The Resolution of a Hartmanis Conjecture

Jin-yi Cai [1] D. Sivakumar [2]

[1]Department of Computer Science, State University of New York at Buffalo, Buffalo, NY 14260.
Research supported in part by NSF grants CCR-9057486 and CCR-9319093, and an Alfred P. Sloan Fellowship.
Email: cai@cs.buffalo.edu

[2]Department of Computer Science, State University of New York at Buffalo, Buffalo, NY 14260.
Research supported in part by NSF grant CCR-9409104.
Email: sivak-d@cs.buffalo.edu

Abstract

Building on the recent breakthrough by Ogihara, we resolve a conjecture made by Hartmanis in 1978 regarding the (non) existence of sparse sets complete for P under logspace many-one reductions. We show that if there exists a sparse hard set for P under logspace many-one reductions, then P = LOGSPACE. We further prove that if P has a sparse hard set under many-one reductions computable in NC^1, then P collapses to NC^1.

1 Introduction

A set S is called *sparse* if there are at most a polynomial number of strings in S up to length n. Sparse sets have been the subject of study in complexity theory for the past 20 years, as they reveal inherent structure and limitations of computation [BH77, HOW92, You92a, You92b]. For instance, it is well known that the class of languages polynomial time Turing reducible (i.e. by Cook reductions) to a sparse set is precisely the class of languages with polynomial size circuits.

One major motivation for the study of sparse sets, and various reducibilities to them, is concerned with the isomorphism conjectures by Berman and Hartmanis. In 1976, they proved that all the natural NP-complete problems (such as those found in [GJ79]) are isomorphic under polynomial time computable functions [BH77]. Based on this evidence they conjectured that all NP-complete problems under polynomial time many-one reducibility (i.e. Karp reductions) are p-time isomorphic. Noting that the densities of any two p-time isomorphic sets are polynomially related, and all known NP-complete sets are exponentially dense, they also conjectured that there are no sparse complete sets for NP. Based on similar evidence for P-complete problems under logspace many-one reducibility, Hartmanis conjectured in 1978 that there are no sparse complete sets for P under logspace many-one reductions [Har78]. It is this conjecture that we address in our paper.

The Berman-Hartmanis isomorphism conjecture has generated a lot of research work in this field. Building on earlier work by Fortune [For79], Mahaney [Mah82] showed that if NP has a sparse complete set under polynomial time many-one reducibility, then P = NP. This is the definitive result concerning the nonexistence of sparse complete sets for NP under Karp reductions. Note that if P = NP, then both conjectures concerning isomorphism and the nonexistence of sparse complete sets for NP are false. Regarding Cook reductions, the famous result by Karp and Lipton [KL82] showed that if NP has a sparse hard set then the polynomial hierarchy collapses to its second level, $\Sigma_2^p = \Pi_2^p$. In the subsequent years, considerable research effort has been devoted to studying variations of this problem; we especially mention the results by Ogihara and Watanabe concerning bounded truth table reductions of NP to sparse sets [OW91]; see [HOW92] or [You92a, You92b] for a survey.

The current paper resolves the Hartmanis conjecture of 1978 in the sense of Mahaney, namely there are no sparse complete sets for P under logspace many-one reductions, if P $\neq$ LOGSPACE. Unlike the NP case, very little progress had been made on this conjecture till very recently. The only known result till last year is due to Hemachandra, Ogihara and Toda [HOT94]. They showed that if P has *polylogarithmically* sparse hard sets, then P = SC, the class of languages recognizable in simultaneous polynomial time and polylogarithmic space. Because of the assumption of *polylogarithmic* sparsity the result leaves an exponential gap. Very recently, Ogihara [Ogi95] made substantial progress toward resolving the Hartmanis conjecture. He showed that if there is a sparse set S that is hard for P under logspace many-one reductions, then P $\subseteq$ DSPACE[$\log^2 n$]. Our work builds on the work of Ogihara.

Our main result of this paper is the following: if there is a sparse set S that is hard for P under logspace many-one reductions, then P = LOGSPACE. In fact, we prove the stronger statement: if there is a sparse set S that is hard for P under many-one reductions, then the P-complete circuit-value problem can be solved by a logspace-uniform family of polynomial size, logarithmic depth circuits that make polynomially many parallel calls to the reduction. Consequently, if P has a sparse hard set under many-one reductions computable in logspace-uniform NC^1, then P equals logspace-uniform NC^1.

An interesting aspect of our work is that the techniques we employ are probabilistic and algebraic in nature, and are influenced by the recent developments in derandomization techniques, especially constructions of small sample spaces, and the theory of finite fields. The proof of our first theorem begins with a crucial observation due to Ogihara. The main ingredient in the resulting simulation is the solution of a system of linear equations over a finite field. We first prove a probabilistic lemma of general interest. Under the assumption of the existence of a sparse set hard for P, we obtain an RNC^2 simulation of P. Using a "small-bias sample space" construction ([NN90, AGHP90]), we derandomize this algorithm to obtain an NC^2 simulation. Finally, exploiting additional algebraic properties of a closely related construction, we arrive at a Vandermonde system. We then solve the system using closed formulae involving the elementary symmetric polynomials over a certain field and discrete Fourier transforms. The final result is a collapse of P to logspace uniform NC^1. In fact, modulo the complexity of the reduction, the resulting simulation can be done in TC^0.

The basic techniques involving derandomization and algebraic computation are rather powerful. There are already a number of extensions, and many additional results will be reported in a subsequent paper. Those results are primarily concerned with various other reducibilities and complexity classes. A number of additional techniques will be needed, including properties of error correcting codes, and a generalization of Mulmuley's NC^2 algorithm [Mul87], combined with an idea of Chistov [Chi85], to compute the rank of a matrix over a finite field. We will show, e.g., if there exists a sparse set hard for P under bounded truth table reductions, then $P = NC^2$. As an indication of the effectiveness of our derandomization and algebraic techniques, we note that it took the research community 10 years to take the similar step from many-one reducibility in Mahaney's result for NP to bounded truth table reducibility in Ogihara-Watanabe's theorem.

2 Preliminaries

All our notations and definitions are standard. We denote by P the class of all languages recognizable in polynomial time by deterministic Turing machines; NP denotes the class of nondeterministic polynomial time languages. The class of all languages recognizable by deterministic Turing machines that use space no more than $O(\log n)$ is denoted LOGSPACE or L; the corresponding nondeterministic class is denoted by NL. In general, DSPACE$[s(n)]$ denotes the class of languages accepted by deterministic Turing machines, which, on inputs of length n, use space no more than $O(s(n))$.

For circuit and parallel complexity, we use the notation SIZE-DEPTH$[s(n), d(n)]$ to denote the class of languages accepted by a uniform family $\{C_n\}_{n=0}^{\infty}$ of bounded fan-in circuits of size $s(n)$ and depth $d(n)$ for inputs of length n. The criterion for uniformity of the circuit family is usually taken to mean that there is a deterministic space $(\log s(n))$-bounded transducer that, on input 0^n, outputs an encoding of the circuit C_n. The class NC^k is defined as SIZE-DEPTH$[\mathrm{poly}(n), \log^k n]$, and $NC = \bigcup_k NC^k$. (Our NC^1 is logspace-uniform NC^1.) The randomized version of NC^k is denoted by RNC^k.

For any language A, let $c_A(n) \doteq \|\{x \in A \mid |x| \leq n\}\|$ denote the *census function* for A.

A is called (polynomially) *sparse* if $c_A(n)$ is bounded by a polynomial in n.

A Boolean circuit C is a directed acyclic graph with ℓ input nodes labeled $1, \ldots, \ell$, and one output node. The interior nodes, called *gates*, are labeled from the set $\{\neg, \wedge, \vee\}$, and are respectively called NOT, AND and OR gates. On any input $x \in \{0,1\}^n$, the output of each gate is defined in the natural way, as also is the output of the circuit. The *circuit-value problem*, abbreviated CVP, of determining whether a Boolean circuit C outputs '1' on input x was shown by Ladner [Lad75] to be complete for P under logspace-computable many-one reductions. Cook [Coo85] defined the notion of NC^1 reducibility, and notes that this problem is complete for P under NC^1 reductions. This reducibility is somewhat subtle technically, so we refer the reader to [Coo85] for details. However, we remark that a consequence of the completeness of CVP is that if $CVP \in \mathrm{NC}^1$, then $\mathrm{P} = \mathrm{NC}^1$.

All logarithms in this paper are to the base 2.

3 An RNC^2 simulation

In this section, we consider the hypothesis that there is a polynomially sparse set S hard for P under logspace (or even NC^2) many-one reductions. Note that the sparse set S need not belong to P itself. (Thus our assumption is even weaker than P-*completeness* as stated in Hartmanis' conjecture.) The framework and basic ideas introduced here are used in all our results.

Following Ogihara [Ogi95] we define the set A of tuples $\langle C, x, I, b\rangle$ where C is a boolean circuit, x is an input to C, I is a subset of the gates, and b is a bit (0 or 1), such that the sum mod 2 of the values of the gates chosen in I from C on input x equals b, i.e.,

$$\bigoplus_{i \in I} g_i(x) = b.$$

Clearly, $A \in \mathrm{P}$ and hence $A \leq_m^L S$. Let f be a logspace computable function such that for all x, $x \in A \iff f(x) \in S$. It is also obvious that $CVP \leq_m^L A$, therefore it suffices to show that A can be solved in RNC^2.

We note that for any C, x, I, exactly one of the bits $b = 0, 1$ satisfies the equation, and thus exactly one of $f(\langle C, x, I, 0\rangle)$ and $f(\langle C, x, I, 1\rangle)$ is a string in S. Moreover, suppose for two *distinct* subsets I and J and some pair of bits b, b', $f(\langle C, x, I, b\rangle) = f(\langle C, x, J, b'\rangle)$, (we are not assuming that the image is in S). In this case, regardless of whether $\bigoplus_{i \in I} g_i(x) = b$ and $\bigoplus_{i \in J} g_i(x) = b'$ are true or not, they hold or fail simultaneously. Thus we have an equation mod 2 on the values of the gates of C on input x, namely

$$\bigoplus_{i \in I \triangle J} g_i(x) = b \oplus b',$$

and $I \triangle J \neq \emptyset$.

Fix any C and x, let n denote the number of nodes in C (including the inputs, output, and the interior gates). Let N denote the largest value of $|f(\langle C, x, I, b\rangle)|$ (over all I and b). Clearly N is polynomially bounded in n. Let $p(n)$ be a polynomial function that bounds $c_S(N)$. Since there are only polynomially many strings in S, some string $w \in S$ must be mapped on by at least $2^n/p(n)$ many subsets I, more precisely, by the tuple $\langle C, x, I, b_I\rangle$, where b_I is the "right value" $b_I = \bigoplus_{i\in I} g_i(x)$. (For notational simplicity we assume $p(n)$ is a power of 2.) As described above, any two such I give rise to an equation mod 2 on the values of the gates of C on input x. The idea now is to choose polynomially many random subsets $I \in \{0,1\}^n$ and compute $f(\langle C, x, I, 0\rangle)$ and $f(\langle C, x, I, 1\rangle)$, collecting as many equations as possible. The following lemma ensures that this process gives us a system of linear equations of sufficiently high rank, even if we restrict attention to a single "popular" $w \in S$ which appears for at least $2^n/p(n)$ many subsets I.

3.1 A probabilistic lemma

Let $B = \{0,1\}^n$ denote the n-dimensional binary cube. With respect to the finite field of two elements $GF(2) = \mathbf{Z_2}$, B is a vector space of dimension n. Let $T \subseteq B$ be an arbitrary subset of the cube. We ask the following question: If we uniformly and independently pick a sequence of m points in B, what can we say about the probability distribution of the dimension of the affine span of those points picked from T as a function of m, n and $|T|$?

Lemma 1 *Suppose $|T| \geq 2^n/k$, where $k = n^{O(1)}$, then for $m = 2kn^2 + n = n^{O(1)}$, if we uniformly and independently pick a sequence of m points in B, the probability that the dimension of the affine span of the points from T is less than $n - \log_2 k$ is at most $e^{-n^2+O(n\log n)}$.*

Proof. Consider any sequence of points of B being picked by the above process. Let us mark any such sequence $p_1, p_2, \ldots, p_m$ by a 0-1 sequence of the same length m according to the following rule: Suppose the subsequence $p_{i_1}, p_{i_2}, \ldots, p_{i_\ell}$ is the intersection of the sequence $\{p_i\}$ with the set T. p_{i_1} is marked 0. For $j > 1$, precisely those points p_{i_j} are marked 1 if the dimension of the affine span of $p_{i_1}, p_{i_2}, \ldots, p_{i_j}$ is greater than that of $p_{i_1}, p_{i_2}, \ldots, p_{i_{j-1}}$. All other points in $\{p_i\}$ are marked 0. This defines a 0-1 sequence σ of length m. We wish to estimate the probability that the number of 1's in σ is small.

The process of uniformly and independently picking a sequence of m points in B induces a probability distribution over the set of 0-1 sequences σ of length m defined as above. Suppose we have picked a sequence $p_1, p_2, \ldots, p_{i-1}$ which intersects with T in a set whose affine span has dimension $< n - \log_2 k$. Then there are at least $|T| - 2^{n-\log_2 k - 1}$ points of T, which, if picked next, would increase the dimension of the affine span of the intersection. This cardinality is $\geq 2^n/k - 2^n/(2k) = 2^n/(2k)$. Hence the conditional probability

$$\Pr[\sigma_i = 1 \mid \text{ the number of 1's in } \sigma_1, \ldots, \sigma_{i-1} < n - \log_2 k] \geq 1/(2k).$$

For any sequence σ with strictly fewer than $n - \log_2 k$ many 1's,

$$\Pr[\sigma] \leq \left(1 - \frac{1}{2k}\right)^{m-(n-\log_2 k)},$$

which is bounded above by e^{-n^2} if $m = 2kn^2 + n$. Therefore,

$$\Pr[\dim(\text{affine span of } \{p_i\}_{i=1}^m \cap T) < n - \log_2 k] \leq \sum_{j < n - \log_2 k} \binom{m}{j} e^{-n^2} < e^{-n^2 + O(n\log n)}. \qquad \square$$

Now by the above lemma, if in parallel we try polynomially many uniformly and independently chosen I, with high probability we will obtain a system of linear equations with rank deficiency at most $\log_2 p(n)$. We now describe how we can use these to determine in NC^2 the outputs of all the gates of C on input x.

Wolog let the rank of the system be $n - \log_2 p(n)$, and let $m(= n^{O(1)})$ denote the number of equations we have. Denote the equations by $E_1, \ldots, E_m$, and for $i \geq 1$, call an equation E_i *useful* if the rank $rk(E_1, \ldots, E_i) > rk(E_1, \ldots, E_{i-1})$. Clearly the number of useful equations is $n - \log_2 p(n)$. Mulmuley [Mul87] gives an algorithm to compute the rank of an $\ell \times n$ matrix, which, for $\ell = n^{O(1)}$, can be implemented by a circuit of depth $O(\log^2 n)$ and size $n^{O(1)}$. For $1 \leq i \leq m$, we compute in parallel $rk(E_1, \ldots, E_i)$, and identify all the useful equations. Now we have $n - \log_2 p(n)$ equations in n variables, with rank $n - \log_2 p(n)$. We apply the same process to the columns, and identify the $(n - \log_2 p(n))$-many useful columns. We rename the variables so that the first $n - \log_2 p(n)$ columns are all useful. For each of the $p(n)$ possible assignments to the last $\log_2 p(n)$ variables, we create in parallel a system of $n - \log_2 p(n)$ equations as an $(n - \log_2 p(n)) \times (n - \log_2 p(n))$ matrix. Each one of these can be solved in $\log^2 n$ depth and $\mathrm{poly}(n)$ size using the algorithm due to Borodin, et al. [BvzGH82]. For each potential solution we get for the gates of the circuit C on input x, we can check its validity using the local information about the circuit C and input x, such as $x_i = 0$, or $x_i = 1$, or $g_j(x) = g_k(x) \wedge g_\ell(x)$, etc. There will be a unique solution that passes all such tests and we will find the output of $C(x)$ in particular. We have proved:

Theorem 2 *If there is a sparse set that is hard for* P *under logspace or* NC^2 *many-one reductions, then* $\mathrm{P} \subseteq \mathrm{RNC}^2$.

4 Deterministic construction

As before we have $B = \{0,1\}^n = \mathbf{Z_2}^n$ considered as an n-dimensional vector space over the finite field $\mathbf{Z_2}$. For each $I \in B$, let $b_I = \bigoplus_{i \in I} g_i(x)$ be the "right value." Then the string $w = f(\langle C, x, I, b_I \rangle) \in S$ and this w is called the color of I. The presumed reduction to the sparse set S gives a coloring of B with at most $p(n)$ colors. Let $D \subseteq B$ be a subset of B of cardinality bounded by a certain polynomial in n. The coloring of B induces a coloring of D, thus D is the union of at most $p(n)$ many color classes:

$$D = C_1 \cup C_2 \cup \ldots \cup C_{p(n)}.$$

Let the affine span of C_i be denoted by $L_i + d_i$, where L_i is a linear subspace, and d_i is a displacement vector. Let $L = L_1 + L_2 + \ldots + L_{p(n)}$ be the sum of the linear subspaces. We call L the span of the color classes. L_i is spanned by differences of vectors in C_i. For some spanning set of vectors of L_i, each vector in the set gives us an equation mod 2 of the values of the gates of C with the given input. If we collect a generating set of vectors for each L_i, together they span L. Thus, if we can construct a set D with polynomial size and with $\dim L \geq n - O(\log n)$ (irrespective of the coloring), we would have succeeded in derandomizing the construction of the last section. That is, by sampling exhaustively in D, we would have obtained a system of linear equations of rank $\geq n - O(\log n)$.

We claim that the above task can be accomplished as follows: given $p(n)$, construct a polynomial sized set D such that for any linear subspace M of B with $\dim M < n - \log_2 p(n)$, and any $p(n)$ displacement vectors $b_1, \ldots, b_{p(n)} \in B$, the union of the $p(n)$ affine subspaces $\bigcup_{i=1}^{p(n)}(M+b_i)$ does not cover the set D. For if so, then no matter what the induced coloring on D is, the span of the color classes L must be of dimension $\geq n - \log_2 p(n)$, simply because the union of at most $p(n)$ affine subspaces $\bigcup_{i=1}^{p(n)}(L+d_i)$ does cover D:

$$\bigcup_{i=1}^{p(n)}(L+d_i) \supseteq \bigcup_{i=1}^{p(n)}(L_i+d_i) \supseteq D.$$

Let $k = 1 + \log_2 p(n) = O(\log n)$. Wolog, we may assume such a linear subspace M has dimension exactly $= n-k$. Any such M can be specified as the null space of a linear system of equations

$$a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n = 0,$$

where $i = 1, \ldots, k$, and the k vectors $\{(a_{i1}, a_{i2}, \ldots a_{in}) \mid i = 1, \ldots, k\}$ are independent vectors in B over $\mathbf{Z_2}$.

Let $m = 2k + \log_2 n + 1 = 2\log_2 p(n) + \log_2 n + 3 = O(\log n)$. The Galois field $\mathbf{F} = GF(2^m)$ has a vector space structure over $GF(2)$ of dimension m. Choose any basis $\{e_1, \ldots, e_m\}$, then for $u = \sum_{i=1}^m u_i e_i$ and $v = \sum_{i=1}^m v_i e_i$ in $\mathbf{F}$, we can define an inner product by letting

$$\langle u, v\rangle = \sum_{i=1}^m u_i v_i,$$

and doing all arithmetic over $\mathbf{Z_2}$.

The set D is defined as follows:

$$D = \{(\langle 1, v\rangle, \langle u, v\rangle, \ldots, \langle u^{n-1}, v\rangle) \mid u, v \in \mathbf{F}\}.$$

Note that $|D| = 2^{2m} = n^{O(1)}$. Now consider any non-zero vector $a = (a_0, a_1, \ldots, a_{n-1}) \in B$ and any $b \in \mathbf{Z_2}$. We wish to estimate the size of the intersection of D with the affine hyperplane $\sum_{i=0}^{n-1} a_i x_i = b$.

Since the inner product $\langle \cdot, \cdot \rangle$ is bilinear over $\mathbf{Z_2}$ we have

$$\sum_{i=0}^{n-1} a_i \langle u^i, v\rangle = \langle \sum_{i=0}^{n-1} a_i u^i, v\rangle.$$

Let $q_a(X)$ denote the polynomial $\sum_{i=0}^{n-1} a_i X^i \in \mathbf{F}[X]$. If u is a root of the polynomial $q_a(X)$, then clearly the inner product $\langle \sum_{i=0}^{n-1} a_i u^i, v\rangle = 0$. Now suppose $u \in \mathbf{F}$ is not a root of $q_a(X)$, then $\sum_{i=0}^{n-1} a_i u^i = q_a(u)$ is a non-zero element in $\mathbf{F}$. It is easy to see that for any non-zero $w \in \mathbf{F}$,

$$\Pr_{v \in \mathbf{F}}[\langle w, v\rangle = 0] = 1/2.$$

Thus,

$$\begin{aligned}
&\Pr_{u,v \in \mathbf{F}}[\sum_{i=0}^{n-1} a_i \langle u^i, v\rangle = 0] \\
= \; &\Pr_{u \in \mathbf{F}}[u \text{ is a root of } q_a(X)] + \Pr_{u \in \mathbf{F}}[u \text{ is not a root of } q_a(X)] \cdot 1/2.
\end{aligned}$$

But $q_a(X)$ is a non-zero polynomial of degree at most $n-1$, thus

$$\Pr_{u \in \mathbf{F}}[u \text{ is a root of } q_a(X)] \le \frac{n-1}{2^m}.$$

Collecting terms, we have

$$\Pr_{u,v \in \mathbf{F}}[\sum_{i=0}^{n-1} a_i \langle u^i, v\rangle = 0] \le \frac{1}{2} + \frac{n-1}{2^{m+1}}.$$

In particular, if $m > \log_2 n$, both affine hyperplanes $\sum_{i=0}^{n-1} a_i x_i = 0, 1$ must intersect our set D.

In general, consider any k linearly independent equations $\sum_{j=0}^{n-1} a_{ij} x_j = b_i$, where $a_{ij}, b_i \in \mathbf{Z_2}$, and $i = 1, \ldots, k$. Denote this affine space by Π. Denote the point in D specified by u, v as $D(u,v)$. We wish to estimate the probability $\Pr_{u,v \in \mathbf{F}}[D(u,v) \in \Pi]$.

Let $Q = \{\sum_{i=1}^{k} \beta_i [\sum_{j=0}^{n-1} a_{ij} X^j] \mid \beta_i \in \mathbf{Z_2}, \text{ but not all } 0\}$ be a set of polynomials. We claim that the cardinality of Q is exactly $2^k - 1$, and none of the polynomials in Q is the zero polynomial. This follows from the fact that the vectors $(a_{i0}, \ldots, a_{i,n-1})$ are independent over $\mathbf{Z_2}$. Let $u \in \mathbf{F}$ be such that no polynomial in Q has u as a root. For such a u,

$$\sum_{j=0}^{n-1} a_{ij} \langle u^i, v\rangle = \langle \sum_{i=0}^{n-1} a_{ij} u^i, v\rangle = b_i,$$

$i = 1, \ldots, k$, is a linear equation system on (the m bits of) v with linearly independent coefficient vectors over $\mathbf{Z_2}$. (For otherwise, a non-zero linear combination of the coefficient vectors of v will be zero, which is precisely the same as u being a root of one of the polynomials in Q.) Thus, the conditional probability for v to satisfy this linear equation system is precisely $1/2^k$. However, since $|Q| = 2^k - 1$, and each polynomial in Q is non-zero and of degree at most $n-1$,

$$\Pr_{u \in \mathbf{F}}[u \text{ is a root of some polynomial in } Q] \le (2^k - 1)(n-1)/2^m.$$

Collecting terms, we obtain

$$\begin{aligned} & \left| \Pr_{u,v \in \mathbf{F}}[D(u,v) \in \Pi] - \frac{1}{2^k} \right| \\ \le & \quad \frac{(2^k-1)(n-1)}{2^m} \left(1 + \frac{1}{2^k}\right) \\ < & \quad \frac{n}{2^{m-k}}, \end{aligned}$$

which by our choice of m and k is bounded above by $1/2^{k+1}$. Thus, in particular,

$$\Pr_{u,v \in \mathbf{F}}[D(u,v) \in \Pi] > 0.$$

Other than linear independence, the coefficient vectors and the right hand side vector $b_1, \ldots, b_k$ in the definition of Π are arbitrary; the total number of the b vectors is $2^k = 2p(n) > p(n)$, and it follows that no linear subspace M of dimension $< n - \log_2 p(n)$ can cover the set D with some $p(n)$ displacements.

Theorem 3 *If there is a sparse set S which is hard for* P *under* NC^2 *many-one reductions, then* $\mathrm{P} = \mathrm{NC}^2$.

5 The Finale: NC^1 Simulation

In this section, we build upon previous ideas to obtain an optimal simulation. We show that if there is a sparse set S that is hard for P under many-one reductions computable in logspace, then P = LOGSPACE. In fact, we prove the following stronger statement:

Theorem 4 *If a sparse set S is hard for* P *under many-one reductions, then the* P*-complete circuit-value problem can be solved by a logspace-uniform family of polynomial size, logarithmic depth circuits that make polynomially many parallel calls to the reduction.*

That is, modulo the complexity of the reduction to the sparse set, the resulting algorithm can be implemented by a uniform circuit of polynomial size and logarithmic depth. It follows that if the reduction itself is computable in logspace-uniform NC^1, then P equals logspace-uniform NC^1.

Proof. (Sketch) It is known that the polynomial $X^{2\cdot 3^\ell}+X^{3^\ell}+1 \in \mathbf{Z_2}[X]$ is an irreducible polynomial over $\mathbf{Z_2}$ for all $\ell \geq 0$ [vL91]. In the following, by a finite field $GF(2^m)$, where $m = 2\cdot 3^\ell$, we refer explicitly to the field $\mathbf{Z_2}[X]/(X^{2\cdot 3^\ell} + X^{3^\ell} + 1)$.

Let S be a sparse set hard for P under logspace-computable many-one reductions. As before, we will consider a refinement of the circuit-value problem. Define

$$L = \left\{ \langle C, x, 1^m, u, v\rangle \mid m = 2\cdot 3^\ell, u, v \in GF(2^m), \sum_{i=0}^{n-1} u^i g_i = v \right\},$$

where C is a boolean circuit and x is an input to C, and where $g_0, \ldots, g_{n-1}$ are 0-1 variables that denote the values of the gates of C on input x. Here exponentiation and summation are carried out in the finite field $GF(2^m)$. It is easy to see that $L \in$ P, since all the required field arithmetic involved in checking $\sum u^i g_i = v$ can be performed in polynomial time.

Clearly $|\langle C, x, 1^m, u, v\rangle|$ is bounded polynomially in n and m. If f is a logspace-computable function that reduces L to S, the bound on the length of queries made by f on inputs of length $|\langle C, x, 1^m, u, v\rangle|$ is some polynomial $q(n, m)$. Let $p(n, m)$ be a polynomial that bounds the number of strings in S of length at most $q(n, m)$. We will choose the smallest m of the form $2\cdot 3^\ell$ such that $2^m/p(n, m) \geq n$. It is clear that $m = O(\log n)$. Let $\mathbf{F}$ denote the finite extension $GF(2^m)$ of $GF(2)$.

Facts. We first collect some facts about implementing the basic operations of $\mathbf{F}$. For each operation, the number of processors needed is at most $n^{O(1)}$.

(1) Finding a primitive element ω that generates the multiplicative group $\mathbf{F}^*$ of $\mathbf{F}$ can be done in logspace by exhaustive search.

(2) Adding two elements $y_1, y_2 \in \mathbf{F}$ is just the bitwise exclusive-or of the representations of y_1 and y_2, and can be done in depth $O(1)$. Adding $n^{O(1)}$-many elements can be done in depth $O(\log n)$.

(3) Using logarithmic space, it is also possible to build the multiplication table for $\mathbf{F}$, so multiplying two elements of $\mathbf{F}$ can be done by a circuit of depth $O(\log\log n)$ and size $(\log n)^{O(1)}$.

(4) Raising the generator ω to any power $i < 2^m$, or computing the discrete logarithm of any element with respect to ω, can be done by table lookup in depth $O(\log \log n)$. The tables themselves can be precomputed using $O(\log n)$ space.

(5) Multiplying $k = n^{O(1)}$ elements of $\mathbf{F}$ can be done in $O(\log n)$ depth. The idea is to use the discrete logarithms of the k elements with respect to the generator ω, and convert multiplications to additions of k $O(\log n)$-bit integers (modulo $2^m - 1$), which can be done in $O(\log n)$ depth using the folklore 3-to-2 trick.

Our parallel algorithm for *CVP* begins by computing $f(\langle C, x, u, v\rangle)$ for all $u, v \in \mathbf{F}$. For every non-zero $u \in \mathbf{F}$, there is a unique element $v_u \in \mathbf{F}$ such that $\langle C, x, u, v_u\rangle \in L$, and therefore $f(\langle C, x, u, v_u\rangle) \in S$. Since $2^m/p(n, m) \geq n$, there is at least one string $w \in S$ such that the number of u satisfying $f(\langle C, x, u, v_u\rangle) = w$ is at least n. Of course, there could be many such w (not necessarily in S), and we don't know which w is a string in S. To handle this, we will assume that every w that has $\geq n$ preimages is a string in S, and attempt to solve for the g_i's. As long as there is at least one $w \in S$ that has $\geq n$ preimages, one of the assumptions must be correct, and we will have the correct solution. Since we know the details of the circuit C, the solutions can be verified, and the incorrect ones weeded out.

Assume, therefore, wolog, that $w \in S$ has $\geq n$ preimages. Let $u_1, u_2, \ldots, u_n$ denote n of them, and let $v_1, v_2, \ldots, v_n$ denote the corresponding v_u's. The equations

$$1g_0 + u_j g_1 + u_j^2 g_2 + \ldots + u_j^{n-1} g_{n-1} = v_j, \qquad j = 1, 2, \ldots, n$$

form an inhomogeneous system of linear equations, where the coefficients u_j^i form a Vandermonde matrix. Since the u_j's are distinct elements of the field $\mathbf{F}$, the system has full rank.

It remains to show how to solve this system of equations in NC^1. While solving general linear equation systems seems to require NC^2, we will arrive at our NC^1 solution via closed formulae.

We omit the details in this extended abstract, but the following closed formula can be shown,

$$g_j = \sum_{i=1}^{n} (-1)^{1+i} \frac{v_i}{\prod_{k \neq i}(u_k - u_i)} P_{n-j-1}(u_1, \ldots, \widehat{u_i}, \ldots, u_n).$$

Here $\widehat{u_i}$ denotes that u_i is missing from the list $u_1, \ldots, u_n$, and P_k denotes the k-th elementary symmetric polynomial, defined as follows:

$$P_0(y_1, \ldots, y_\ell) = 1; \qquad P_k(y_1, \ldots, y_\ell) = \sum_{\substack{I \subseteq [\ell] \\ |I| = k}} \prod_{i \in I} y_i, \qquad k > 0.$$

By Facts (3) and (5), computing $v_i/(\prod_{k \neq i}(u_k - u_i))$ in NC^1 is fairly straightforward. Hence it suffices to show how to compute the polynomials $P_k(u_1, \ldots, \widehat{u_i}, \ldots, u_n)$, in logspace-uniform NC^1. A folklore theorem indicates that this can be done in non-uniform NC^1. For our application, however, the uniformity is crucial.

It is easy to see that for $y_1, \ldots, y_\ell \in \mathbf{F}$, $P_k(y_1, \ldots, y_\ell)$ equals $P_k(y_1, y_2, \ldots, y_\ell, 0, 0, \ldots, 0)$ for any number of extra zeroes. Let $r = |\mathbf{F}^*|$, the number of elements in the multiplicative group of $\mathbf{F}$. We will give an NC^1 algorithm to compute the elementary symmetric polynomial of r elements, not necessarily distinct, from the finite field $\mathbf{F}$. By appending $r - \ell$ zeroes, we can then compute $P_k(y_1, y_2, \ldots, y_\ell)$.

For $0 < k \leq r$, the value of the elementary symmetric polynomial $P_k(y_1, y_2, \ldots, y_r)$ is the coefficient of X^{r-k} in $h(X) \doteq \prod_{i=1}^{r}(X + y_i) - X^r$. Note that, given any $\alpha \in \mathbf{F}$, $h(\alpha)$ can be evaluated in NC^1, by Facts (2) and (5).

If we write $h(X)$ as $\sum_{i=0}^{r-1} a_i X^i$, the coefficient $a_i = P_{r-i}(y_1, \ldots, y_r)$ for $0 \leq i < r$. The idea now is to choose α's carefully from $\mathbf{F}$, compute $h(\alpha)$ and compute the coefficients a_i by interpolation. If we choose ω to be a primitive element of order r in $\mathbf{F}^*$, the powers of ω, namely $1 = \omega^0, \omega^1, \omega^2, \ldots, \omega^{r-1}$, run through the elements of $\mathbf{F}^*$. For $0 \leq i < r$, let $b_i = h(\omega^i)$. The relationship between the pointwise values (b_i's) and the coefficients (a_i's) of $h(X)$ can be written as:

$$\begin{pmatrix} b_0 \\ b_1 \\ \vdots \\ b_{r-1} \end{pmatrix} = \begin{pmatrix} 1 & \omega^0 & \omega^{0\cdot 2} & \ldots & \omega^{0\cdot(r-1)} \\ 1 & \omega^1 & \omega^{1\cdot 2} & \ldots & \omega^{1\cdot(r-1)} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & \omega^{r-1} & \omega^{(r-1)\cdot 2} & \ldots & \omega^{(r-1)\cdot(r-1)} \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_{r-1} \end{pmatrix}.$$

The matrix Ω is the Discrete Fourier Transform matrix, and is a Vandermonde matrix. Since the powers of ω are all distinct, Ω is invertible, and one can compute the coefficients a_i by $(a_0, \ldots, a_{r-1})^T = \Omega^{-1}(b_0, \ldots, b_{r-1})^T$. The crucial advantage over the earlier Vandermonde system is that with this particular choice of Ω, the matrix Ω^{-1} has a simple explicit form:

$$\Omega_{ij}^{-1} = 1/(\Omega_{ij}) = \omega^{-ij}.$$

Computing the coefficients of $h(X)$ is now simply a matrix-vector multiplication. Theorem 4 is proven. □

Corollary 5 *If there is a sparse set S that is hard for* P *under logspace-computable many-one reductions, then* P = LOGSPACE.

Corollary 6 *If there is a sparse set S that is hard for* P *under many-one reductions computable in logspace-uniform* NC^1*, then* P *equals logspace-uniform* NC^1*.*

Corollary 7 *If there is a set S with census function bounded by $2^{(\log n)^a}$ that is hard for* P *under many-one reductions computable in space $(\log n)^b$, then* $\mathrm{P} \subseteq \mathrm{DSPACE}[(\log n)^c]$*, where $c = \max\{a, b\}$.*

Acknowledgment

We thank Mitsu Ogihara for showing us his work in a Rochester-Buffalo joint complexity seminar. The resolution of Hartmanis' conjecture would not have been possible without the breakthrough of Ogihara. We also thank Allan Borodin for discussion of circuit complexity of elementary symmetric polynomials.

References

[AGHP90] N. Alon, O. Goldreich, J. Hastad, and R. Peralta. Simple constructions of almost k-wise independent random variables. In *Proc. 31st FOCS*, pages 544–553, 1990.

[BH77] L. Berman and J. Hartmanis. On isomorphisms and density of NP and other complete sets. *SIAM Journal on Computing*, 6:305–321, 1977. A preliminary version appeared in STOC 1976.

[BvzGH82] A. Borodin, J. von zur Gathen, and J. Hopcroft. Fast parallel matrix and GCD computations. *Information and Control*, 52:241–256, 1982.

[Chi85] A.L. Chistov. Fast parallel calculation of the rank of matrices over a field of arbitrary characteristic. In *Proc. 5th FCT*, Lecture Notes in Computer Science, pages 63–69. Springer-Verlag, 1985.

[Coo85] S. Cook. A taxonomy of problems with fast parallel algorithms. *Information and Control*, 64:2–22, 1985.

[For79] S. Fortune. A note on sparse complete sets. *SIAM Journal on Computing*, 8(3):431–433, 1979.

[GJ79] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Co., New York, 1979.

[Har78] J. Hartmanis. On log-tape isomorphisms of complete sets. *Theoretical Computer Science*, 7(3):273–286, 1978.

[HOT94] L. Hemachandra, M. Ogiwara, and S. Toda. Space-efficient recognition of sparse self-reducible languages. *Computational Complexity*, 4:262–296, 1994.

[HOW92] L. Hemachandra, M. Ogiwara, and O. Watanabe. How hard are sparse sets. In *Proc. 7th Structures*, pages 222–238, 1992.

[KL82] R. Karp and R. Lipton. Turing machines that take advice. *L'Enseignement Mathématique*, 28(2):191–209, 1982. A preliminary version appeared in STOC 1980.

[Lad75] R. Ladner. The circuit value problem is log space complete for P. *SIGACT News*, 7(1):18–20, 1975.

[Mah82] S. Mahaney. Sparse complete sets for NP: Solution of a conjecture of berman and hartmanis. *J. Comp. Sys. Sci.*, 25(2):130–143, 1982. A preliminary version appeared in FOCS 1980.

[Mul87] K. Mulmuley. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. *Combinatorica*, 7(1):101–104, 1987.

[NN90] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. In *Proc. 22nd STOC*, pages 213–223, 1990.

[Ogi95] M. Ogihara. Sparse hard sets for P yield space-efficient algorithms. Technical Report 569, Dept. of CS, University of Rochester, January 1995.

[OW91] M. Ogiwara and O. Watanabe. On polynomial-time bounded truth-table reducibility of NP sets to sparse sets. *SIAM Journal on Computing*, 20(3):471–483, 1991. A preliminary version appeared in STOC 1990.

[vL91] J.H. van Lint. *Introduction to Coding Theory.* Springer-Verlag, 1991.

[You92a] P. Young. How reductions to sparse sets collapse the polynomial-time hierarchy: A primer (Part I). *SIGACT News*, 23(3):107–117, 1992.

[You92b] P. Young. How reductions to sparse sets collapse the polynomial-time hierarchy: A primer (Part II). *SIGACT News*, 23(4):83–94, 1992.