is looking at the object

  (WORLD::ORANGE 204.0638 22.035358).**

**The robot is near the object

  (WORLD::ORANGE 204.0638 22.035358).**

**The robot is going to examine object.**

**OBJECT FOUND IS A UXO.**

 I searched.

 I am near a UXO.

 I am picking up the UXO.

**The robot is picking up the UXO.**

 I picked up the UXO.
```

I am holding the UXO.

I am turning to Z4.

The robot is looking towards WORLD:Z4.

I turned to Z4.

I am looking at Z4.

I am going to Z4.

The robot is going to WORLD:Z4.

The robot is in WORLD:Z4 at the point: (35.81, 3.69).

I went to Z4.

I am in Z4.

I am dropping the UXO.

The robot is dropping the UXO.

I dropped the UXO.

I am empty handed.

I am turning to Z1.

The robot is looking towards WORLD:Z1.

I turned to Z1.

I am looking at Z1.

I am going to Z1.

The robot is going to WORLD:Z1.

The robot is in WORLD:Z1 at the point: (38.00, 3.69).

I went to Z1.

I am in Z1.

I am searching.

```
**The robot is searching for a UXO . . .**
**Object found at: (49.88, 109.08).**
**The robot is going near the object
  (WORLD::WHITE 49.881214 109.07636) . . .**
**The robot is looking at the object
  (WORLD::WHITE 49.881214 109.07636).**
**The robot is near the object
  (WORLD::WHITE 49.881214 109.07636).**
**The robot is going to examine object.**
**OBJECT FOUND IS NOT A UXO.**
**The robot is searching for a UXO . . .**
**Object found at: (84.77, 125.84).**
**The robot is going near the object
  (WORLD::ORANGE 84.7668 125.843216) . . .**
**The robot is looking at the object
  (WORLD::ORANGE 84.7668 125.843216).**
**The robot is near the object
  (WORLD::ORANGE 84.7668 125.843216).**
**The robot is going to examine object.**
**OBJECT FOUND IS A UXO.**
 I searched.
 I am near the UXO.
 I am picking up the UXO.
**The robot is picking up the UXO.**
```

```
 I picked up the UXO.
 I am holding the UXO.
 I am turning to Z4.
**The robot is looking towards WORLD:Z4.**
 I turned to Z4.
 I am looking at Z4.
 I am going to Z4.
**The robot is going to WORLD:Z4.**
**The robot is in WORLD:Z4 at the point: (19.69, 30.14).**
**<<<THE BATTERY IS LOW>>>.**
```

The battery goes low while Cassie is about to drop the UXO. However, Cassie believes that dropping the UXO has higher priority over recharging the battery if in Z4.

```
 I went to Z4.
 I am in Z4
 and the battery is low.
 I am dropping the UXO.
**The robot is dropping the UXO.**
 I dropped the UXO.
 I am empty handed.
```

Now, Cassie can attend to the low-battery.

```
 I am turning to Z3.
**The robot is looking towards WORLD:Z3.**
 I turned to Z3.
 I am looking at Z3.
```

```
 I am going to Z3.

**The robot is going to WORLD:Z3.**

**The robot is in WORLD:Z3 at the point: (19.69, -5.00).**

 I went to Z3.

 I am in Z3.

 I am recharging the battery.

**The robot is recharging the battery.**

 I recharged the battery.

 the battery is full.
```

Having recharged the battery, Cassie now resumes the act of clearing the field.

```
 I am turning to Z1.

**The robot is looking towards WORLD:Z1.**

 I turned to Z1.

 I am looking at Z1.

 I am going to Z1.

**The robot is going to WORLD:Z1.**

**The robot is in WORLD:Z1 at the point: (38.00, 5.00).**

 I went to Z1.

 I am in Z1.

 I am searching.

**The robot is searching for a UXO . . .**

**Object found at: (30.00, 151.88).**

**The robot is going near the object

  (WORLD::WHITE 30.004522 151.88202) . . .**
```

```
**The robot is looking at the object
  (WORLD::WHITE 30.004522 151.88202).**
**The robot is near the object
  (WORLD::WHITE 30.004522 151.88202).**
**The robot is going to examine object.**
**OBJECT FOUND IS NOT A UXO.**
**The robot is searching for a UXO . . .**
**Object found at: (25.77, 166.43).**
**The robot is going near the object
  (WORLD::WHITE 25.7723 166.42965) . . .**
**The robot is looking at the object
  (WORLD::WHITE 25.7723 166.42965).**
**The robot is near the object (WORLD::WHITE 25.7723 166.42965).**
**The robot is going to examine object.**
**OBJECT FOUND IS NOT A UXO.**
**The robot is searching for a UXO . . .**
**Object found at: (76.54, 183.86).**
**The robot is going near the object
  (WORLD::WHITE 76.54233 183.85709) . . .**
**The robot is looking at the object
  (WORLD::WHITE 76.54233 183.85709).**
**The robot is near the object
  (WORLD::WHITE 76.54233 183.85709).**
**The robot is going to examine object.**
```

```
**OBJECT FOUND IS NOT A UXO.**

**The robot is searching for a UXO . . .**

**<<<THE BATTERY IS LOW>>>.**
```

The battery goes low again. However, this time, recharging the battery is more important than searching.

```
 the battery is low.

 I am turning to Z3.

**The robot is looking towards WORLD:Z3.**

 I turned to Z3.

 I am looking at Z3.

 I am going to Z3.

**The robot is going to WORLD:Z3.**

**The robot is in WORLD:Z3 at the point: (19.69, -5.00).**

 I went to Z3.

 I am in Z3.

 I am recharging the battery.

**The robot is recharging the battery.**

 I recharged the battery.

 the battery is full.

 I am turning to Z1.

**The robot is looking towards WORLD:Z1.**

 I turned to Z1.

 I am looking at Z1.

 I am going to Z1.
```

The robot is going to WORLD:Z1.

The robot is in WORLD:Z1 at the point: (38.00, 5.00).

I went to Z1.

I am in Z1.

I am searching.

The robot is searching for a UXO . . .

Z1 is cleared.

I am turning to Z2.

The robot is looking towards WORLD:Z2.

I turned to Z2.

I am looking at Z2.

I am going to Z2.

The robot is going to WORLD:Z2.

The robot is in WORLD:Z2 at the point: (38.00, -5.00).

I went to Z2.

I am in Z2.

Chapter 12

Conclusions and Future Work

In this chapter, we summarize the system proposed in the dissertation, point out its contributions to research on cognitive robotics, and discuss possible future directions.

12.1 Summary

This dissertation presents a theory of an embodied cognitive agent reasoning and acting in time. An agent reasoning and acting in time is one that interleaves reasoning, action execution, perception, and interaction with other agents while maintaining a sense of the passage of time and a record of what has happened and is happening. The theory could be viewed as consisting of three components that interact with each other at various levels. These are the logical $\mathcal{FOCS}$ system developed in Chapters 3 and 4, the meta-logical theory of subjective time presented in Chapters 5 through 8, and the theory of sequential acting and interrupt handling laid out in Chapters 9 and 10.

12.1.1 $\mathcal{FOCS}$

$\mathcal{FOCS}$ is a first-order sorted logic viewed as a calculus of states. The prominence of *states* in the logic is based on their primacy for an embodied cognitive agent reasoning and acting in time. In particular, for such an agent, the experienced world is a sequence of temporal frames (NOW-MTFs) each corresponding to a set of states that hold at any "now". Based on its experience of states holding the agent should be able to derive the occurrence of *events*—temporal entities that are never directly experienced by the agent since they are not located at any particular "now".

States are proposition-like entities that homogeneously hold over time. The traditional propositional connectives and quantifiers have been extended in $\mathcal{FOCS}$ to cover the domain of states. Adopting an Allen-style interval calculus, the core $\mathcal{FOCS}$ ontology partitions states into four categories: eternal, temporary, permanent, and co-permanent. Eternal states are states that always hold or never hold; they play the role of propositions in the logic. Temporary states are ones that freely start and cease to hold. They feature prominently in the overall theory since they are the only sort of states that change unrestrictedly as time goes by. Permanent states are states which, once started, never cease to hold. Though of no practical value, co-permanent states complete the matrix-like categorization of states; they are states that, once they cease, never start again.

In addition to states, a supervenient sort of events is introduced. Events are unitary situations, exhibiting a type-token distinction, that heterogeneously occur over time. Two general sorts of events are distinguished: punctual events and durative events. An agent never experiences an event, whether punctual or durative; it may only infer its occurrence. Punctual events, which are mainly state transitions, are determined to have occurred when a state switches from not holding to holding (or vice versa). Durative events are complex temporal entities made up of a state's starting to hold, its holding for a while, and its ceasing to hold. Expanding the traditional linguistic notion of telicity, $\mathcal{FOCS}$ includes axioms characterizing five different types of events. The following are examples of these types.

1. $\overset{\rightarrow\bullet}{\mathsf{telic}}$: John ran to the store. (The event is complete when and only when John stops running and is at the store)

2. $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$: John pushed the rock down the hill into the river. (The event is complete when and only when the rock is in the river; John may have stopped pushing long before that)

3. $\overset{\longleftarrow}{\mathsf{atelic}}$: John ran toward the store. (The event is complete any time John stops running before reaching the store)

4. $\overset{\longrightarrow}{\mathsf{atelic}}$: John ran past the store. (The event is complete any time John stops running after reaching the store)

5. $\overleftrightarrow{\mathsf{atelic}}$: John ran. (No restrictions on when the event is complete)

For each type of event, $\mathcal{FOCS}$ includes an axiom allowing the agent to infer the occurrence of an event based solely on its knowledge of states holding and not holding—the only thing that it can experience.

12.1.2 Subjective Time

An agent's subjective sense of time involves two main components: a representation of "now" that continuously changes reflecting the progression of time, and a feel for how much time has passed. Due to the granular vagueness of the concept of "now", a mereological meet semi-lattice of intervals represents the agent's sense of "now" at different levels of granularities. A meta-logical variable, NOW, that assumes values from amongst the set of time-denoting $\mathcal{FOCS}$ terms, represents the agent's notion of "now" at the finest level of granularity; its value is the smallest element of the lattice of "now"s at any time. Values of NOW form a chain ordered by the temporal precedence relation, and the progression of time is modeled by introducing a new term at the end of the chain and making it the new value of NOW.

The agent's sense of how much time has passed is based on an internal sub-conscious clock, the pacemaker, that primarily gives the agent a feel for the duration of NOW-intervals. Two durations may *feel* somewhat different but, nevertheless, be indistinguishable at the level of conscious reasoning. Therefore, a quantized representation of durations was introduced to ground the amount-denoting $\mathcal{FOCS}$ terms. Amounts are quantized into intervals each corresponding to one half-order of magnitude (Hobbs, 2000). Both the use of half-order of magnitudes for quantization and the pacemaker for the sense of time are not central to the theory; they could be replaced by other, more suitable alternatives if needed.

The theory makes the reasonable assumption that the value of NOW changes (reflecting the progression of time) when and only when there is a detectable state change in the agent's environment. At the current state of the theory, the environment does not include states of the agent's mind; thus, pure reasoning does not take time as far as the theory is concerned. However, reasoning may take time if it involves performing some actions (mainly, sensory actions) in order to add a missing link to a chain of inference. For NOW to change, then, the agent must perceive, proprioceive, or be told about some state change. We have outlined a number of principles that govern state change as determined by each of these means (in addition to inference). Based on these principles, a collection of algorithms was introduced to account for what happens when time passes. In addition to changing NOW, other components of the system need to be updated—most importantly, the agent's belief space, since the passage of time always involves new beliefs about the detected changes. It should also be noted that the algorithms are general enough to account for multiple simultaneous changes including onsets and/or cessations of various states.

In addition to updating NOW and the agent's belief space, a set of modality registers are updated whenever a change is determined through perception or proprioception. Modality registers contain $\mathcal{FOCS}$ propositions representing what each agent modality is being used for or what it perceptually conveys about the environment. The main utility of these registers is in the smooth, reasoning-free projection of bodily and perceived states as time passes by. We have outlined an algorithm (`setup_new_MTF`) that is invoked every time NOW changes and that is responsible for extending continuing states over the new NOW. The algorithm makes use of the agent's knowledge of how much time has passed and typical durations of states in order to defeasibly determine whether a state continues to persist. A number of results were presented demonstrating the agent's *temporal rationality*—that its beliefs about whether a state holds "now", or whether it has moved into the past, are justified by what it knows, what it perceives, what it proprioceives, and what it feels regarding how much time has passed.

Interleaving reasoning and acting in time raises two problems with reasoning about "now". The problem of *the unmentionable now* emerges when we attempt to represent rules or plans of action that intuitively involves reference to future, yet unknown, "now"s. The solution proposed by the dissertation builds "now" into the inference and acting systems revealing the intimate involvement of time in the agent architecture. The problem of *the fleeting now* emerges when the agent is reasoning about "now", when the very process of reasoning results in "now" moving. Basically, the agent is interested in whether some state holds "now". However, since reasoning and sensory acts take time, whatever conclusion it makes will be, strictly speaking, about a different "now". The proposed solution is based on the simple intuition that the concept of "now" is vague as to the size of the interval it represents. The agent wonders whether the state holds, not at the sharp moment of experience, but over a broader "now", an interval whose duration is comparable to the typical duration of the state. Such an interval may still be "now", relative to some coarse level of granularity, at the time of the conclusion. Whether this is the case depends on the amount of time it takes the agent to reach a conclusion—something that it can determine, given its sense of time.

12.1.3 Action Execution

To execute a sequence of acts, an agent must start executing one step when and only when the previous step has been completed. We have argued that it is often the case that determining when an act is complete requires reasoning, not just monitoring of bodily activity. Thus, the apparently simple notion of sequential acting is, after all, not as simple as it may seem. Using the newly-introduced cascading mechanism, an agent performs a sequence of acts by starting to perform the first act and forming the belief that, when it is complete, it should perform the rest of sequence. Determining when an act is complete depends on what type of act it is. Acts that fall under any of the telicity-based categories discussed above are determined to be complete given the $\mathcal{FOCS}$ axioms characterizing their completion conditions. However, we have shown that it is often possible to side-step the lengthy reasoning process required to determine completion by recognizing the notion of the *goal* of an act—a state that signals its completion. Control acts that determine how a collection of acts are to be executed (for example, sequentially, iteratively, etc.) are determined to be complete by the very process of executing them. We have provided operational semantics for the execution and completion of those acts.

With cascades, a sequence of acts is broken up in such a way that, as the agent is performing a step in the sequence, it is aware and can reason about what is left to be done. This feature allows the agent to smoothly weave reasoning about interrupts and errors into action execution. The interrupt handling system presented in Chapter 10 is based on a dynamic system of priorities among acts. Priorities are represented by explicit beliefs of the agent about which acts have higher priority in different situations. This explicit representation of context-sensitive priorities distinguishes the system proposed from other systems found in the AI literature. The acting executive handles interrupts by taking priorities into account in deciding what to do next. Every time the agent decides to perform some new set of acts, it reasons about the priorities of these acts together with anything else that it might be in the midst of doing. Thus, the agent may decide to suspend some of its on-going activities if they do not have top priority. Resumption is possible in most cases (see below for limitations) given a careful use of cascading.

12.2 Future Directions

The research presented in the dissertation paves the road for further extensions, refinements, and investigations. Some of the possible future directions are discussed below.

12.2.1 Durations and the Sense of Time

Reasoning about durations as presented here is confined to reasoning about their order with respect to the relation $<_Q$. This is sufficient for the purposes of this dissertation. However, representing and reasoning about durations has a broader range of applications in temporal reasoning (Allen and Kautz, 1988, for example). For example, consider an agent that knows that a bus arrives every 10 minutes and waits at the bus stop for 3 minutes. Given that a bus was at the bus stop at some time, t_1, was there a bus at t_2? In general, to answer this question, the agent needs to know the amount of time separating t_1 and t_2. Such knowledge may be explicitly conveyed to the agent by direct assertion. But, in general, it need not. For example, the agent might know that there is some t_3 such that $t_1 \prec t_3 \prec t_2$, and that the amount of time between t_1 and t_3 is 5 minutes and that between t_3 and t_2 is also 5 minutes. In such a situation, the agent should be capable of answering the question. This may only be achieved, however, if the agent can engage in some arithmetical reasoning about durations—something that is not accommodated by our theory.

Even if one were to dismiss the kind of reasoning required to cope with the situation presented above, consider a case where both t_1 and t_2 are NOW-intervals. In such a case, the duration between t_1 and t_2 is given by $d_t(\Phi(t_1), \Phi(t_2))$ (see Definition 7.2). This quantity exists at the PML and may be efficiently computed there without the need for any arithmetical reasoning. What is needed then is to allow such a computation as part of the reasoning process. This may indeed be achieved if one thinks of performing this computation as a sensory act that may be executed to add a missing link to a chain of reasoning. How this exactly is to be done requires further investigation.

In addition to temporal reasoning, the agent's sense of time may be used in planning and plan execution (Vere, 1983; Ambros-Ingerson and Steel, 1988, for example). Consider the example of bringing the secretary from Chapter 9. The plan was for the agent to call the secretary. Having called her, the agent would just wait there until the secretary arrives. Practically, however, there should be limits on this waiting period. In particular, a more reasonable plan for bringing the secretary (Gloria) may be given by the following iterative act (see Section 4.1 for a discussion of sniterate:

$$\mathsf{sniterate}(\{\langle\neg\mathsf{Here}(\text{Gloria}),\mathsf{cascade}(\mathsf{Say}(\text{Come here Gloria}),\mathsf{Wait}(q))\rangle\},\mathsf{Say}(\text{Hi Gloria}))$$

Thus, the agent calls the secretary, waits for some time, and then, if she has not yet arrived, calls her again. The term q above may be the typical duration it takes the secretary to arrive after being called, a period after which it may be assumed that the action has failed and needs to be repeated.[1] The waiting action makes use of the pacemaker to determine when a period of time equivalent to $[\![q]\!]$ has elapsed. In general, using its sense of time, the agent may impose certain constraints on when it should achieve some goals, and an action may be determined by the agent to have failed if it continues past a dead-line. How to effectively integrate the sense of time with action execution is again worth further investigation.

12.2.2 Interrupts and Modalities

Section 12.1.3 reviewed our current theory of interrupt handling. The basic idea is for the agent to reason about the priorities of acts it is performing or about to perform in order to figure out which need to wait, which are to be suspended, and which should be initiated immediately. If A is the set of those acts, the agent computes sets $A_\top$ and $A_\bot$ such that the prioritized execution of acts in A amounts to performing everything in $A_\top$, and *then* prioritizing and performing everything in $A_\bot$.

The presence of cascades in A makes things more complicated. Since, for a cascade, it is the priority of the first act that matters, then it is that act's priority that gets compared to the priorities of other acts in $\mathcal{A}$. In the current theory (see Chapter 10), if the first act of a cascade has top priority, it is considered a member of $A_\top$, with the rest of the cascade in $A_\bot$. The problem with this is that, even though the rest of the cascade should be initiated once its first act completes, it has to wait for other unrelated acts in $A_\top$ to finish. In fact, the problem is not only confined to the case of cascades. Consider the example from Chapter 10, of the agent that decides to recharge its battery while carrying a block to some table and talking to some other agent, Stu.

[1] Needless to say, the plan given above is still naive; the agent might try another way to bring the secretary if calling does not work for a couple of times.

- $A_\top = \{\mathsf{TalkTo(Stu)}, \mathsf{Recharge(battery)}\}$.

- $A_\bot = \{\mathsf{GoTo(table)}\}$

Both Recharge(battery) and GoTo(table) use $\mathsf{LOCOMOTION}_{prop}$ and VISION_{prop}, while TalkTo(Stu) uses $\mathsf{AUDITION}_{prop}$ (see Section 5.2.3). Intuitively, the agent should start moving to the table once the battery has been recharged. However, given the current scheme, it will only move to the table when both Recharge(battery) and TalkTo(Stu) finish.

The proposed revision of our current theory requires the agent to reason about its modalities. For an agent with n proprioceptual modalities ($\mu_1, \ldots, \mu_n$), instead of a pair of sets $A_\top$ and $A_\bot$, there should be a set of cascades, $\{C_i\}_{i=1}^{n}$, one for each modality, μ_i. Each of these cascades contains exactly two elements: $\top_i$ and $\mathsf{p\text{-}DoAll}(\bot_i)$. $\top_i$ is the highest-priority act that uses μ_i, and $\bot_i$ is the set of acts that need to wait for $\top_i$ to complete before they can be initiated. Note that members of $\bot_i$ are either low-priority acts using μ_i or tails of cascades for which $\top_i$ is the head. Actually, things are more complicated since the same act may occupy different modalities and need not be top-priority with respect to all of them.

12.2.3 Resumption

Suppose the agent is performing a sequence of acts, $\mathsf{cascade}(ec_1, \ldots, ec_n)$, when an interrupt, $\mathfrak{I}$, occurs requiring some reaction $\mathfrak{R}(\mathfrak{I})$. The view of interrupt handling implied by the current theory reduces the issue to that of inserting $\mathfrak{R}(\mathfrak{I})$ somewhere within the sequence $\langle ec_1, \ldots, ec_n \rangle$ based on priority. Although, in many cases, such a view suffices, it is generally not sufficient for appropriately handling interrupts. The sequence $ec_1, \ldots, ec_n$ is not just an arbitrary collection of acts to be performed in some accidental order; such a sequence is often a plan for achieving a particular goal. Thus, the execution of each act is generally dependent on the state brought about by the execution of previous acts in the sequence. The introduction of some arbitrary act, $\mathfrak{R}(\mathfrak{I})$, within such a structure may disturb the dependencies among the acts in the original sequence, rendering the execution of those that follow $\mathfrak{R}(\mathfrak{I})$ either useless or impossible. Evidently, the problem of interrupt handling is not simply that of appropriately scheduling a reaction. Another main aspect of the problem is figuring out how to resume the interrupted behavior.

In general, suppose that $\mathfrak{R}(\mathfrak{I})$ gets performed just before an act ec_i ($1 < i \leq n$) in the sequence

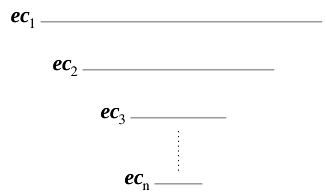

Figure 12.1: A nest of acts.

$ec_1, \ldots, ec_n$. Currently, the system supports two types of resumption: (i) resuming exactly where it left off, at ec_i; and (ii) resuming at some ec_j, $i < j \leq n$, in case the performance of $\mathfrak{R}(\mathfrak{I})$ results in the goals of all the acts ec_k, $i \leq k < j$, being achieved (see the description of cascades in Section 9.7). The final case which the current theory does not support is if the agent needs to resume at some ec_j, $1 \leq j < i$. This may happen if performing $\mathfrak{R}(\mathfrak{I})$ undoes the effects of some early actions, thereby changing the state of the world in such a way that resuming at $\mathfrak{R}(\mathfrak{I})_i$ may not result in achieving the goal of the original sequence, or may even be impossible.

An important feature that our agent seems to be lacking is a *deep* understanding of what it is doing. Not only should the agent know that it is performing a certain sequence of acts (which it currently does), but it should also be aware of the more global goals and behaviors of which the sequence is a low-level implementation. It is only through such awareness that the agent may determine how to resume an interrupted sequence.

12.2.4 “What am I doing?”

One of the primary constraints on our theory of agents is natural language competence. In the context of an acting agent, a very possible question by a human operator is “what are you doing?”. Answering such a question is not as easy as it may seem. Suppose that the agent is performing a composite act ec_1. To perform ec_1, the agent carries out a sequence of acts whose first act is ec_2. ec_2 is itself a composite act that involves performing yet another composite act ec_3. This nesting of acts grows deeper until a primitive act, ec_n, is performed (see Figure 12.1). How should the agent answer the question “what are you doing?”, if asked at a time at which it is performing ec_n?[2] All the acts that ec_n is nested within are possible answers. As far as natural-language interaction is concerned,

[2]Putting aside the issue of concurrent acting.

the best answer is a pragmatic issue that is beyond the scope of this research. However, the agent may need to ask itself the corresponding question, "what am I doing?", as a step of reasoning in the service of acting. For example, to correctly resume a suspended sequence, the agent may need to have a global understanding of why it is performing the sequence—an answer could be found higher in the nest of acts. Also, the priority of an act may depend on what it is performed for—again, something that depends on top-level acts the agent is performing.

Making the agent aware of such a temporal-causal nesting of its acts might not be hard to incorporate in the existing theory/system. However, some issues need to be settled first. What exactly is the relation between acts in the nest? If a low-level act is interrupted, is a high-level act also interrupted? Or is it still on-going in some broad sense (see (Galton, 1984) on the *broad sense* of the progressive)?

12.2.5 Modality, Causality, and Auto-epistemic Logic

At various points, we have identified certain deficiencies in the $\mathcal{FOCS}$ system. For example, in Section 3.7, it was pointed out that, in order to properly account for intuitions about the different sorts of states, $\mathcal{FOCS}$ ought to include some modal notion of possibility. The main point is a theoretical one; in a modal setting,[3] we can provide sufficient conditions for a state to be eternal, permanent, co-permanent, or temporary. Given such conditions, rather than posit, we can *prove* that the complements of eternal, permanent, co-permanent, and temporary states are eternal, co-permanent, permanent, and temporary, respectively. We can also prove similar results for complex (conjunctive) states.

Another hole in the $\mathcal{FOCS}$ system is its lack of any account of causality, as pointed out in Section 4.6.4.[4] The problem was identified in the discussion of Axioms **AE14–AE18** that allow an agent to infer that a given durative event is of a particular category. Consider, again, the example of bringing the secretary. The agent calls the secretary and according to the axioms (**AE15**, in particular), once she arrives, it identifies this calling event as one of *its* bringing the secretary. The problem, however, is that the secretary might have arrived for some reason other than her heeding the call (for example,

[3]Recall that we can continue to live without modal operators if we introduce reified possible worlds into the $\mathcal{FOCS}$ ontology.

[4]This is not a problem with $\mathcal{FOCS}$ proper, but the particular $\mathcal{FOCS}$ theory presented here.

she might have just been passing by). Intuitively, the agent's conclusion is not unreasonable, but *logically* it is not totally sound. The main problem of course is coming up with an epistemological theory that accounts for conditions allowing cognitive agents to recognize causality. Evidently, this involves at least logic, philosophy, and psychology.

At different points, we have also noted that the system might use some auto-epistemic enhancements. For example, in Section 4.6.4, it was pointed out that the antecedents of Axioms **AE14**–**AE18** contain an important conjunct asserting that certain events do not occur. Although such a conjunct may be easily stated in the axioms, it is generally not possible to infer it, rendering the axioms unusable. Instead of a proposition asserting that an event does not occur, all that is needed is one asserting that the agent does not know of any such occurrence. For *that* to be inferred, some form of auto-epistemic reasoning seems inevitable.

The need for auto-epistemic reasoning was also witnessed in the discussion of priorities in Section 10.3. In particular, it was noted that the the definition of the meta-logical relation $>_{pr}$, which involves the agent's lack of knowledge, may actually be represented in the logic if it includes auto-epistemic features. In addition, the agent's defeasible assumptions about the persistence of states based on their typical durations and its sense of time may be represented in an auto-epistemic logic rather being proceduralized in algorithm `setup_new_MTF`. Needless to say, we would have to allow the agent to interleave PML computations based on the pacemaker within the reasoning process (see Section 12.2.1).

12.2.6 Implications for Linguistics and Psychology

In addition to possible future direction within AI proper, some of the ideas presented in the dissertation have implications for research in linguistics and psychology. In particular, the analysis of telicity in Chapter 4 presents distinctions that seem to have been overlooked by the research on linguistic aspect. It would be interesting to investigate whether the five-way analysis of telicity presented here has linguistic significance—for example, whether it is possible to come up with linguistic tests distinguishing $\overset{\rightarrow\bullet}{\mathsf{telic}}$, $\overset{\rightarrow\ldots\bullet}{\mathsf{telic}}$, $\overleftarrow{\mathsf{atelic}}$, $\overrightarrow{\mathsf{atelic}}$, and $\overleftrightarrow{\mathsf{atelic}}$ sentences; or whether these classes are morpho-syntactically marked in any languages.

The discussion of typical durations and backward projection in Chapter 7.2 also raises some questions for psychologists of time. To recapitulate, there are at least three questions:

1. What kinds of mental representations and processes are involved in reasoning about typical durations of states?
2. As regards the backward projection factor, what are the biases of human subjects as to how long a perceived state has been holding?
3. What are the factors determining those biases?

12.2.7 Implementation and Testing

Finally, it remains to be said that one important task for the future is to fully-implement and test the theories presented here. As I pointed out in Chapter 11, the current implementation is only partial, and that extensive and careful testing is needed to validate it.

Appendix A

The $\mathcal{FOCS}$ System

A.1 Time

A.1.1 Axioms

- **AT1.** $Tv \prec Tv' \Rightarrow \neg[Tv' \prec Tv]$.
- **AT2.** $[Tv_1 \prec Tv_2 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_1 \prec Tv_3$.
- **AT3.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow \neg[Tv \prec Tv']$.
- **AT4.** $[Tv \sqsubseteq Tv' \wedge Tv' \sqsubseteq Tv] \Rightarrow \mathsf{Equiv}(Tv, Tv')$.
- **AT5.** $[Tv_1 \sqsubseteq Tv_2 \wedge Tv_2 \sqsubseteq Tv_3] \Rightarrow Tv_1 \sqsubseteq Tv_3$.
- **AT6.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow Tv \sqsubseteq Tv'$.
- **AT7** $Tv_1 \sqsubset Tv_2 \Leftrightarrow [Tv_1 \sqsubseteq Tv_2 \wedge \neg\mathsf{Equiv}(Tv_1, Tv_2)]$
- **AT8.** $[Tv_1 \supset\subset Tv_2] \Leftrightarrow [Tv_1 \prec Tv_2 \;\; \wedge \;\; \neg\exists Tv_3[Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2]]$.
- **AT9.** $[Tv_1 \prec Tv_2 \wedge Tv_3 \sqsubseteq Tv_1] \Rightarrow Tv_3 \prec Tv_2$.
- **AT10.** $[Tv_1 \prec Tv_2 \wedge Tv_3 \sqsubseteq Tv_2] \Rightarrow Tv_1 \prec Tv_3$.
- **AT11.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_4 \wedge Tv_4 \prec Tv_3] \Rightarrow Tv_4 \sqsubseteq Tv_1$.

A.1.2 Theorems

- **TT1.** $Tv_1 \prec Tv_2 \Rightarrow [\neg[Tv_2 \sqsubseteq Tv_1] \wedge \neg[Tv_1 \sqsubseteq Tv_2]]$

- **TT2.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow [Tv_2 \sqsubset Tv_1 \wedge Tv_3 \sqsubset Tv_1]$

- **TT3.** $Tv \supset\subset Tv' \Rightarrow \neg[Tv' \supset\subset Tv]$.

- **TT4.** $[Tv_1 \supset\subset Tv_2 \wedge Tv_2 \supset\subset Tv_3] \Rightarrow \neg[Tv_1 \supset\subset Tv_3]$.

- **TT5.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow \neg[Tv \supset\subset Tv']$.

A.2 States

A.2.1 General

Axioms

- **AS1.** $[\forall Tv' \; [Tv' \sqsubset Tv \Rightarrow \mathsf{Holds}(Sv, Tv')]] \Rightarrow \mathsf{Holds}(Sv, Tv)$.

- **AS2.** $\mathsf{Holds}(Sv, Tv) \Rightarrow \; \forall Tv'[Tv' \sqsubset Tv \Rightarrow \mathsf{Holds}(Sv, Tv')]$.

- **AS3.** $\mathsf{MHolds}(Sv, Tv) \Rightarrow [\mathsf{Holds}(Sv, Tv) \wedge \neg\exists Tv'[\mathsf{Holds}(Sv, Tv') \wedge Tv \sqsubset Tv']]$

- **AS4.** $\mathsf{Holds}(\dot{\neg}\, Sv, Tv) \Leftrightarrow \forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \neg\mathsf{Holds}(Sv, Tv')]$

- **AS5.** $\mathsf{Holds}(Sv_1 \,\dot{\wedge}\, Sv_2, Tv) \Leftrightarrow [\mathsf{Holds}(Sv_1, Tv) \wedge \mathsf{Holds}(Sv_2, Tv)]$

- **AS6.** $\mathsf{Holds}(Sv_1 \,\dot{\vee}\, Sv_2, Tv) \Leftrightarrow \mathsf{Holds}(\dot{\neg}\, [\dot{\neg}\, Sv_1 \,\dot{\wedge}\dot{\neg}\, Sv_2], Tv)$

- **AS7.** $\mathsf{Holds}(\dot{\forall}\, x[Sv], Tv) \Leftrightarrow \forall x[\mathsf{Holds}(Sv, Tv)]$

- **AS8.** $\mathsf{Holds}(\dot{\exists}\, x[Sv], Tv) \Leftrightarrow \forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge \exists x[\mathsf{Holds}(Sv, Tv'')]]]$

Theorems

- **TS1.** $\mathsf{Holds}(\dot{\neg}\, Sv, Tv) \Rightarrow \neg\mathsf{Holds}(Sv, Tv)$

- **TS2.** $\mathsf{Holds}(Sv, Tv) \Leftrightarrow \mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tv)$

- **TS3.** $\neg\mathsf{Holds}(Sv, Tv) \Leftrightarrow \exists Tv'[Tv' \sqsubseteq Tv \wedge \mathsf{Holds}(\dot{\neg}\, Sv, Tv')]$

- **TS4.** $\mathsf{Holds}(Sv_1 \dot{\vee} Sv_2, Tv) \Leftrightarrow$
$$\forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]$$

- **TEES.** $\exists Tv\mathsf{Holds}(Sv_1^e \dot{\wedge} Sv_2^e, Tv) \Rightarrow \forall Tv\mathsf{Holds}(Sv_1^e \dot{\wedge} Sv_2^e, Tv)$

- **TEPS.** $\forall Tv[\mathsf{Holds}(Sv_1^e \dot{\wedge} Sv_2^p, Tv) \Rightarrow$
$$\forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(Sv_1^e \dot{\wedge} Sv_2^p, Tv')]]$$

- **TPPS.** $\forall Tv[\mathsf{Holds}(Sv_1^p \dot{\wedge} Sv_2^p, Tv) \Rightarrow$
$$\forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(Sv_1^p \dot{\wedge} Sv_2^p, Tv')]]$$

- **TEcPS.** $\forall Tv[\mathsf{Holds}(Sv_1^e \dot{\wedge} Sv_2^{cp}, Tv) \Rightarrow$
$$\forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(Sv_1^e \dot{\wedge} Sv_2^{cp}, Tv')]]$$

- **TcPcPS.** $\forall Tv[\mathsf{Holds}(Sv_1^{cp} \dot{\wedge} Sv_2^{cp}, Tv) \Rightarrow$
$$\forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(Sv_1^{cp} \dot{\wedge} Sv_2^{cp}, Tv')]]$$

A.2.2 Eternal States

Axioms

- **AES1.** $\exists Tv[\mathsf{Holds}(S^e v, Tv)] \Rightarrow \forall Tv[\mathsf{Holds}(S^e v, Tv)]$.

Theorems

- **TES3.** $\exists Tv\mathsf{Holds}(\dot{\neg}\, Sv^e, Tv) \Rightarrow \forall Tv\mathsf{Holds}(\dot{\neg}\, Sv^e, Tv)$

- **TES4.** $\mathsf{Holds}(\dot{\neg}\, S^e v, Tv) \Leftrightarrow \neg\mathsf{Holds}(S^e v, Tv)$

A.2.3 Permanent States

Axioms

- **APS1.** $\mathsf{Holds}(S^p v, Tv) \Rightarrow \forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(S^p v, Tv')]$

Theorems

- **TPS1.** $\forall Tv[\mathsf{Holds}(\dot{\neg}\, Sv^p, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]]$

A.2.4 Co-Permanent States

Axioms

- **AcPS1.** $\forall Tv[\mathsf{Holds}(Sv^{cp}v, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(Sv^{cp}v, Tv')]]$

Theorems

- **TcPS1.** $\forall Tv[\mathsf{Holds}(\dot{\neg}\, S^{cp}, Tv) \Rightarrow \forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(\dot{\neg}\, S^{cp}, Tv')]]$

A.3 Events

A.3.1 General

Axioms

- **AE1.** $[\mathsf{Occurs}(Ev, Tv_1) \wedge \mathsf{Occurs}(Ev, Tv_2)] \Rightarrow \mathsf{Equiv}(Tv_1, Tv_2)$

- **AE2.** $[\mathsf{Occurs}(Ev_1, Tv) \wedge \mathsf{Equiv}(Ev_1, Ev_2)] \Rightarrow \mathsf{Occurs}(Ev_2, Tv)$.

- **AE3.** $[\mathsf{Cat}(Ev_1, ECv) \wedge \mathsf{Equiv}(Ev_1, Ev_2)] \Rightarrow \mathsf{Cat}(Ev_2, ECv)$.

- **AE4.** $\mathsf{Occurs}(\mathsf{Clos}(\dot{E}v_1, \dot{E}v_2), Tv) \Leftrightarrow$
$$\exists S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, S^t v)$$
$$\wedge \mathsf{Occurs}(\dot{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\dot{E}v_2, Tv_2)$$
$$\wedge \mathsf{Covers}(Tv, Tv_1, Tv_2)]$$

- **AE5.** $\forall\, \overline{\dot{E}\dot{v}}\, [\exists\, \dot{E}v_1, \dot{E}v_2\, [\mathsf{Equiv}(\overline{\dot{E}\dot{v}}, \mathsf{Clos}(\dot{E}v_1, \dot{E}v_2))]]$

- **AE6.** $\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, S^t v) \Leftrightarrow \mathsf{Cat}(\mathsf{Clos}(\dot{E}v_1, \dot{E}v_2), \mathsf{PO}(S^t v))$

- **AE7.** $[\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, S^t v) \wedge \mathsf{Cat}(\mathsf{Clos}(\dot{E}v_1, \dot{E}v_2), \overline{\dot{E}C\dot{v}})] \Rightarrow$
$$\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\dot{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\dot{E}v_2, Tv_2)$$

$$Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow$$

$$\wedge \mathsf{Holds}(\mathsf{Prog}(\overset{\bullet-\bullet}{ECv}), Tv_3)]$$

- **AE8.** $\mathsf{Holds}(\mathsf{Complete}(Ev), Tv) \Rightarrow \exists Tv'[Tv' \prec Tv \wedge \mathsf{Occurs}(Ev, Tv')]$

- **AE9.** $\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev})))$.

- **AE10.** $\forall \overset{\bullet-\bullet}{Ev}\ [\exists \overset{\bullet}{E}v_1, Tv_1[\mathsf{Cat}(\overset{\bullet}{E}v_1, \uparrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1)] \Leftrightarrow$

  $\exists \overset{\bullet}{E}v_2, Tv_2[\mathsf{Cat}(\overset{\bullet}{E}v_2, \downarrow \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)]]$

- **AE11.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev_1}, \mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev}))) \wedge \mathsf{Cat}(\overset{\bullet-\bullet}{Ev_2}, \mathsf{PO}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev})))$

  $\wedge \mathsf{Occurs}(\overset{\bullet-\bullet}{Ev_1}, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet-\bullet}{Ev_2}, Tv_2)] \Rightarrow$

  $\mathsf{Equiv}(\overset{\bullet-\bullet}{Ev_1}, \overset{\bullet-\bullet}{Ev_2})$

- **AE12.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \mathsf{Holds}(\mathsf{Int}(\overset{\bullet-\bullet}{Ev}), Tv)] \Rightarrow \mathsf{Holds}(\mathsf{Prog}(\overset{\bullet-\bullet}{ECv}), Tv)$

- **AE13.** $[\mathsf{Occurs}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), Tv) \wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev}))] \Rightarrow$

  $[\mathsf{Equiv}(\overset{\bullet-\bullet}{Ev}, \mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2)) \Leftrightarrow \exists Tv'[\mathsf{Holds}(\mathsf{Complete}(\overset{\bullet-\bullet}{Ev}), Tv')]]$

- **AE14.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overset{\rightarrow\bullet}{\mathsf{telic}}\ (\overset{\bullet-\bullet}{ECv}, Sv)$

  $\wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow Sv)$

  $\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_2)$

  $\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_2)] \Rightarrow$

  $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{ECv})$

- **AE15.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overset{\rightarrow\ldots\bullet}{\mathsf{telic}}\ (\overset{\bullet-\bullet}{ECv}, Sv)$

  $\wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow Sv)$

  $\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3)$

  $\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_3) \wedge [Tv_2 \prec Tv_3 \vee \mathsf{Equiv}(Tv_2, Tv_3)]] \Rightarrow$

  $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{ECv})$

- **AE16.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{ECv}) \wedge \overset{\longleftarrow}{\mathsf{atelic}}\ (\overset{\bullet-\bullet}{ECv}, Sv)$

  $\wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev}))$

  $\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$

  $\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_2)] \Rightarrow$

  $\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{ECv})$

- **AE17.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{EC}v) \wedge \overrightarrow{\mathsf{atelic}}\ (\overset{\bullet-\bullet}{EC}v, Sv)$

$$\wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev})) \wedge \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow Sv)$$

$$\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3)$$

$$\wedge \mathsf{NoOcc}(\uparrow Sv, Tv_1, Tv_3) \wedge [Tv_3 \prec Tv_2 \vee \mathsf{Equiv}(Tv_2, Tv_3)]] \Rightarrow$$

$$\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{EC}v)$$

- **AE18.** $[\mathsf{Cat}(\overset{\bullet-\bullet}{Ev}, \overset{\bullet-\bullet}{EC}v) \wedge \overleftrightarrow{\mathsf{atelic}}\ (\overset{\bullet-\bullet}{EC}v) \wedge \mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, \mathsf{Int}(\overset{\bullet-\bullet}{Ev}))] \Rightarrow$

$$\mathsf{Cat}(\mathsf{Clos}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2), \overset{\bullet-\bullet}{EC}v)$$

Theorems

- **TE1.** $\mathsf{Occurs}(Ev, Tv) \Rightarrow \forall Tv'[Tv' \sqsubset Tv \Rightarrow \neg\mathsf{Occurs}(Ev, Tv')]$.

- **TE2.** $\mathsf{Occurs}(Ev, Tv) \Rightarrow \forall Tv'[Tv \sqsubset Tv' \Rightarrow \neg\mathsf{Occurs}(Ev, Tv')]$.

- **TE3.** $\exists \overset{\bullet-\bullet}{Ev}\ [\mathsf{Occurs}(\overset{\bullet-\bullet}{Ev}, Tv)] \Leftrightarrow$

$$\exists\ \overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, S^t v)$$

$$\wedge \mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$$

$$\wedge \mathsf{Covers}(Tv, Tv_1, Tv_2)]$$

- **TE4.** $\mathsf{OCPair}(\overset{\bullet}{E}v_1, \overset{\bullet}{E}v_2, S^t v) \Rightarrow$

$$\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\overset{\bullet}{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{E}v_2, Tv_2)$$

$$Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow$$

$$\wedge \mathsf{Holds}(\mathsf{Prog}(\mathsf{PO}(S^t v)), Tv_3)]$$

A.3.2 Onsets and Cessations

Axioms

- **AOC1.** $\mathsf{Cat}(\overset{\bullet}{E}v, \uparrow Sv) \Leftrightarrow \mathsf{Cat}(\overset{\bullet}{E}v, \downarrow(\neg Sv))$.

- **AOC2.** $\mathsf{Cat}(\overset{\bullet}{E}v, \downarrow Sv) \Leftrightarrow \mathsf{Cat}(\overset{\bullet}{E}v, \uparrow(\neg Sv))$.

- **AOC3.** $[\mathsf{Occurs}(\overset{\bullet}{E}v, Tv)\ \wedge\ \mathsf{Cat}(\overset{\bullet}{E}v, \uparrow Sv)]\ \Rightarrow$

$$\exists Tv'[Tv \supset\subset Tv' \wedge \mathsf{Holds}(Sv, Tv')]$$

- **AOC4.** $[\mathsf{Occurs}(\dot{E}v,Tv) \ \wedge \ \mathsf{Cat}(\dot{E}v,\uparrow Sv)] \ \Rightarrow$
$$\exists Tv'[Tv' \supset\subset Tv \wedge \mathsf{Holds}(\neg Sv,Tv')]$$

- **AOC5.** $[\neg\mathsf{Holds}(Sv,Tv_1) \ \wedge \ \mathsf{Holds}(Sv,Tv_2) \ \wedge \ Tv_1 \prec Tv_2] \ \Rightarrow$
$$\exists \dot{E}v, Tv_3[\mathsf{Cat}(\dot{E}v,\uparrow Sv) \ \wedge \ \mathsf{Occurs}(\dot{E}v,Tv_3)$$
$$\wedge \ Tv_1 \prec Tv_3 \ \wedge \ Tv_3 \prec Tv_2].$$

Theorems

- **TOC1.** $\mathsf{Cat}(\dot{E}v,\uparrow Sv) \Leftrightarrow \mathsf{Cat}(\dot{E}v,\uparrow(\neg\neg Sv)).$

- **TOC2.** $\mathsf{Cat}(\dot{E}v,\downarrow Sv) \Leftrightarrow \mathsf{Cat}(\dot{E}v,\downarrow(\neg\neg Sv)).$

- **TOC3.** $[\mathsf{Occurs}(\dot{E}v,Tv) \ \wedge \ \mathsf{Cat}(\dot{E}v,\downarrow Sv)] \ \Rightarrow$
$$\exists Tv'[Tv' \supset\subset Tv \wedge \mathsf{Holds}(Sv,Tv')]$$

- **TOC4.** $[\mathsf{Occurs}(\dot{E}v,Tv) \ \wedge \ \mathsf{Cat}(\dot{E}v,\downarrow Sv)] \ \Rightarrow$
$$\exists Tv[Tv \supset\subset Tv' \wedge \mathsf{Holds}(\neg Sv,Tv')]$$

- **TOC5.** $[\mathsf{Holds}(Sv,Tv_1) \ \wedge \ \neg\mathsf{Holds}(Sv,Tv_2) \ \wedge \ Tv_1 \prec Tv_2] \ \Rightarrow$
$$\exists \dot{E}v, Tv_3[\mathsf{Cat}(\dot{E}v,\downarrow Sv) \ \wedge \ \mathsf{Occurs}(\dot{E}v,Tv_3)$$
$$\wedge \ Tv_1 \prec Tv_3 \ \wedge \ Tv_3 \prec Tv_2].$$

- **TOC6.** $[\mathsf{Occurs}(\dot{E}v_1,Tv_1) \ \wedge \ \mathsf{Cat}(\dot{E}v_1,\uparrow S^p v)] \ \Rightarrow$
$$\forall \dot{E}v_2, Tv_2[[Tv_1 \prec Tv_2 \ \wedge \ \mathsf{Occurs}(\dot{E}v_2,Tv_2)] \ \Rightarrow \ \neg\mathsf{Cat}(\dot{E}v_2,\uparrow S^p v)]$$

- **TOC7.** $[\mathsf{Occurs}(\dot{E}v_1,Tv_1) \ \wedge \ \mathsf{Cat}(\dot{E}v_1,\downarrow S^{cp} v)] \ \Rightarrow$
$$\forall \dot{E}v_2, Tv_2[[Tv_1 \prec Tv_2 \ \wedge \ \mathsf{Occurs}(\dot{E}v_2,Tv_2)] \ \Rightarrow \ \neg\mathsf{Cat}(\dot{E}v_2,\downarrow S^{cp} v)]$$

- **TOC8.** $[\mathsf{Occurs}(\dot{E}v_1,Tv_1) \wedge \mathsf{Cat}(\dot{E}v_1,\uparrow S^t v)$
$$\wedge \ \mathsf{Occurs}(\dot{E}v_2,Tv_2) \ \wedge \ \mathsf{Cat}(\dot{E}v_2,\uparrow S^t v) \ \wedge \ Tv_1 \prec Tv_2] \Rightarrow$$
$$\exists \dot{E}v_3, Tv_3[\mathsf{Occurs}(\dot{E}v_3,Tv_3) \ \wedge \ \mathsf{Cat}(\dot{E}v_3,\downarrow S^t v)$$
$$\wedge \ Tv_1 \prec Tv_3 \ \wedge \ Tv_3 \prec Tv_2]$$

- **TOC9.** $[\mathsf{Occurs}(\dot{E}v_1,Tv_1) \wedge \mathsf{Cat}(\dot{E}v_1,\downarrow S^t v)$
$$\wedge \ \mathsf{Occurs}(\dot{E}v_2,Tv_2) \ \wedge \ \mathsf{Cat}(\dot{E}v_2,\downarrow S^t v) \ \wedge \ Tv_1 \prec Tv_2] \Rightarrow$$

$$\exists\, \overset{\bullet}{E}v_3, Tv_3[\mathsf{Occurs}(\overset{\bullet}{E}v_3, Tv_3) \;\wedge\; \mathsf{Cat}(\overset{\bullet}{E}v_3, \uparrow S^t v)$$

$$\wedge\, Tv_1 \prec Tv_3 \;\wedge\; Tv_3 \prec Tv_2]$$

Appendix B

Proofs of Theorems from Chapter 3

Following are the proofs of a selected subset of the theorems appearing in Chapter 3. Unless the particular rules used in moving from one step of the proof to the next are not obvious, I use "FOPC" to refer the use of the standard introduction and elimination rules of inference of first-order predicate calculus.

Proof of TT1

I break the proof of **TT1** down by proving two simpler lemmas.

- **LT1.** $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_2 \sqsubseteq Tv_1]$

Proof.

1. $Tv_1 \prec Tv_2$ (Assumption)
2. $Tv_2 \sqsubseteq Tv_1$ (Assumption)
3. $Tv_1 \prec Tv_2 \wedge Tv_2 \sqsubseteq Tv_1$ (1, 2, $\wedge$-introduction)
4. $[Tv_1 \prec Tv_2 \wedge Tv_2 \sqsubseteq Tv_1] \Rightarrow Tv_2 \prec Tv_2$ (**AT9**, $\forall$-elimination)
5. $Tv_2 \prec Tv_2$ (3, 4, $\Rightarrow$-elimination)
6. $\mathsf{Equiv}(Tv_2, Tv_2) \Rightarrow \neg[Tv_2 \prec Tv_2]$ (**AT3**, $\forall$-elimination)

7. $\neg[Tv_2 \prec Tv_2]$ (6, reflexivity of Equiv, $\Rightarrow$-elimination)

8. $\neg[Tv_2 \sqsubseteq Tv_1]$ (2, 5, 7, $\neg$-introduction)

9. $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_2 \sqsubseteq Tv_1]$ (1, 8, $\Rightarrow$-introduction)

Q.E.D.

- **LT2.** $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_1 \sqsubseteq Tv_2]$

Proof.

1. $Tv_1 \prec Tv_2$ (Assumption)

2. $Tv_1 \sqsubseteq Tv_2$ (Assumption)

3. $Tv_1 \prec Tv_2 \wedge Tv_1 \sqsubseteq Tv_2$ (1, 2, $\wedge$-introduction)

4. $[Tv_1 \prec Tv_2 \wedge Tv_1 \sqsubseteq Tv_2] \Rightarrow Tv_1 \prec Tv_1$ (**AT10**, $\forall$-elimination)

5. $Tv_1 \prec Tv_1$ (3, 4, $\Rightarrow$-elimination)

6. $\mathsf{Equiv}(Tv_1, Tv_1) \Rightarrow \neg[Tv_1 \prec Tv_1]$ (**AT3**, $\forall$-elimination)

7. $\neg[Tv_1 \prec Tv_1]$ (6, reflexivity of Equiv, $\Rightarrow$-elimination)

8. $\neg[Tv_1 \sqsubseteq Tv_2]$ (2, 5, 7, $\neg$-introduction)

9. $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_1 \sqsubseteq Tv_2]$ (1, 8, $\Rightarrow$-introduction)

Q.E.D.

Given **LT1** and **LT2**, **TT1** readily follows.

Proof of TT2

Similar to the proof of **TT1**, we prove **TT2** by proving the two lemmas **LT3** and **LT4**.

- **LT3.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_2 \sqsubset Tv_1$

  Proof.

1. $Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3$ (Assumption)
2. $\neg[Tv_2 \sqsubset Tv_1]$ (Assumption)
3. $Tv_2 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
4. $\mathsf{Equiv}(Tv_2, Tv_1)$ (Follows from 2, 3, and **AT7**)
5. $Tv_1 \sqsubseteq Tv_2$ (4, **AT6**, $\Rightarrow$-elimination)
6. $Tv_3 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
7. $Tv_3 \sqsubseteq Tv_1 \wedge Tv_1 \sqsubseteq Tv_2$ (5, 6, $\wedge$-introduction)
8. $Tv_3 \sqsubseteq Tv_2$ (7, **AT5**, $\Rightarrow$-elimination)
9. $Tv_2 \prec Tv_3$ (1, $\wedge$-elimination)
10. $\neg[Tv_3 \sqsubseteq Tv_2]$ (9, **LT1**, $\Rightarrow$-elimination)
11. $Tv_2 \sqsubset Tv_1$ (2, 8, 10, $\neg$-elimination)
12. $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_2 \sqsubset Tv_1$ (1, 11, $\Rightarrow$-introduction)

Q.E.D.

- **LT4.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_3 \sqsubset Tv_1$

  Proof.

  1. $Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3$ (Assumption)
  2. $\neg[Tv_3 \sqsubset Tv_1]$ (Assumption)
  3. $Tv_3 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
  4. $\mathsf{Equiv}(Tv_3, Tv_1)$ (Follows from 2, 3, and **AT7**)
  5. $Tv_1 \sqsubseteq Tv_3$ (4, **AT6**, $\Rightarrow$-elimination)
  6. $Tv_2 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
  7. $Tv_2 \sqsubseteq Tv_1 \wedge Tv_1 \sqsubseteq Tv_3$ (5, 6, $\wedge$-introduction)
  8. $Tv_2 \sqsubseteq Tv_3$ (7, **AT5**, $\Rightarrow$-elimination)
  9. $Tv_2 \prec Tv_3$ (1, $\wedge$-elimination)

10. $\neg[Tv_2 \sqsubseteq Tv_3]$ (9, **LT2**, $\Rightarrow$-elimination)

11. $Tv_3 \sqsubset Tv_1$ (2, 8, 10, $\neg$-elimination)

12. $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_3 \sqsubset Tv_1$

(1, 11, $\Rightarrow$-introduction)

Q.E.D.

TT2 follows directly from **LT3** and **LT4**.

Proof of TS2

- $\mathsf{Holds}(Sv, Tv) \Leftrightarrow \mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tv)$

$\Rightarrow$

1. $\neg\mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tc)$ (Assumption)
2. $\neg\forall Tv[Tv \sqsubseteq Tc \Rightarrow \neg\mathsf{Holds}(\dot{\neg}\, Sv, Tv)]$ (1, **AS4**, FOPC)
3. $\neg\forall Tv[Tv \sqsubseteq Tc \Rightarrow \neg\forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \neg\mathsf{Holds}(Sv, Tv')]]$ (2, **AS4**, FOPC)
4. $\exists Tv[Tv \sqsubseteq Tc \wedge \forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \neg\mathsf{Holds}(Sv, Tv')]]$ (3, FOPC)
5. $\exists Tv[Tv \sqsubseteq Tc \wedge \neg\mathsf{Holds}(Sv, Tv)]$ (4, **AT6**, FOPC)
6. $\neg\forall Tv[Tv \sqsubseteq Tc \Rightarrow \neg\mathsf{Holds}(Sv, Tv)]$ (5, FOPC)
7. $\neg\mathsf{Holds}(Sv, Tc)$ (6, **AS2**, FOPC)
8. $\neg\mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tc) \Rightarrow \neg\mathsf{Holds}(Sv, Tc)$ (1, 7, FOPC)
9. $\mathsf{Holds}(Sv, Tc) \Rightarrow \mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tc)$ (8, FOPC)
10. $\forall Tv[\mathsf{Holds}(Sv, Tv) \Rightarrow \mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tv)]$ (9, $\forall$-introduction)

$\Leftarrow$

1. $\mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv, Tc)$ (Assumption)
2. $\neg\mathsf{Holds}(\dot{\neg}\, Sv, Tc)$ (1, **TS1**, FOPC)
3. $\neg\neg\mathsf{Holds}(Sv, Tc)$ (2, **TS1**, FOPC)
4. $\mathsf{Holds}(Sv, Tc)$ (3, FOPC)

5. $\mathsf{Holds}(\overset{\bullet}{\neg}\overset{\bullet}{\neg}\, Sv, Tc) \Rightarrow \mathsf{Holds}(Sv, Tc)$ (1, 4, FOPC)

6. $\forall Tv[\mathsf{Holds}(\overset{\bullet}{\neg}\overset{\bullet}{\neg}\, Sv, Tv) \Rightarrow \mathsf{Holds}(Sv, Tv)]$ (5, $\forall$-introduction)

Q.E.D.

Proof of TS3

- $\neg\mathsf{Holds}(Sv, Tv) \Leftrightarrow \exists Tv'[Tv' \sqsubseteq Tv \wedge \mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')]$

$\Rightarrow$

1. $\neg\mathsf{Holds}(Sv, Tc)$ (Assumption)
2. $\neg\mathsf{Holds}(\overset{\bullet}{\neg}\overset{\bullet}{\neg}\, Sv, Tc)$ (1, **TS2**, FOPC)
3. $\neg\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \neg\mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')]$ (2, **AS4**, FOPC)
4. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')]$ (3, FOPC)
5. $\neg\mathsf{Holds}(Sv, Tc) \Rightarrow \exists Tv'[Tv' \sqsubseteq Tc \wedge \mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')]$ (1, 4, FOPC)
6. $\forall Tv[\neg\mathsf{Holds}(Sv, Tv) \Rightarrow \exists Tv'[Tv' \sqsubseteq Tv \wedge \mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')]]$ (5, $\forall$-introduction)

$\Leftarrow$

1. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')]$ (Assumption)
2. $\neg\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \neg\mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')]$ (1, **AS4**, FOPC)
3. $\neg\mathsf{Holds}(\overset{\bullet}{\neg}\overset{\bullet}{\neg}\, Sv, Tc)$ (2, **TS2**, FOPC)
4. $\neg\mathsf{Holds}(Sv, Tc)$ (3, **TS2**, FOPC)
5. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')] \Rightarrow \neg\mathsf{Holds}(Sv, Tc)$ (1, 4, FOPC)
6. $\forall Tv[\exists Tv'[Tv' \sqsubseteq Tv \wedge \mathsf{Holds}(\overset{\bullet}{\neg}\, Sv, Tv')] \Rightarrow \neg\mathsf{Holds}(Sv, Tv)]$ (5, $\forall$-introduction)

Q.E.D.

Proof of TES3

- $\exists Tv\mathsf{Holds}(\overset{\bullet}{\neg}\, Sv^e, Tv) \Rightarrow \forall Tv\mathsf{Holds}(\overset{\bullet}{\neg}\, Sv^e, Tv)$

1. $\exists Tv \mathsf{Holds}(\dot{\neg}\, Sv^e, Tv)$ (Assumption)
2. $\exists Tv \neg \mathsf{Holds}(Sv^e, Tv)$ (1, **TS1**, FOPC)
3. $\neg \forall Tv \mathsf{Holds}(Sv^e, Tv)$ (2, FOPC)
4. $\neg \forall Tv \mathsf{Holds}(\dot{\neg}\, Sv^e, Tv)$ (Assumption)
5. $\exists Tv \neg \mathsf{Holds}(\dot{\neg}\, Sv^e, Tv)$ (4, FOPC)
6. $\exists Tv[\exists Tv'[Tv' \sqsubseteq Tv \wedge \mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv^e, Tv')]]$ (5, **TS3**, FOPC)
7. $\exists Tv[\exists Tv'[Tv' \sqsubseteq Tv \wedge \mathsf{Holds}(Sv^e, Tv')]]$ (6, **TS2**, FOPC)
8. $\exists Tv \mathsf{Holds}(Sv^e, Tv)$ (7, FOPC)
9. $\forall Tv \mathsf{Holds}(Sv^e, Tv)$ (8, **AES1**, FOPC)
10. $\forall Tv \mathsf{Holds}(\dot{\neg}\, Sv^e, Tv)$ (3, 4, 9, FOPC)
11. $\exists Tv \mathsf{Holds}(\dot{\neg}\, Sv^e, Tv) \Rightarrow \forall Tv \mathsf{Holds}(\dot{\neg}\, Sv^e, Tv)$ (1, 10, FOPC)

Q.E.D.

Proof of TPS1

- $\forall Tv[\mathsf{Holds}(\dot{\neg}\, Sv^p, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]]$

1. $\mathsf{Holds}(\dot{\neg}\, Sv^p, Tc)$ (Assumption)
2. $\neg \mathsf{Holds}(Sv^p, Tc)$ (1, **TS1**, FOPC)
3. $\neg \forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]$ (Assumption)
4. $\exists Tv'[Tv' \prec Tc \wedge \neg \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]$ (3, FOPC)
5. $\exists Tv'[Tv' \prec Tc \wedge \exists Tv''[Tv'' \sqsubseteq Tv' \wedge \mathsf{Holds}(\dot{\neg}\dot{\neg}\, Sv^p, Tv'')]]$ (4, **TS3**, FOPC)
6. $\exists Tv'[Tv' \prec Tc \wedge \exists Tv''[Tv'' \sqsubseteq Tv' \wedge \mathsf{Holds}(Sv^p, Tv'')]]$ (5, **TS2**, FOPC)
7. $\exists Tv''[Tv'' \prec Tc \wedge \mathsf{Holds}(Sv^p, Tv'')]$ (6, **AT9**, FOPC)
8. $Tc' \prec Tc \wedge \mathsf{Holds}(Sv^p, Tc')]$ (7, $\exists$-elimination)
9. $Tc' \prec Tc$ (8, FOPC)
10. $\mathsf{Holds}(Sv^p, Tc')$ (8, FOPC)

11. $\forall Tv[Tc' \prec Tv \Rightarrow \mathsf{Holds}(Sv^p, Tv)]$ (10, **APS1**, FOPC)

12. $\mathsf{Holds}(Sv^p, Tc)$ (9, 11, FOPC)

13. $\forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]$ (2, 3, 12, FOPC)

14. $\mathsf{Holds}(\dot{\neg}\, Sv^p, Tc) \Rightarrow \forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]$ (1, 13, FOPC)

15. $\forall Tv[\mathsf{Holds}(\dot{\neg}\, Sv^p, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(\dot{\neg}\, Sv^p, Tv')]]$ (14, $\forall$-introduction)

Q.E.D.

Proof of TEcPS

- $\forall Tv[\mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tv) \Rightarrow \forall Tv'[Tv \prec Tv' \Rightarrow \mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tv')]]$

1. $\mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tc)$ (Assumption)

2. $\mathsf{Holds}(Sv_1^e, Tc) \wedge \mathsf{Holds}(Sv_1^{cp}, Tc)$ (1, **AS5**, FOPC)

3. $\forall Tv' \mathsf{Holds}(Sv_1^e, Tv') \wedge \mathsf{Holds}(Sv_1^{cp}, Tc)$ (2, **AES1**, FOPC)

4. $\forall Tv' \mathsf{Holds}(Sv_1^e, Tv') \wedge \forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(Sv_2^{cp}, Tv')]$ (3, **AcPS1**, FOPC)

5. $\forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(Sv_1^e, Tv')] \wedge \forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(Sv_2^{cp}, Tv')]$ (4, FOPC)

6. $\forall Tv'[Tv' \prec Tc \Rightarrow [\mathsf{Holds}(Sv_1^e, Tv') \wedge \mathsf{Holds}(Sv_1^{cp}, Tv')]]$ (5, FOPC)

7. $\forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tv')]$ (6, **AS5**, FOPC)

8. $\mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tc) \Rightarrow \forall Tv'[Tv' \prec Tc \Rightarrow \mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tv')]$ (1, 7, FOPC)

9. $\forall Tv[\mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tv) \Rightarrow \forall Tv'[Tv' \prec Tv \Rightarrow \mathsf{Holds}(Sv_1^e \,\dot{\wedge}\, Sv_2^{cp}, Tv')]]$

(8, $\forall$-introduction)

Q.E.D.

Proof of TS4

- **TS4.** $\forall Tv[\mathsf{Holds}(Sv_1 \,\dot{\vee}\, Sv_2, Tv) \Leftrightarrow$

$\forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]]$

$\Rightarrow$

1. $\mathsf{Holds}(Sv_1 \dot{\vee} Sv_2, Tc)$ (Assumption)
2. $\mathsf{Holds}(\dot{\neg}\,[\dot{\neg}\, Sv_1 \dot{\wedge} \dot{\neg}\, Sv_2], Tc)$ (1, **AS6**, FOPC)
3. $\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \neg\mathsf{Holds}(\dot{\neg}\, Sv_1 \dot{\wedge} \dot{\neg}\, Sv_2, Tv')]$ (2, **AS4**, FOPC)
4. $\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \neg[\mathsf{Holds}(\dot{\neg}\, Sv_1, Tv') \wedge \mathsf{Holds}(\dot{\neg}\, Sv_2, Tv')]]$ (3, **AS5**, FOPC)
5. $\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \neg\forall Tv''[Tv'' \sqsubseteq Tv' \Rightarrow [\neg\mathsf{Holds}(Sv_1, Tv'') \wedge \neg\mathsf{Holds}(Sv_2, Tv'')]]]$

   (4, **AS4**, FOPC)
6. $\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge \neg[\neg\mathsf{Holds}(Sv_1, Tv'') \wedge \neg\mathsf{Holds}(Sv_2, Tv'')]]]$

   (5, FOPC)
7. $\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]$

   (6, FOPC)
8. $\mathsf{Holds}(Sv_1 \dot{\vee} Sv_2, Tc) \Rightarrow$

   $\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]$

   (1, 7, FOPC)
9. $\forall Tv[\mathsf{Holds}(Sv_1 \dot{\vee} Sv_2, Tv) \Rightarrow$

   $\forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]]$

   (8, $\forall$-introduction)

$\Leftarrow$

1. $\neg\mathsf{Holds}(Sv_1 \dot{\vee} Sv_2, Tc)$ (Assumption)
2. $\neg\mathsf{Holds}(\dot{\neg}\,[\dot{\neg}\, Sv_1 \dot{\wedge} \dot{\neg}\, Sv_2], Tc)$ (1, **AS6**, FOPC)
3. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \mathsf{Holds}(\dot{\neg}\dot{\neg}\,[\dot{\neg}\, Sv_1 \dot{\wedge} \dot{\neg}\, Sv_2], Tv')]$ (2, **TS3**, FOPC)
4. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \mathsf{Holds}(\dot{\neg}\, Sv_1 \dot{\wedge} \dot{\neg}\, Sv_2, Tv')]$ (3, **TS2**, FOPC)
5. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \mathsf{Holds}(\dot{\neg}\, Sv_1, Tv') \wedge \mathsf{Holds}(\dot{\neg}\, Sv_2, Tv')]$ (4, **AS5**, FOPC)
6. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \forall Tv''[Tv'' \sqsubseteq Tv' \Rightarrow [\neg\mathsf{Holds}(Sv_1, Tv'') \wedge \neg\mathsf{Holds}(Sv_2, Tv'')]]]$

   (5, **AS4**, FOPC)
7. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \forall Tv''[Tv'' \sqsubseteq Tv' \Rightarrow \neg[\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]$

   (6, FOPC)

8. $\exists Tv'[Tv' \sqsubseteq Tc \wedge \neg\exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]$

(7, FOPC)

9. $\neg\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow \neg\exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]]$

(8, FOPC)

10. $\forall Tv'[Tv' \sqsubseteq Tc \Rightarrow$

$$\neg\exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]] \Rightarrow$$

$$\mathsf{Holds}(Sv_1 \overset{\bullet}{\vee} Sv_2, Tc)$$

(1, 9, FOPC)

11. $\forall Tv[\forall Tv'[Tv' \sqsubseteq Tv \Rightarrow$

$$\neg\exists Tv''[Tv'' \sqsubseteq Tv' \wedge [\mathsf{Holds}(Sv_1, Tv'') \vee \mathsf{Holds}(Sv_2, Tv'')]]] \Rightarrow$$

$$\mathsf{Holds}(Sv_1 \overset{\bullet}{\vee} Sv_2, Tv)]$$

(10, $\forall$-introduction)

Q.E.D.

Appendix C

Proofs of Theorems from Chapter 4

Proof of TOC3

- $\forall Sv, \overset{\bullet}{E}v, Tv[\mathsf{Occurs}(\overset{\bullet}{E}v, Tv) \ \wedge \ \mathsf{Cat}(\overset{\bullet}{E}v, \downarrow Sv)] \ \Rightarrow \exists Tv'[Tv' \supset\subset Tv \wedge \mathsf{Holds}(Sv, Tv')]$

1. $\mathsf{Occurs}(\overset{\bullet}{E}c, Tc) \ \wedge \ \mathsf{Cat}(\overset{\bullet}{E}c, \downarrow Sc)$ (Assumption)
2. $\mathsf{Occurs}(\overset{\bullet}{E}c, Tc) \ \wedge \ \mathsf{Cat}(\overset{\bullet}{E}c, \uparrow\neg Sc)$ (1, **AOC2**, FOPC)
3. $\exists Tv'[Tv' \supset\subset Tc \wedge \mathsf{Holds}(\neg\neg Sc, Tv')]$ (2, **AOC4**, FOPC)
4. $\exists Tv'[Tv' \supset\subset Tc \wedge \mathsf{Holds}(Sc, Tv')]$ (3, **TS2**, FOPC)
5. $[\mathsf{Occurs}(\overset{\bullet}{E}c, Tc) \ \wedge \ \mathsf{Cat}(\overset{\bullet}{E}v, \downarrow Sc)] \Rightarrow \exists Tv'[Tv' \supset\subset Tc \wedge \mathsf{Holds}(Sc, Tv')]$ (1, 4, FOPC)
6. $\forall Sv, \overset{\bullet}{E}v, Tv[\mathsf{Occurs}(\overset{\bullet}{E}v, Tv) \ \wedge \ \mathsf{Cat}(\overset{\bullet}{E}v, \downarrow Sv)] \ \Rightarrow \exists Tv'[Tv' \supset\subset Tv \wedge \mathsf{Holds}(Sv, Tv')]$

   (5, $\forall$-introduction)

Q.E.D.

Proof of TOC5

- $\forall Sv, Tv_1, Tv_2[\mathsf{Holds}(Sv, Tv_1) \ \wedge \ \neg\mathsf{Holds}(Sv, Tv_2) \ \wedge \ Tv_1 \prec Tv_2] \ \Rightarrow$
  $\exists \overset{\bullet}{E}v, Tv_3[\mathsf{Cat}(\overset{\bullet}{E}v, \downarrow Sv) \ \wedge \ \mathsf{Occurs}(\overset{\bullet}{E}v, Tv_3)$
  $\wedge \ Tv_1 \prec Tv_3 \ \wedge \ Tv_3 \prec Tv_2]$

1. $\mathsf{Holds}(Sc, Tc_1) \ \wedge\ \neg\mathsf{Holds}(Sc, Tc_2) \ \wedge\ Tc_1 \prec Tc_2$ (Assumption)

2. $\mathsf{Holds}(\neg\neg Sc, Tc_1) \ \wedge\ \neg\mathsf{Holds}(Sc, Tc_2) \ \wedge\ Tc_1 \prec Tc_2$ (1, **TS2**, FOPC)

3. $\neg\mathsf{Holds}(\neg Sc, Tc_1) \ \wedge\ \neg\mathsf{Holds}(Sc, Tc_2) \ \wedge\ Tc_1 \prec Tc_2$ (2, **TS1**, FOPC)

4. $\neg\mathsf{Holds}(\neg Sc, Tc_1) \ \wedge\ \exists Tv[Tv \sqsubseteq Tc_2 \wedge \mathsf{Holds}(\neg Sc, Tv)] \ \wedge\ Tc_1 \prec Tc_2$ (3, **TS3**, FOPC)

5. $\neg\mathsf{Holds}(\neg Sc, Tc_1) \ \wedge\ Tc_3 \sqsubseteq Tc_2 \wedge \mathsf{Holds}(\neg Sc, Tc_3) \ \wedge\ Tc_1 \prec Tc_2$ (4, $\exists$-elimination)

6. $\neg\mathsf{Holds}(\neg Sc, Tc_1) \ \wedge\ \mathsf{Holds}(\neg Sc, Tc_3) \ \wedge\ Tc_1 \prec Tc_3$ (5, **AT9**, FOPC)

7. $\exists \dot{E}v, Tv_3[\mathsf{Cat}(\dot{E}v, \uparrow\neg Sc) \ \wedge\ \mathsf{Occurs}(\dot{E}v, Tv_3) \ \wedge\ Tc_1 \prec Tv_3 \ \wedge\ Tv_3 \prec Tc_3]$

   (6, **AOC5**, FOPC)

8. $\exists \dot{E}v, Tv_3[\mathsf{Cat}(\dot{E}v, \downarrow Sc) \ \wedge\ \mathsf{Occurs}(\dot{E}v, Tv_3) \ \wedge\ Tc_1 \prec Tv_3 \ \wedge\ Tv_3 \prec Tc_3]$

   (7, **AOC2**, FOPC)

9. $[\mathsf{Holds}(Sc,Tc_1) \ \wedge\ \neg\mathsf{Holds}(Sc, Tc_2) \ \wedge\ Tc_1 \prec Tc_2] \ \Rightarrow$
   $\exists \dot{E}v, Tv_3[\mathsf{Cat}(\dot{E}v, \downarrow Sc) \ \wedge\ \mathsf{Occurs}(\dot{E}v, Tv_3)$
   $\wedge\ Tc_1 \prec Tv_3 \ \wedge\ Tv_3 \prec Tc_2]$

   (1, 8, FOPC)

10. $\forall Sv, Tv_1, Tv_2[\mathsf{Holds}(Sv,Tv_1) \ \wedge\ \neg\mathsf{Holds}(Sv, Tv_2) \ \wedge\ Tv_1 \prec Tv_2] \ \Rightarrow$
    $\exists \dot{E}v, Tv_3[\mathsf{Cat}(\dot{E}v, \downarrow Sv) \ \wedge\ \mathsf{Occurs}(\dot{E}v, Tv_3)$
    $\wedge\ Tv_1 \prec Tv_3 \ \wedge\ Tv_3 \prec Tv_2]$

    (9, $\forall$-introduction)

Q.E.D.

Proof of TE3

Below I only prove the $\Rightarrow$ direction. The $\Leftarrow$ direction follows immediately form **AE4** and the fact that Clos forms a durative event.

- $\forall Tv[\exists \dot{\bar{\dot{E}}}v\ [\mathsf{Occurs}(\dot{\bar{\dot{E}}}v, Tv)] \Leftrightarrow$
  $\exists \dot{E}v_1, \dot{E}v_2, S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, S^t v)$
  $\wedge \mathsf{Occurs}(\dot{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\dot{E}v_2, Tv_2)$
  $\wedge \mathsf{Covers}(Tv, Tv_1, Tv_2)]]$

1. $\exists \overset{\bullet-\bullet}{Ev}\, [\mathsf{Occurs}(\overset{\bullet-\bullet}{Ev}, Tc)]$ (Assumption)

2. $\mathsf{Occurs}(\overset{\bullet-\bullet}{Ec}, Tc)$ (1, $\exists$-elimination)

3. $\exists \overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2\, [\mathsf{Equiv}(\overset{\bullet-\bullet}{Ec}, \mathsf{Clos}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2))]$ (2, **AE5**, FOPC)

4. $\mathsf{Equiv}(\overset{\bullet-\bullet}{Ec}, \mathsf{Clos}(\overset{\bullet}{Ec}_1, \overset{\bullet}{Ec}_2))$ (3, $\exists$-elimination)

5. $\mathsf{Occurs}(\mathsf{Clos}(\overset{\bullet}{Ec}_1, \overset{\bullet}{Ec}_2), Tc)$ (2, 4, **AE2**, FOPC)

6. $\exists S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\overset{\bullet}{Ec}_1, \overset{\bullet}{Ec}_2, S^t v) \wedge \mathsf{Occurs}(\overset{\bullet}{Ec}_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{Ec}_2, Tv_2)$
$\wedge \mathsf{Covers}(Tc, Tv_1, Tv_2)]$

(5, **AE4**, FOPC)

7. $\exists \overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_2, Tv_2)$
$\wedge \mathsf{Covers}(Tc, Tv_1, Tv_2)]$

(6, $\exists$-introduction)

8. $\exists \overset{\bullet-\bullet}{Ev}\, [\mathsf{Occurs}(\overset{\bullet-\bullet}{Ev}, Tc)] \Rightarrow$
$\exists \overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v)$
$\wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_2, Tv_2)$
$\wedge \mathsf{Covers}(Tc, Tv_1, Tv_2)]$

(1, 7, FOPC)

9. $\forall Tv[\exists \overset{\bullet-\bullet}{Ev}\, [\mathsf{Occurs}(\overset{\bullet-\bullet}{Ev}, Tv)] \Rightarrow$
$\exists \overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v, Tv_1, Tv_2[\mathsf{OCPair}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v)$
$\wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_2, Tv_2)$
$\wedge \mathsf{Covers}(Tv, Tv_1, Tv_2)]]$

(8, $\forall$-introduction)

Q.E.D.

Proof of TE4

- $\forall \overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v[\mathsf{OCPair}(\overset{\bullet}{Ev}_1, \overset{\bullet}{Ev}_2, S^t v) \Rightarrow$
$\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\overset{\bullet}{Ev}_1, Tv_1) \wedge \mathsf{Occurs}(\overset{\bullet}{Ev}_2, Tv_2)$
$Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow$

$\wedge \mathsf{Holds}(\mathsf{Prog}(\mathsf{PO}(S^t v)), Tv_3)]]$

1. $\mathsf{OCPair}(\dot{E}c_1, \dot{E}c_2, S^t c)$ (Assumption)

2. $\mathsf{Cat}(\mathsf{Clos}(\dot{E}c_1, \dot{E}c_2), \mathsf{PO}(S^t c))$ (1, **AE6**, FOPC)

3. $\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\dot{E}c_1, Tv_1) \wedge \mathsf{Occurs}(\dot{E}c_2, Tv_2)$

$$Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow$$

$$\wedge \mathsf{Holds}(\mathsf{Prog}(\mathsf{PO}(S^t c)), Tv_3)]$$

(2, **AE7**, FOPC)

4. $\mathsf{OCPair}(\dot{E}c_1, \dot{E}c_2, S^t c) \Rightarrow$

$$\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\dot{E}c_1, Tv_1) \wedge \mathsf{Occurs}(\dot{E}c_2, Tv_2)$$

$$Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow$$

$$\wedge \mathsf{Holds}(\mathsf{Prog}(\mathsf{PO}(S^t c)), Tv_3)]$$

(1, 3, FOPC)

5. $\forall \dot{E}v_1, \dot{E}v_2, S^t v[\mathsf{OCPair}(\dot{E}v_1, \dot{E}v_2, S^t v) \Rightarrow$

$$\forall Tv_1, Tv_2, Tv_3[[\mathsf{Occurs}(\dot{E}v_1, Tv_1) \wedge \mathsf{Occurs}(\dot{E}v_2, Tv_2)$$

$$Tv_1 \prec Tv_3 \wedge Tv_3 \prec Tv_2] \Rightarrow$$

$$\wedge \mathsf{Holds}(\mathsf{Prog}(\mathsf{PO}(S^t v)), Tv_3)]]$$

(4, ∀-introduction)

Q.E.D.

Appendix D

Temporal Progression Algorithms

Algorithm $\pi(\tau)$

1. If $\tau \notin \mathcal{T}$, then

   1a. If $\mathcal{A}(\tau)$ is defined, then return $\mathcal{A}(\tau)$.

   1b. FAIL.

2. If $\tau = $*NOW, then return *COUNT.

3. If $\tau = $*NOW$_i$, for some $i \in \mathbb{N}$, then return $\mathcal{A}(\tau)$.

4. If τ is a transition interval the return 0.

5. If τ is an atomic interval, then FAIL.

6. Return $\sum_{\Phi(t_i) \in \text{ Span}(t)} \pi(t_i)$

Algorithm $\eta(n \in \mathbb{N})$

1. If $n = 0$ then, return 0.

2. Return $1 + \text{round}(\log_{\sqrt{10}}(n))$.

Algorithm $\rho(n \in \mathbb{N})$

1. $h \longleftarrow \eta(n)$.

2. If there is $q \in \Psi(\mathcal{Q})$ such that $\mathcal{A}(q) = h$, then return q.

3. Pick some $q \in \mathcal{Q}$, such that $q \notin \Psi(\mathcal{Q})$.

4. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle q, h\rangle\}$

5. $min \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h' < h\}$.

6. $max \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h < h']\}$.

7. If *min* is not empty, then $\beta \longleftarrow \beta \cup \{q_{gmin} <_{\mathcal{Q}} q\}$, where q_{gmin} is the greatest element of the linearly-ordered poset $\langle min, <_{\mathcal{Q}}\rangle$.

8. If *max* is not empty, then $\beta \longleftarrow \beta \cup \{q <_{\mathcal{Q}} q_{lmax}\}$, where q_{lmax} is the smallest element of the linearly-ordered poset $\langle max, <_{\mathcal{Q}}\rangle$.

9. return q.

Algorithm $\texttt{state\_present}(s, t)$

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s, q)$.

2. For every $t' \in \lceil {*}\mathsf{NOW} \rceil$

    3. If $\beta \not\vdash \neg\mathsf{Holds}(s, t')$

       and $\delta(t') = q'$ or $\beta \vdash \delta(t') <_{\mathcal{Q}} q'$,

       then $\beta \longleftarrow \beta \cup \{t' \sqsubset t\}$.

Algorithm $\texttt{state\_query}(s)$

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s, q)$.

2. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.

3. $\beta \longleftarrow \beta \cup \{\text{*NOW} \sqsubset t, \mathsf{Dur}(t, q')\}$.

4. Initiate deduction for $\mathsf{Holds}(s,t)$.

Algorithm `initialize_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.

2. $\mathsf{COUNT} \longleftarrow 0$.

3. $\mathsf{NOW} \longleftarrow t$.

Algorithm `move_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.

2. $\beta \longleftarrow \beta \cup \{\text{*NOW} \prec t, \mathsf{Dur}(\text{*NOW}, \rho(\text{*COUNT}))\}$.

3. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle \text{*NOW}, \text{*COUNT} \rangle\}$.

4. $\mathsf{COUNT} \longleftarrow 0$.

5. $\mathsf{NOW} \longleftarrow t$.

Algorithm `setup_new_MTF`$(S^{\downarrow} \subseteq \Psi(\mathrm{TEMP}), t_{\mathrm{tr}})$

1. `move_NOW`

2. For all $\mu \in \mathcal{M}_{\mathrm{prop}}$

    3. If there are s and t such that $*\mu = \mathsf{MHolds}(s,t)$, then $\beta \longleftarrow \beta \cup \{\text{*NOW} \sqsubset t\}$.

4. For all $\mu \in \mathcal{M}_{\mathrm{per}}$

    5. For all s and t such that $\mathsf{MHolds}(s,t) \in *\mu$

        6. $\beta \longleftarrow \beta \cup \{\text{*NOW} \sqsubset t\}$.

7. For all $t \in \lceil$*$\mathsf{NOW}_{i-1}\rceil \setminus \{$*$\mathsf{NOW}_{i-1}\}$

   8. If $\eta(\pi(t)) \leq \mathcal{A}(\delta(t))$, then $\beta \longleftarrow \beta \cup \{$*$\mathsf{NOW} \sqsubset t\}$.

   9. Else $\beta \longleftarrow \beta \cup \{t \prec$*$\mathsf{NOW}\}$.

10. $\beta \longleftarrow \beta \cup \{t_{\mathrm{tr}} \prec$*$\mathsf{NOW}\}$.

11. For all $s \in S^{\downarrow}$ $\beta \longleftarrow \beta \cup \{\neg\mathsf{Holds}(s,$*$\mathsf{NOW})\}$.

12. For every state interval $t \in \Phi($*$\mathsf{NOW}_{i-1})$,

   13. If $\beta \vdash$*$\mathsf{NOW} \sqsubset t$, then $\beta \longleftarrow \beta \cup \{$*$\mathsf{NOW} \sqsubset t\}$.

   14. Else, if $\beta \vdash$*$\mathsf{NOW} \not\sqsubset t$, then `event_cease`$(s,t,t_{\mathrm{tr}})$,

       where s is the state with which t is associated.

   15. Else, if $\eta(\pi(t)) \leq \mathcal{A}(q)$, where $\mathsf{SDur}(s,q)$ and $\mathsf{MHolds}(s,t)$, then

       $\beta \longleftarrow \beta \cup \{$*$\mathsf{NOW} \sqsubset t\}$.

   16. Else `event_cease`(s,t,t_{tr}),

       where s is the state with which t is associated.

Algorithm `event_start`(s,t,t_{tr})

1. Pick some $\dot{e} \in \dot{\mathcal{E}}$, such that $\dot{e} \notin \Psi(\dot{\mathcal{E}})$.

2. $\beta \longleftarrow \beta \cup \{\mathsf{Cat}(\dot{e}, \uparrow s), \mathsf{Occurs}(\dot{e}, t_{\mathrm{tr}}), t_{\mathrm{tr}} \supset\subset t\}$.

Algorithm `event_cease`(s,t,t_{tr})

1. Pick some $\dot{e} \in \dot{\mathcal{E}}$, such that $\dot{e} \notin \Psi(\dot{\mathcal{E}})$.

2. $\beta \longleftarrow \beta \cup \{\mathsf{Cat}(\dot{e}, \downarrow s), \mathsf{Occurs}(\dot{e}, t_{\mathrm{tr}}), t \supset\subset t_{\mathrm{tr}}\}$.

Algorithm `state_change`$(S^{\uparrow} \subseteq \textsc{Temp}, S^{\downarrow} \subseteq \textsc{Temp})$

1. $P_{\text{new}} \longleftarrow \{\}$.

2. Pick some $t_{\text{tr}} \in \mathcal{T}$, such that $t_{\text{tr}} \notin \Psi(\mathcal{T})$.

3. $\beta \longleftarrow \beta \cup \{\text{*NOW} \prec t_{\text{tr}}\}$.

4. For all $s_i \in S^{\uparrow}$

   5. If $\beta \vdash \mathsf{Holds}(s_i, \text{*NOW})$ then `start_ceive`(s_i, t_i),

      where t_i is the state interval associated with s_i such that $\beta \vdash \text{*NOW} \sqsubset t_i$

   6. else

      6a. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.

      6b. If $\beta \vdash \neg\mathsf{Holds}(s_i, \text{*NOW})$ then `state_start`$(s_i, t_i, t_{\text{tr}})$.

      6c. Else, `state_persist`(s_i, t_i).

      6d. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{Mholds}(s_i, t_i)\}$.

7. For all $s_i \in S^{\downarrow}$, `cease_perceive`(s_i).

8. `setup_new_MTF`$(\{\}, t_{\text{tr}})$.

9. `Forward`(P_{new}).

Algorithm `start_ceive`(s, t)

1. `start_proprioceive`(s, t).

2. `start_perceive`(s, t).

Algorithm `start_proprioceive`(s, t)

1. If $\text{Mod}_{\text{prop}}(s) \neq \{\}$, then

   for all $\mu \in \text{Mod}_{\text{prop}}(s)$, $\mu \longleftarrow \mathsf{MHolds}(s, t)$.

Algorithm $\texttt{start\_perceive}(s,t)$

1. If $\text{Mod}_{\text{per}}(s)$ is defined, then

   $\text{Mod}_{\text{per}}(s) \longleftarrow *\text{Mod}_{\text{per}}(s) \cup \{\mathsf{MHolds}(s,t)\}$.

Algorithm $\texttt{cease\_perceive}(s)$

1. If $\text{Mod}_{\text{per}}(s)$ is defined, then

   $\text{Mod}_{\text{per}}(s) \longleftarrow *\text{Mod}_{\text{per}}(s) \setminus \{\mathsf{MHolds}(s,t)\}$

Algorithm $\texttt{state\_start}(s,t,t_{\text{tr}})$

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t)\}$.

2. $\texttt{event\_start}(s,t,t_{\text{tr}})$.

3. $\texttt{start\_ceive}(s,t)$.

Algorithm $\texttt{state\_persist}(s,t)$

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t)\}$.

2. Pick some $t' \in \mathcal{T}$, such that $t' \notin \Psi(\mathcal{T})$.

3. $\texttt{event\_start}(s,t,t')$.

4. $\texttt{start\_perceive}(s,t)$.

Algorithm $\texttt{assert\_persist}(S \subseteq \text{TEMP} \cup \text{PERM})$

1. $P_{\text{new}} \longleftarrow \{\}$.

2. For all $s_i \in S$

   3. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.

4. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i,t_i), {}^*\mathsf{NOW} \sqsubset t_i\}$.

5. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i,t_i)\}$.

6. Pick some $t_i' \in \mathcal{T}$, such that $t_i' \notin \Psi(\mathcal{T})$.

7. $\texttt{event\_start}(s_i,t_i,t_i')$.

8. $\texttt{Forward}(P_{\text{new}})$.

Algorithm $\texttt{assert\_start}(S \subseteq \text{TEMP} \cup \text{PERM})$

1. Pick some $t_{\text{tr}} \in \mathcal{T}$, such that $t_{\text{tr}} \notin \Psi(\mathcal{T})$.

2. $\beta \longleftarrow \beta \cup \{{}^*\mathsf{NOW} \prec t_{\text{tr}}\}$.

3. $\texttt{setup\_new\_MTF}(\{\}, t_{\text{tr}})$.

4. $P_{\text{new}} \longleftarrow \{\}$.

5. For all $s_i \in S$

   6. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.

   7. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i,t_i), {}^*\mathsf{NOW} \sqsubset t_i\}$.

   8. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i,t_i)\}$.

   9. $\texttt{event\_start}(s_i,t_i,t_{\text{tr}})$.

10. $\texttt{Forward}(P_{\text{new}})$.

Algorithm $\texttt{assert\_cease}(S \subseteq \text{TEMP})$

1. Pick some $t_{\text{tr}} \in \mathcal{T}$, such that $t_{\text{tr}} \notin \Psi(\mathcal{T})$.

2. $\texttt{setup\_new\_MTF}(S, t_{\text{tr}})$.

3. $P_{\text{new}} \longleftarrow \{\}$.

4. For all $s_i \in S$, $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\neg\mathsf{Holds}(s_i, {}^*\mathsf{NOW})\}$.

5. $\texttt{Forward}(P_{\text{new}})$.

Algorithm `initialize`$(S \subseteq \textsc{Temp})$

1. $P_{\text{new}} \longleftarrow \{\}$.
2. `initialize_NOW`.
3. For all $s_i \in S$
    4. Pick some $t_i \in \mathcal{T}$ such that $t_i \notin \Psi(\mathcal{T})$.
    5. Pick some $t'_i \in \mathcal{T}$, such that $t'_i \notin \Psi(\mathcal{T})$.
    6. `state_start`(s_i, t_i, t'_i).
    7. $\beta \longleftarrow \beta \cup \{\text{*NOW} \sqsubset t_i\}$.
    8. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
9. `Forward`(P_{new}).

Algorithm `Forward`$(P \subseteq \Psi(\mathcal{P}))$

1. $P_{\text{inf}} \longleftarrow$ `Forward_old`(P).
2. For every $p \in P \cup P_{\text{inf}}$
    3. If $p = \mathsf{Holds}(s, \text{*NOW})$, for some $s \in \Psi(\textsc{Temp})$, then
        4. If there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s, t)$, then `state_present`$(s, t)$.
        5. Else
            5a. Pick some $t \in \mathcal{T}$ such that $t \notin \Psi(\mathcal{T})$.
            5b. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s, t), \text{*NOW} \sqsubset t\}$.
            5c. `state_present`(s, t).

Bibliography

Ali, S. S. and Shapiro, S. C. (1993). Natural language processing using a propositional semantic network with structured variables. *Minds and Machines*, 3(4):421–451.

Allen, J. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11).

Allen, J. (1984). Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154.

Allen, J. and Hayes, P. (1985). A common-sense theory of time. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 528–531, Los Altos, CA. Morgan Kaufmann.

Allen, J., Hendler, J., and Tate, A., editors (1990). *Readings in Planning*. Morgan Kaufmann, San Mateo, CA.

Allen, J. F. and Kautz, H. A. (1988). A model of naive temporal reasoning. In (Hobbs and Moore, 1988), pages 261–268.

Almeida, M. (1987). *Reasoning about the Temporal Structure of Narratives*. PhD thesis, Department of Computer Science, State University of New York at Buffalo. Techincal report 87–10.

Almeida, M. (1995). Time in narratives. In Duchan, J., Bruder, G., and Hewitt, L., editors, *Deixis in Narrative: A Cognitive Science Approach*, pages 159–189. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.

Almeida, M. and Shapiro, S. C. (1983). Reasoning about the temporal structure of narrative texts. In *Proceedings of the Fifth Annual Meeting of the Cognitive Science Society*, pages 5–9, Rochester, NY.

Ambros-Ingerson, J. and Steel, S. (1988). Integrating planning, execution and monitoring. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 83–88, San Mateo, CA. Morgan Kaufmann.

Artale, A. and Franconi, E. (1998). A temporal description logic for reasoning about actions and plans. *Journal of Artificial Intelligence Research*, 9:463–506.

Bacchus, F., Tenenberg, J., and Koomen, J. (1991). A non-reified temporal logic. *Artificial Intelligence*, 52:87–108.

Bach, E. (1986). The algebra of events. *Linguistics and Philosophy*, 9:5–16.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. The MIT Press, Cambridge, MA.

Bettini, C. and De Sibi, R. (1999). Symbolic representation of user-defined time granularities. In *Proceedings of the 6th International Workshop on Temporal Representation and Reasoning (TIME'99)*, Orlando, FL. IEEE Computer Society Press.

Bettini, C., Wang, X. S., and Jajodia, S. (1996). A general framework and reasoning models for time granularity. In *Proceedings of the 3rd International Workshop on Temporal Representation and Reasoning (TIME'96)*, Key West, FL. IEEE Computer Society Press.

Brown, F. M., editor (1987). *The Frame Problem in Artificial Intelligence : Proceedings of the 1987 Workshop*, Los Altos, CA. Morgan Kaufmann.

Cartwright, R. (1975). Scattered objects. In Lehrer, K., editor, *Analysis and Metaphysics: Essays in honor of R. M. Chisholm*, pages 153–171. D. Reidel Publishing Company, Dordrecht, Holland.

Chen, X. and De Giacomo, G. (1999). Reasoning about nondeterministic and concurrent actions: A process algebra approach. *Artificial Intelligence*, 107:63–98.

Chisholm, R. M. (1990). Events without times: An essay on ontology. *Noûs*, 24:413–428.

Comrie, B. (1976). *Aspect*. Cambridge University Press, Cambridge.

Cravo, M. and Martins, J. (1993). SNePSwD: A newcomer to the SNePS family. *Journal of Experimental and Theoretical Artificial Intelligence*, 5(2–3):135–148.

Cresswell, M. J. (1990). *Entities and Indices*. Kluwer Academic Publishers, Dordrecht.

Cukierman, D. and Delgrande, J. P. (1998). Expressing time intervals and repetition within a formalization of calendars. *Computational Intelligence*, 14(4):563–597.

Dahl, Ö. (1981). On the definition of the telic-atelic (bounded-nonbounded) distinction. In Tedeschi, P. and Zaenen, A., editors, *Tense and Aspect*, volume 14 of *Syntax and Semantics*, pages 79–90. Academic Press, New York.

Davidson, D. (1963). Actions, reasons, and causes. *The Journal of Philosophy*, 60(23):685–700.

Davidson, D. (1967). The logical form of action sentences. In Recher, N., editor, *The Logic of Decision and Action*, pages 81–95. University of Pittsburgh Press, Pittsburgh, PA.

Davidson, D. (1970). Events as particulars. *Noûs*, 4:25–32.

Davis, E. (1992). Semantics for tasks that can be interrupted or abandoned. In Hendler, J., editor, *The 1st International Conference on Artificial Intelligence Planning Systems (AIPS '92)*, pages 37–44, San Francisco, CA. Morgan Kaufmann.

Davis, E. (1994). Knowledge preconditions for plans. *Journal of Logic and Computation*, 4(5):721–766.

De Eugenio, B. (1998). An action representation formalism to interpret natural language instructions. *Computational Intelligence*, 14(1):89–113.

De Giacomo, G., Lespérance, Y., and Levesque, H. (2000). *ConGolog*, a concurrent programming language based on the situation calculus. *Artificial Intelligence*, 121(1–2):109–169.

De Giacomo, G. and Levesque, H. (1998). An incremental interpreter for high-level programs with sensing. Technical report, Department of Computer Science, University of Toronto.

De Giacomo, G., Reiter, R., and Soutchanski, M. (1998). Execution monitoring of high-level robot programs. In Cohn, A. G., Schubert, L. K., and Shapiro, S. C., editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR '98)*, pages 453–464, San Francisco, CA. Morgan Kaufmann.

De Swart, H. (1998). Aspect shift and coercion. *Natural Language and Linguistic Theory*, 16:347–385.

Dean, T. and Kanazawa, K. (1988). Probabilistic temporal reasoning. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 524–528, Los Altos, CA. Morgan Kaufmann.

Declerck, R. (1979). Aspect and the bounded/unbounded (telic/atelic) distinction. *Linguistics*, 17:761–794.

Declerck, R. (1989). Boundedness and the structure of situations. *Leuvense Bijdragen*, 78:275–304.

Depraetere, I. (1995). On the necessity of distinguishing between (un)boundedness and (a)telicity. *Linguistics and Philosophy*, 18:1–19.

Desclés, J.-P. (1989). State, event, process, and topology. *General Linguistics*, 29(3):159–200.

Devlin, K. (1991). *Logic and Information*. Cambridge University Press, Cambridge.

Dix, J., Kraus, S., and Subrahmanian, V. (2001). Temporal agent programs. *Artificial Intelligence*, 127(1):87–135.

Dowden, B. (1998). What is essential to being a clock. In Fieser, J., editor, *Internet Encyclopedia of Philosophy <http://www.utm.edu/research/iep/t/time>*.

Dowty, D. (1977). Toward a semantic analysis of verb aspect and the English 'imperfective' progressive. *Linguistics and Philosophy*, 1:45–77.

Duchan, J., Bruder, G., and Hewitt, L., editors (1995). *Deixis in Narrative: A Cognitive Science Perspective*. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.

Elgot-Drapkin, J. J. and Perlis, D. (1990). Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2:75–98.

Eschenbach, C., Habel, C., and Smith, B., editors (1994). *Topological Foundations of Cognitive Science*. Papers from the Workshop at the FISI-CS, Buffalo, NY.

Euzanet, J. (1995). An algebraic approach to granularity in qualitative time and space representation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 894–900, San Mateo, CA. Morgan Kaufmann.

Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24:85–168.

Freksa, C. (1992). Temporal reasoning based on semi-intervals. *Artificial Intelligence*, 54:199–227.

Friedman, W. (1990). *About Time: Inventing the Fourth Dimension*. MIT Press, Cambridge, MA.

Gabbay, D., Hogger, C., and Robinson, J., editors (1994). *Nonmonotonic Reasoning and Uncertain Reasoning*, volume 3 of *The Handbook of Logic in Artificial Intelligence and Logic Programming*. Oxford University Press, Oxford.

Gabbay, D., Hogger, C. J., and Robinson, J. A., editors (1995). *Epistemic and Temporal Reasoning*, volume 4 of *Handbook of Logic in Artificial Intelligence and Logic Programming*. Clarendon Press, Oxford.

Gabbay, D. and Moravcsik, J. (1979). Verbs, events, and the flow of time. In Rohrer, C., editor, *Time, Tense, and Qunatifiers*, pages 59–83. Niemeyer, Tübingen.

Galton, A. (1984). *The Logic of Aspect*. Clarendon Press, Oxford.

Galton, A. (1987a). The logic of occurrence. In (Galton, 1987c), pages 169–196.

Galton, A. (1987b). Temporal logic and computer science: An overview. In (Galton, 1987c), pages 1–52.

Galton, A., editor (1987c). *Temporal Logics and Their Applications*. Academic Press.

Galton, A. (1990). A critical examination of Allen's theory of action and time. *Artificial Intelligence*, 42:159–188.

Galton, A. (1995). Time and change for AI. In (Gabbay et al., 1995), pages 175–240.

Garey, H. (1957). Verbal aspect in French. *Language*, 33:91–110.

Georgeff, M. and Lansky, A. (1987). Reactive reasoning and planning. In *Proceedings of the 6th National Conference on Artificial Intelligence*, pages 677–682, Los Altos, CA. Morgan Kaufmann.

Giunchiglia, E. and Lifschitz, V. (1995). Dependent fluents. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, volume 2, pages 1964–1969, San Mateo, CA. Morgan Kaufmann.

Givón, T. (1979). *On Understanding Grammar*. Academic Press, New York, NY.

Goldman, A. (1970). *A Theory of Human Action*. Princeton University Press, Princeton, NJ.

Gurney, J., Perlis, D., and Purang, K. (1997). Interpreting presuppositions using active logic: From contexts to utterances. *Computational Intelligence*, 5(4).

Habel, C. (1994). Discreteness, finiteness, and the structure of topological spaces. In (Eschenbach et al., 1994).

Hanks, S. and McDermott, D. (1987). Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33:379–412.

Hayes, P. (1985). The second naive physics manifesto. In Hobbs, J. and Moore, R., editors, *Formal Theories of the Commonsense World*, pages 1–36. Ablex Publishing Corporation, Norwood, NJ.

Herweg, M. (1991a). A critical examination of two classical approaches to aspect. *Journal of Semantics*, 8:363–402.

Herweg, M. (1991b). Perfective and imperfective aspect and the theory of events and states. *Linguistics*, 29:969–1010.

Hexmoor, H. (1995). *Representing and Learning Routine Activities*. PhD thesis, Department of Computer Science, State University of New York at Buffalo, Buffalo, NY.

Hexmoor, H., Lammens, J., and Shapiro, S. C. (1993). Embodiment in GLAIR: a grounded layered architecture with integrated reasoning for autonomous agents. In Dankel, II, D. D. and Stewman, J., editors, *Proceedings of The Sixth Florida AI Research Symposium (FLAIRS 93)*, pages 325–329. The Florida AI Research Society.

Hexmoor, H. and Shapiro, S. C. (1997). Integrating skill and knowledge in expert agents. In Feltovich, P. J., Ford, K. M., and Hoffman, R. R., editors, *Expertise in Context*, pages 383–404. AAAI Press/MIT Press, Menlo Park, CA / Cambridge, MA.

Hirst, G. (1991). Existence assumptions in knowledge representation. *Artificial Intelligence*, 49:199–242.

Hobbs, J. (1985). Granularity. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 432–435, Los Altos, CA. Morgan Kaufmann.

Hobbs, J. (2000). Half orders of magnitude. In Obrst, L. and Mani, I., editors, *Papers from the Workshop on Semantic Approximation, Granularity, and Vagueness*, pages 28–38. A Workshop of the Seventh International Conference on Principles of Knowledge Representation and Reasoning.

Hobbs, J. and Moore, R., editors (1988). *Formal Theories of the Commonsense World*. Ablex Publishing Corporation, Norwood, NJ.

Hughes, G. E. and Cresswell, M. J. (1996). *A New Introduction to Modal Logic*. Routledge, London and New York.

Huntley, M. (1984). The semantics of english imperatives. *Linguistics and Philosophy*, 7:103–133.

Ismail, H. O. and Shapiro, S. C. (1999). Cascaded acts: Conscious sequential acting for embodied agents. Technical Report 99-10, Department of Computer Science and Engineering, State University of New York at Buffalo.

Ismail, H. O. and Shapiro, S. C. (2000a). Conscious error recovery and interrupt handling. In Arabnia, H. R., editor, *Proceedings of the International Conference on Artificial Intelligence (IC-AI'2000)*, pages 633–639, Las Vegas, NV. CSREA Press.

Ismail, H. O. and Shapiro, S. C. (2000b). Two problems with reasoning and acting in time. In Cohn, A., Giunchiglia, F., and Selman, B., editors, *Principles of Knowledge Representation* *and Reasoning: Proceedings of the Seventh International Conference (KR 2000)*, pages 355–365, San Francisco, CA. Morgan Kaufmann.

Israel, D. (1995). Process logics of action. Draft. Avialable at <http://www.ai.sri.com/˜israel/>.

Jackendoff, R. (1990). *Semantic Structures*. MIT Press, Cambridge, MA.

Jackendoff, R. (1991). Parts and boundaries. *Cognition*, 41:9–45.

Jackendoff, R. (1996). The proper treatment of measuring out, telicity, and perhaps even quantification in english. *Natural Language and Linguistic Theory*, 14:305–354.

Johnson, F. L. and Shapiro, S. C. (2000a). Automatic belief revision in SNePS. In Baral, C. and Truszczynski, M., editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning (NMR2000)*. Unpaginated, 5 pages.

Johnson, F. L. and Shapiro, S. C. (2000b). Formalizing a deductively open belief space. Technical Report 2000-02, Department of Computer Science and Engineering, State University of New York at Buffalo.

Johnson, F. L. and Shapiro, S. C. (2000c). Implementing integrity constraints in an existing belief revision system. In Baral, C. and Truszczynski, M., editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning (NMR2000)*. Unpaginated, 8 pages.

Kamp, H. (1971). Formal properties of 'now'. *Theoria*, 37:227–273.

Kamp, H. (1979). Some remarks on the logic of change. Part I. In Rohrer, C., editor, *Time, Tense, and Quantifiers*, pages 135–174. Niemeyer, Tübingen.

Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer Academic, Dordrecht.

Kanazawa, K. (1992). *Reasoning about Time and Probability*. PhD thesis, Brown University, Department of Computer Science, Providence, RI. Technical Report CS-92-61.

Kaplan, D. (1979). On the logic of demonstratives. *Journal of Philosophical Logic*, 8:81–98.

Kautz, H. (1986). The logic of persistence. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 401–405, Los Altos, CA. Morgan Kaufmann.

Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4:1–45.

Konolige, K. (1994). Autoepistemic logic. In (Gabbay et al., 1994), pages 217–295.

Koomen, J. (1991). Reasoning about recurrence. *International Journal of Intelligent Systems*, 6:461–496.

Kowalski, R. and Sergot, M. (1986). A logic-based calculus of events. *New Generation Computing*, 4:67–95.

Krifka, M. (1989). Nominal reference, temporal constitution and quantification in event semantics. In Bartsch, R., van Benthem, J., and van Emde Boas, P., editors, *Semantics and Contextual Expression*, pages 75–116. Foris Publications, Dordrecht, Holland.

Krifka, M. (1998). The origins of telicity. In Rothstein, S., editor, *Events and Grammar*, pages 197–235. Kluwer Academic Publishers, Dordrecht.

Kumar, D. (1994). *From Beliefs and Goals to Intentions and Actions: An Amalgamated Model of Inference and Acting*. PhD thesis, Department of Computer Science, State University of New York at Buffalo, Buffalo, NY. Technical Report 94-04.

Kumar, D. (1996). The SNePS BDI architecture. *Decision Support Systems*, 16:3–19.

Kumar, D. and Shapiro, S. C. (1994a). Acting in service of inference (and vice versa). In Dankel, II, D. D., editor, *Proceedings of the Seventh Florida Artificial Intelligence Research Symposium*, pages 207–211, St. Petersburg, FL. The Florida AI Research Society.

Kumar, D. and Shapiro, S. C. (1994b). The OK BDI architecture. *International Journal on Artificial Intelligence Tools*, 3(3):349–366.

Ladkin, P. (1986). Primitives and units for time specification. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 354–359, Los Altos, CA. Morgan Kaufmann.

Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press, Chicago and London.

Lakoff, G. and Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. Basic Books, New York, NY.

Landman, F. (1992). The progressive. *Natural Language Semantics*, 1:1–32.

Langacker, R. (1987). *Foundations of Cognitive Grammar*, volume 1. Stanford University Press, Stanford, CA.

Leonard, H. and Goodman, N. (1940). The calculus of individuals and its uses. *The Journal of Symbolic Logic*, 5(2):45–55.

Lespérance, Y. and Levesque, H. (1994). An argument for indexical representations in temporal reasoning. In Elio, R., editor, *Proceedings of the Tenth Biennial Conference of the Canadian Society for Computational Studies of Intelligence*, pages 271–277, Banff, Canada.

Lespérance, Y. and Levesque, H. (1995). Indexical knowledge and robot action– a logical account. *Artificial Intelligence*, 73(1–2):69–115.

Lespérance, Y., Tam, K., and Jenkin, M. (1998). Reactivity in a logic-based robot programming framework. In *Cognitive Robotics: Papers from the 1998 AAAI Fall Symposium*, pages 98–105, Menlo Park, CA. AAAI Press. Technical report FS-98-02.

Levesque, H. and Reiter, R. (1998). High-level robotics: Beyond planning. In *Cognitive Robotics: Papers from the 1998 AAAI Fall Symposium*, pages 106–108, Menlo Park, CA. AAAI Press. Technical report FS-98-02.

Levesque, H., Reiter, R., Lespérance, Y., Lin, F., and Scherl, R. (1997). GOLOG: A logic programming language for dynamic domains. *The Journal of Logic Programming*, 31(1–3):59–83.

Levin, I. and Zakay, D., editors (1989). *Time and Human Cognition: A Life-Span Perspective*. Number 59 in Advances in Psychology. North-Holland, Amsterdam.

Lewis, D. (1973). Causation. *The Journal of Philosophy*, 70(17):556–567.

Lifschitz, V. (1986). Point-wise circumscription: Preliminary report. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 406–410, Los Altos, CA. Morgan Kaufmann.

Lifschitz, V. (1987). Formal theories of action. In (Brown, 1987), pages 35–58.

Lifschitz, V. (1994). Circumscription. In (Gabbay et al., 1994), pages 297–352.

Lifschtiz, V. and Rabinov, A. (1989). Things that change by themselves. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 864–867, Los Altos, CA. Morgan Kaufmann.

Lin, F. (1995). Embracing causality in specifying the indirect effects of actions. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 1985–1991.

Link, G. (1987). Algebraic semantics of event structures. In Groenendijk, J., Stokhof, M., and Veltman, F., editors, *Proceedings of the Sixth Amsterdam Collouium*, pages 243–262, Amsterdam. ITLI. Reprinted as chapter 11 of (Link, 1998).

Link, G. (1998). *Algebraic Semantics in Language and Philosophy*. CSLI Publications, Stanford, CA.

Maida, A. S. and Shapiro, S. C. (1982). Intensional concepts in propositional semantic networks. *Cognitive Science*, 6:291–330.

Mani, I. (1998). A theory of granularity and its application to problems of polysemy and underspecification of meaning. In Cohn, A. G., Schubert, L. K., and Shapiro, S. C., editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR '98)*, pages 245–255, San Francisco, CA.

Martins, J. and Shapiro, S. C. (1988). A model for belief revision. *Artificial Intelligence*, 35(1):25–79.

McCann, H. (1998). *The Works of Agency: On Human Action, Will, and Freedom*. Cornell University Press, Ithaca, NY.

McCarthy, J. (1980). Circumscription— A form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39.

McCarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, D. and Michie, D., editors, *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, Edinburgh, Scotland.

McDermott, D. (1982). A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101–155.

McDermott, D. and Doyle, J. (1980). Non-monotonic logic I. *Artificial Intelligence*, 13(1–2):41–72.

McKay, D. P. and Shapiro, S. C. (1981). Using active connection graphs for reasoning with recursive rules. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligenc*, pages 368–374, Los Altos, CA. Morgan Kaufmann.

McTaggart, J. E. (1908). The unreality of time. *Mind*, 68:457–474.

Michaelis, L. (1994). The ambiguity of the English present perfect. *Journal of Linguistics*, 30:111–157.

Millikan, R. G. (1990). The myth of the essential indexial. *Noûs*, 24(5):723–734.

Mitchell, J. C. (1996). *Foundations for Programming Languages*. The MIT Press, Cambridge, MA.

Moens, M. and Steedman, M. (1988). Temporal ontology and temporal reference. *Computational Linguistics*, 14:15–28.

Montague, R. (1960). On the nature of certain philosophical entities. *The Monist*, 53:159–194.

Moore, R. C. (1984). Possible-world semantics for autoepistemic logic. In *Proceedings of the 1984 Non-monotonic Reasoning Workshop*, pages 344–354, New Paltz, NY. American Association for Artificial Intelligence.

Moore, R. C. (1985). Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94.

Moore, R. C. (1988). A formal theory of knowledge and action. In (Hobbs and Moore, 1988).

Morgenstern, L. (1996). The problem with solutions to the frame problem. In Ford, K. and Pylyshyn, Z., editors, *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*. Ablex, Norwood.

Mourelatos, A. P. D. (1978). Events, processes, and states. *Linguistics and Philosophy*, 2:415–434.

Myers, K. and Smith, D. (1988). The persistence of derived information. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 496–500, Los Altos, CA. Morgan Kaufmann.

Nirkhe, M., Kraus, S., Miller, M., and Perlis, D. (1997). How to (plan to) meet a deadline between *now* and *then*. *Journal of Logic and Computation*, 7(1):109–156.

Ortiz, C. (1999). A commonsense language for reasoning about causation and rational action. *Artificial Intelligence*, 111:73–130.

Pagan, F. G. (1981). *Formal Specification of Programming Languages : A Panoramic Primer*. Prentice-Hall, Englewood Cliffs, N.J.

Parsons, T. (1989). The progressive in English: Events, states, and processes. *Linguistics and Philosophy*, 12:213–241.

Partee, B. H. (1984). Nominal and temporal anaphora. *Linguistics and Philosophy*, 7:243–286.

Pastor, M. and Artieda, J., editors (1996). *Time, Internal Clocks and Movement*. Number 115 in Advances in Psychology. Elsevier, Amsterdam.

Perlis, D., Elgot-Drapkin, J. J., and Miller, M. (1991). Stop the world—I want to think. *International Journal of Intelligent Systems*, 6:443–456.

Perry, J. (1979). The problem of the essential indexical. *Noûs*, 13(1):3–21.

Pfeifer, K. (1989). *Actions and Other Events: The Unifier-Multiplier Controversy*. Peter Lang Publishing, Inc., New York, NY.

Pianesi, F. and Varzi, A. (1996a). Events, topology, and temporal relations. *The Monist*, 79(1):89–116.

Pianesi, F. and Varzi, A. (1996b). Refining temporal reference in event structures. *Notre Dame Journal of Formal Logic*, 37(1):71–83.

Pinto, J. A. (1994). *Temporal Reasoning in the Situation Calculus*. PhD thesis, Department of Computer Science, University of Toronto, Toronto, ON.

Poynter, D. (1989). Judging the duration of time intervals: A process of remembering segments of experience. In (Levin and Zakay, 1989), pages 305–331.

Prior, A. (1967). *Past, Present and Future*. Clarendon Press, Oxford.

Prior, A. N. (1968). "Now". *Noûs*, 2(2):101–119.

Pustejovsky, J. (1991). The syntax of event structure. *Cognition*, 41:47–81.

Pylyshyn, Z., editor (1986). *The Robot's Dilemma : The Frame Problem in Artificial Intelligence*. Ablex, Norwood, NJ.

Quine, W. V. O. (1960). *Word and Object*. The MIT Press, Cambridge, MA.

Rapaport, W. J. (1978). Meinongian theories and a Russellian paradox. *Noûs*, 12:153–180.

Reichenbach, H. (1947). *Elements of Symbolic Logic*. Macmillan Co., New York, NY.

Reiter, R. (1991). The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In Lifschitz, V., editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, San Diego, CA.

Reiter, R. (1998). Sequential, temporal GOLOG. In Cohn, A. G., Schubert, L. K., and Shapiro, S. C., editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR '98)*, pages 547–556, San Francisco, CA. Morgan Kaufmann.

Rescher, N. and Garson, J. (1968). Topological logic. *The Journal of Symbolic Logic*, 33(4):537–548.

Rosch, E. (1978). Principles of categorization. In Rosch, E. and Lloyd, B. B., editors, *Cognition and Categorization*, pages 27–48. Lawrence Erlbaum Associates, Hillsdale, NJ.

Rumelhart, D., Lindsay, P., and Norman, D. (1972). A process model for long-term memory. In (Tulving and Donaldson, 1972), pages 198–248.

Sadri, F. (1987). Three recent approaches to temporal reasoning. In (Galton, 1987c), pages 121–168.

Schein, B. (1993). *Plurals and Events*. The MIT Press, Cambridge, MA.

Schmiedel, A. (1990). A temporal terminological logic. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 640–645, San Mateo, CA. Morgan Kaufmann.

Schubert, L. (2000). The situations we talk about. In Minker, J., editor, *Logic-Based Artificial Intelligence*. Kluwer Academic Publishers, Dordrecht. In press.

Schubert, L. and Hwang, C. H. (2000). Episodic logic meets Little Red Riding Hood: A comprehensive natural representation for language understanding. In Iwańska, Ł. M. and Shapiro, S. C., editors, *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*, pages 111–174. AAAI Press/The MIT Press, Menlo Park, CA.

Shanahan, M. (1995). A circumscriptive calculus of events. *Artificial Intelligence*, 77:249–284.

Shanahan, M. (1997). *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. MIT Press, Cambridge, MA.

Shanahan, M. (1998). Reinventing shakey. In *Cognitive Robotics: Papers from the 1998 AAAI Fall Symposium*, pages 125–135, Menlo Park, CA. AAAI Press. Technical report FS-98-02.

Shapiro, S. C. (1989). The CASSIE projects: An approach to natural language competence. In *Proceedings of the 4th Portugese Conference on Artificial Intelligence*, pages 362–380, Lisbon, Portugal. Springer-Verlag.

Shapiro, S. C. (1991). Cables, paths, and ”subconscious” reasoning. In Sowa, J. F., editor, *Principles of Semantic Networks: Explorations in the Representation of knowledge*, pages 137–156. Morgan Kaufmann, San Mateo, CA.

Shapiro, S. C. (1993). Belief spaces as sets of propositions. *Journal of Experimental and Theoretical Artificial Intelligence*, 5:225–235.

Shapiro, S. C. (1996). Formalizing English. *International Journal of Expert Systems*, 9(1):151–171.

Shapiro, S. C. (1998). Embodied Cassie. In *Cognitive Robotics: Papers from the 1998 AAAI Fall Symposium*, pages 136–143, Menlo Park, CA. AAAI Press. Technical report FS-98-02.

Shapiro, S. C. (2000). SNePS: A logic for natural language understanding and commonsense reasoning. In Iwańska, Ł. M. and Shapiro, S. C., editors, *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*, pages 175–195. AAAI Press/The MIT Press, Menlo Park, CA.

Shapiro, S. C., Kumar, D., and Ali, S. (1989). A propositional network approach to plans and plan recognition. In *Proceedings of the 1988 Workshop on Plan Recognition*, pages 21–41, Los Altos, CA. Morgan Kaufmann.

Shapiro, S. C. and McKay, D. P. (1980). Inference with recursive rules. In *Proceedings of the First Annual National Conference on Artificial Intelligence*, pages 151–153, Los Altos, CA. Morgan Kaufmann.

Shapiro, S. C. and Rapaport, W. J. (1987). SNePS considered as a fully intensional propositional semantic network. In Cercone, N. and McCalla, G., editors, *The Knowledge Frontier*, pages 263–315. Springer-Verlag, New York.

Shapiro, S. C. and Rapaport, W. J. (1991). Models and minds: Knowledge representation for natural-language competence. In Cummins, R. and Pollock, J., editors, *Philosophy and AI: Essays at the Interface*, pages 215–259. MIT Press, Cambridge, MA.

Shapiro, S. C. and Rapaport, W. J. (1992). The SNePS family. *Computers and mathematics with applications*, 23(2–5):243–275. Reprinted in F. Lehman, ed. *Semantic Networks in Artificial Intelligence*, pages 243–275. Pergamon Press, Oxford, 1992.

Shapiro, S. C. and Rapaport, W. J. (1995). An introduction to a computational reader of narratives. In *Deixis in Narrative: A Cognitive Science Perspective*, pages 79–105. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.

Shapiro, S. C. and the SNePS Implementation Group (1999). *SNePS 2.5 User's Manual*. Department of Computer Science and Engineering, State University of New York at Buffalo.

Shoham, Y. (1985). Ten requirements for a theory of change. *New Generation Computing*, 3:467–477.

Shoham, Y. (1986). Chronological ignorance: Time, nonmonotonicity, necessity and causal theories. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 389–393, Los Altos, CA. Morgan Kaufmann.

Shoham, Y. (1987). Temporal logics in AI: Semantical and ontological considerations. *Artificial Intelligence*, 33:89–104.

Shoham, Y. and McDermott, D. (1988). Problems in formal temporal reasoning. *Artificial Intelligence*, 36:49–61.

Smith, B. and Varzi, A. (2000). Fiat and bona fide boundaries. *Philosophy and Phenomenological Research*, 60(2):401–420.

Smith, C. S. (1997). *The Parameter of Aspect*. Kluwer Academic Publishers, Dordrecht, second edition.

Smith, C. S. (1999). Activities: States or events? *Linguistics and Philosophy*, 22:479–508.

Steedman, M. (2000). *The Productions of Time*. Draft. Available at <*http://www.cogsci.ed.ac.uk/˜steedman/papers.html*>.

Stein, L. A. (1990). An atemporal frame problem. *International Journal of Expert Systems*, 3(4):371–381.

Stein, L. A. and Morgenstern, L. (1994). Motivated action theory: A formal theory of causal reasoning. *Artificial Intelligence*, 71:1–42.

Stoljar, D. (1997). The deflationary theory of truth. In *Stanford Encyclopedia of Philosophy* <*http://plato.stanford.edu/entries/truth-deflationary/*>. The Metaphysics Research Lab, Stanford University.

Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12:49–100.

Talmy, L. (2000). *Toward a Cognitive Semantics*. MIT Press.

Tam, K., Lloyd, J., Lespérance, Y., Levesque, H., Lin, F., Marcu, D., Reiter, R., and Jenkin, M. (1997). Controlling autonomous robots with GOLOG. In Sattar, A., editor, *10th Australian Joint Conference on Artificial Intelligence, AI'97*, Advanced Topics in Artificial Intelligence, pages 1–12, Berlin; New York. Springer Verlag.

Taylor, B. (1977). Tense and continuity. *Linguistics and Philosophy*, 1:199–220.

ter Meulen, A. (1997). *Representing Time in Natural Language*. The MIT Press, Cambridge, MA, second edition.

Thielscher, M. (1999). From situation calculus to fluent calculus: State update axioms as a solution to the inferential frame problem. *Artificial Intelligence*, 111:277–299.

Thomson, J. J. (1971). The time of a killing. *The Journal of Philosophy*, 68(5):115–132.

Traverso, P. and Spalazzi, L. (1995). A logic for acting, sensing, and planning. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, volume 2, pages 1941–1947, San Mateo, CA. Morgan Kaufmann.

Tulving, E. (1972). Episodic and semantic memory. In (Tulving and Donaldson, 1972), pages 382–404.

Tulving, E. and Donaldson, W., editors (1972). *Organization of Memory*. Academic Press, New York, NY.

van Benthem, J. (1983). *The Logic of Time*. D. Reidel Publisihng Company, Dordrecht, Holland.

van Benthem, J. (1995). Temporal logic. In (Gabbay et al., 1995), pages 241–350.

Varzi, A. (1996). Parts, wholes, and part-whole relations: The prospects of mereotopology. *Data and Knowledge Engineering*, 20:259–286.

Vendler, Z. (1957). Verbs and times. *The Philosophical Review*, 66(2):143–160.

Vere, S. A. (1983). Planning in time: Windows and durations for activities and goals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5:246–267.

Verkuyl, H. (1989). Aspectual classes and aspectual composition. *Linguistics and Philosophy*, 12:39–94.

Vlach, F. (1993). Temporal adverbials, tenses and the perfect. *Linguistics and Philosophy*, 16:231–283.

Williams, B. (1986). Doing time: Putting qualitative reasoning on firmer ground. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 105–112.

Woods, W. (1975). What's in a link: Foundations of semantic networks. In Bobrow, D. and Collins, A., editors, *Representation and Understanding: Studies in Cognitive Science*, pages 35–81. Academic Press, New York.