

Sparse Hard Sets for P Yield Space-Efficient Algorithms

Mitsunori Ogiwara*
Department of Computer Science
University of Rochester
Rochester, NY 14627, USA

Abstract

In 1978, Hartmanis conjectured that there exist no sparse complete sets for P under logspace many-one reductions. In this paper, in support of the conjecture, it is shown that if P has sparse hard sets under logspace many-one reductions, then $P \subseteq \text{DSPACE}[\log^2 n]$.

1 Introduction

In 1978, Hartmanis [7] conjectured that no P-complete sets under logspace many-one reductions can be polynomially sparse; i.e., for any set $A \in P$ to which every set in P is logspace many-one reducible, the function that maps n to the number of elements in A of length up to n is not bounded by a polynomial in n . The conjecture is interesting and fascinating. If the conjecture is true, then $L \neq P$, because $L = P$ implies that every nonempty finite set is P-complete. So, since it is widely assumed that L is different from P, one might believe the validity of the conjecture. Such a reasoning may well be fallacious: proving this conjecture is *at least* as hard as proving $L \neq P$, and therefore, it may well be the case that even though $L \neq P$, P has polynomially sparse complete sets. In order to support the conjecture, one would perhaps need to show a result in the other direction; that is, if the conjecture *does not* hold some 'implausible' event occurs. Such an implausible event would be the collapse of P to a (presumably) much smaller class. As sets in P already have time-efficient recognition algorithms, this should mean that P has 'space-efficient' algorithms, e.g., P is included in $\text{DSPACE}[\log^k n]$ for some k .

The conjecture is reminiscent of the celebrated Berman-Hartmanis conjecture [2] that all NP-complete sets under polynomial-time many-one reduction are polynomially isomorphic. If the Berman-Hartmanis conjecture is true, then $P \neq \text{NP}$ and polynomially sparse sets cannot be NP-complete. A result to support this conjecture was obtained by Mahaney [10]. He showed that if there is a polynomially sparse hard set for NP, then $P = \text{NP}$; that is, unless NP collapses to the seemingly small class P, NP cannot have sparse complete sets.

*A.k.a Mitsunori Ogiwara.

In contrast with the “sparse hard set problem for NP,” not much work has been done on the Hartmanis’ conjecture on P—we could call it “the sparse hard set problem for P.” The only paper we are aware of is by Hemaspaandra, Ogihara, and Toda [8], who prove that P cannot have *poly-logarithmic* sparse hard sets unless P is included in SC, the class of sets recognized simultaneously in polynomial-time and in poly-logarithmic space. As one can easily see, the result is still too weak because the poly-logarithmic sparsity is a much more stringent condition than polynomial sparsity.

In this paper, we give the first solution to the sparse hard set problem for P by showing that unless $P \subseteq \text{DSPACE}[\log^2 n]$, the Hartmanis’ conjecture holds.

Theorem 1 *There exist no sparse P-hard sets under logspace many-one reductions unless $P \subseteq \text{DSPACE}[\log^2 n]$.*

Let us say a few words about the proof. Assuming the existence of a sparse P-hard set, we are able to reduce, in a space-efficient manner, any instance of the circuit value problem to a circuit value problem of a circuit consisting exclusively of ‘parity’ gates. However, this restricted circuit value problem is known to be in $\text{DSPACE}[\log^2 n]$.

Readers familiar with the class $\oplus\text{L}$ [3], a logarithmic space-bounded version of $\oplus\text{P}$ [11, 6], will recognize our reduction to map an instance of the circuit value problem to a problem in $\oplus\text{L}$, and recall that $\oplus\text{L} \subseteq \text{DSPACE}[\log^2 n]$ (see [3] and [1]).

The paper is organized as follows. In section 2, we will define the basic notation and the circuit value problems. In section 3, we prove our main theorem.

2 Circuit Value Problems

A Boolean circuit is a directed acyclic graph C with labeled nodes. Nodes in C with indegree 0 are called *input gates* while the other gates are called *interior gates*. Input gates in C have distinct labels from $\{1, \dots, n\}$, where n is the number of input gates in C . There is one designated node in C with outdegree 0, which is called the *output gate*. Each interior gate is labeled by a Boolean function chosen from $\{\neg, \wedge, \vee\}$. A gate labeled by \neg is called a *NOT* gate and has indegree 1. A gate labeled by \wedge (or \vee) is called an *AND* gate (an *OR* gate, respectively) and has indegree ≥ 2 . A gate g is said to be a *direct input* to a gate g' if there is an arc from g to g' in C .

A Boolean circuit is said to be of *bounded fan-in* if every gate has indegree ≤ 2 . It is said to be of *unbounded fan-in* if some gate may have indegree > 2 . A Boolean circuit is encoded by its adjacency matrix and the labels of the gates, where we always assume that the output gate is the last node and for every i , the i -th input gate is the i -th node.

Let C be a Boolean circuit of m gates and n inputs and let $x = x_1 \cdots x_n \in \{0, 1\}^n$. For each $i, 1 \leq i \leq m$, let g_i denote the i -th gate in C . For $i, 1 \leq i \leq m$, the output of g_i in C on input x , denoted by $C[x, i]$, is determined inductively as follows:

- If g_i is an input gate labeled by j , then $C[x, i] = x_j$.
- If g_i is a NOT gate whose unique direct input is g_j , then $C[x, i] = \neg(C[x, j])$.

- If g_i is an AND gate and its direct inputs are g_{j_1}, \dots, g_{j_k} , then $C[x, i] = C[x, j_1] \wedge \dots \wedge C[x, j_k]$.
- If g_i is an OR gate and its direct inputs are g_{j_1}, \dots, g_{j_k} , then $C[x, i] = C[x, j_1] \vee \dots \vee C[x, j_k]$.

The output of C on input x , denoted by $C(x)$, is $C[x, m]$.

The circuit value problem (CVP) is the problem of deciding whether a bounded fan-in Boolean circuit C outputs 1 on input x . Ladner [9] showed that CVP is complete for P under logspace many-one reductions. A circuit C is *topologically sorted* if for every i, j , if g_i is a direct input gate of g_j , then $i < j$. One can easily observe that the construction by Ladner can be used to show that the topologically sorted version of the problem, TSCVP, is complete under logspace many-one reductions. We will identify TSCVP with the set of all strings $C\#x$ such that C is a topologically sorted Boolean circuit of n inputs, $|x| = n$, and $C(x) = 1$.

Parity function, denoted by \oplus , is one that maps a binary string to the parity of the number of 1's in it. We also view \oplus as the function that maps a natural number n to $(n \bmod 2)$. A parity (or, exclusive-or) gate is a gate of unbounded fan-in that, given binary bits a_1, \dots, a_n as inputs, computes $\oplus(a_1 \dots a_n)$. By convention, we will use both $\oplus(a_1, \dots, a_n)$ and $a_1 \oplus \dots \oplus a_n$ to denote $\oplus(a_1 \dots a_n)$. A parity circuit is an unbounded fan-in circuit in which all the gates compute \oplus . The *parity circuit problem* is defined as a variation of the circuit value problem, in which it is asked whether a parity circuit outputs 1 on a specified input. We define Parity-CVP to be the set corresponding to the problem: the set of all $C\#x$ such that C is a parity circuit and, on input x , outputs 1.

A set L is in $\oplus L$ [3] if there exists a logarithmic space-bounded nondeterministic Turing machine N such that for every x , $x \in L$ if and only if the number of accepting computation paths of N on x is odd.

Proposition 2 Parity-CVP is in $\oplus L$.

Proof Let C be a parity circuit of m gates g_1, \dots, g_m and n input gates and let $x = x_1 \dots x_n$. Note for any $a, b, c \in \{0, 1\}$, that $\oplus(a, b, c) = \oplus(\oplus(a, b), c) = \oplus(a, \oplus(b, c))$.

For each $i, 1 \leq i \leq m$, let $\mu(i)$ denote the number of paths in C on x from some input gate with value 1 to the gate g_i . We claim for any $i, 1 \leq i \leq m$, that $C[x, i] = \oplus(\mu(i))$. This is proven by induction.

For the base case, let g_i be an input gate. Then $C[x, i] = 1$ if and only if $x_i = 1$. Trivially, there is exactly one path from the gate to itself. So, the claim holds.

For the induction step, let g_i be an interior gate and let g_{h_1}, \dots, g_{h_l} be an enumeration of all direct inputs to g_i . Clearly, $\mu(i) = \sum_{j=1}^l \mu(h_j)$. Suppose that the claim holds for h_1, \dots, h_l , i.e., $C[x, h_j] = \oplus(\mu(h_j))$ for all $j, 1 \leq j \leq l$. By definition, $C[x, i] = \oplus(C[x, h_1] + \dots + C[x, h_l])$. So, $C[x, i] = \oplus(\sum_{j=1}^l \oplus(\mu(h_j))) = \oplus(\sum_{j=1}^l \mu(h_j)) = \oplus(\mu(i))$. Thus, the claim holds for g_i . Hence the claim holds for every gate.

Now, noting that Boolean circuits are acyclic, it is easy to construct a nondeterministic machine witnessing Parity-CVP $\in \oplus L$. Our machine, on input $C\#x$, guesses a sequence

g_{i_1}, \dots, g_{i_h} of at most m gates and accepts if and only if the sequence is a path from an input gate outputting 1 to the output gate. The verification can be done sequentially, so the machine has only to store consecutive two elements in the sequence. So, it can be logarithmic space-bounded. Clearly, the number of accepting computation paths of the machine on $C\#x$ is equal to the number of paths in C on x from some input gate with value 1 to the output gate. So, the machine witnesses that Parity-CVP $\in \oplus\text{L}$. This proves the proposition. \square

Proposition 3 [3, 1] $\oplus\text{L} \subseteq \text{DSPACE}[\log^2 n]$.

Here we provide a sketch of the proof, which is reminiscent of Savitch's theorem [12]. Let N be a logarithmic space-bounded nondeterministic machine N witnessing that $L \in \oplus\text{L}$ and x be an input to N . For two IDs I and J of N on x and a natural number t , define $Q(I, J, t)$ to be the parity of the number of computation paths of N on x from I to J of length at most 2^t . Define $Q(I, I, t) = 1$ for every I and t . Let $m = \mathcal{O}(\log |x|)$ be a natural number such that N on x runs for at most 2^m steps and let I_{ini} be the unique start ID of N on x . We may assume that there is a unique accepting ID of N on x . Let I_{acc} denote this ID. Clearly, $x \in L$ if and only if $Q(I_{ini}, I_{acc}, m) = 1$. Note for every I, J , and $t > 0$, that

$$Q(I, J, t) = \bigoplus \left(\sum_K Q(I, K, t-1) \cdot Q(K, J, t-1) \right),$$

where K ranges over all IDs of N on x . This suggests the following recursive procedure to evaluate $Q(I, J, t)$:

- If $I = J$, then return 1.
- If $I \neq J$ and $t = 0$, then compute and return $Q(I, J, 0)$ by simulating one move of N on x at ID I .
- If $I \neq J$ and $t > 0$, then set c to 0 and for each K , set c to $(c + Q(I, K, t-1) \cdot Q(K, J, t-1))$ modulo 2.

If we run this procedure to evaluate $Q(I_{ini}, I_{acc}, m)$, then the recursion depth is $m = \mathcal{O}(\log |x|)$. Since each ID is encoded as a string of length $\mathcal{O}(\log |x|)$, the evaluation requires $\mathcal{O}(\log^2 |x|)$ space, and thus, $L \in \text{DSPACE}[\log^2 n]$.

From the above two propositions, we immediately obtain the following:

Proposition 4 Parity-CVP $\in \text{DSPACE}[\log^2 n]$.

3 Proof of Theorem 1

We repeat the statement of the theorem.

Theorem 1 *There exist no sparse P-hard sets under logspace many-one reductions unless $\text{P} \subseteq \text{DSPACE}[\log^2 n]$.*

Suppose that there exists a sparse P-hard set under logspace many-one reductions. Then we show that $P \subseteq \text{DSPACE}[\log^2 n]$, in particular, TSCVP is in $\text{DSPACE}[\log^2 n]$.

We will make use of the following set A : A is the set of all strings of the form $C\#x\#I\#b$ such that

- $C\#x$ is an instance of TSCVP, i.e., C is a topologically sorted Boolean circuit with m gates and n inputs and $x \in \{0, 1\}^n$,
- I is a nonempty subset of $\{1, \dots, m\}$ encoded as the enumeration of its elements in increasing order,
- $b \in \{0, 1\}$, and
- $\bigoplus_{i \in I} C[x, i] = b$.

Clearly, $A \in P$. So, by our supposition, A is logspace many-one reducible to a sparse set S via some function f . Note that for a sufficiently large m and every legitimate $C\#x\#I\#b$, it holds that $|C\#x\#I\#b| \leq 2|C\#x|$. Since S is sparse, this implies that for every $C\#x$, the number of $y \in S$ such that $y = f(C\#x\#I\#b)$ for some I and b is bounded by $2^{d \log |C\#x|}$ for some constant d .

Let $C\#x$ be fixed, whose membership in TSCVP we are testing. Let g_1, \dots, g_m be the gates of C , where g_1, \dots, g_n are the input gates and g_m is the output gate. Let $\ell = |C\#x|$ and $e = \lceil d \log \ell \rceil$. As we have already fixed C and x , we will simply use $I\#b$ to denote $C\#x\#I\#b$ by dropping C and x . By the above observation, the number of $y \in S$ such that $y = f(I\#b)$ for some $I \in \{1, \dots, m\}$ and $b \in \{0, 1\}$ is less than 2^e .

Now we introduce the notion of good gates and bad gates. Let \mathcal{T} be the set of all nonempty subsets of $\{1, \dots, m\}$ of size at most e . Let $i \in \{n+1, \dots, m\}$. We say that g_i is *good* if there exist $I, J \in \mathcal{T}$ and $b, c \in \{0, 1\}$ such that

$$f(I\#b) = f(J\#c) \text{ and } i = \max(I \Delta J), \quad (1)$$

where $I \Delta J$ denotes the symmetric difference of I and J . Otherwise, g_i is called *bad*. Intuitively, an interior gate g_i is good if we can easily find a set of gates g_j, \dots, g_k such that the parity of the output of these gates is equal to the output of g_i , and thus, the evaluation of g_i can be reduced to the evaluation of g_j, \dots, g_k .

The outline of the main steps of the proof are as follows: (1) First, we show that there are very few bad gates, (2) then construct a parity circuit D whose inputs are x and the bad gates and whose interior gates are good gates, and (3) then show that for some assignment of values to the bad gates in D , the value of each gate in D is equal to the value of the corresponding gate in C . Then, use the fact that D can be computed in polylog space.

Claim 1 *The number of bad gates is at most e .*

Proof Assume that there are $e+1$ bad gates and let $g_{h_1}, \dots, g_{h_{e+1}}$ be an enumeration of $e+1$ bad gates. Let \mathcal{R} be the set of all nonempty subsets of $\{h_1, \dots, h_{e+1}\}$ of size at most

e. Note for any $I \in \mathcal{R}$, that exactly one of $f(I\#0)$ or $f(I\#1)$ is in S because exactly one of $I\#0$ or $I\#1$ is in A . So, let b_I be the unique $b \in \{0, 1\}$ such that $f(I\#b) \in S$. Note also, for any distinct $I, J \in \mathcal{R}$, that $f(I\#b_I) \neq f(J\#b_J)$. Otherwise, g_k with $k = \max(I \Delta J)$, which is bad by our assumption, is good, a contradiction. Since there are $2^{e+1} - 2 \geq 2^e$ elements in \mathcal{R} , we can collect 2^e elements in S , which contradicts the assumption that there are less than 2^e elements in S we see as the image of f . This proves the claim. \square

Now let g_{h_1}, \dots, g_{h_q} be the enumeration of all bad gates and let $H = \{h_1, \dots, h_q\}$, where $q \leq e$. For each good g_i , let $(I(i), b(i), J(i), c(i))$ be the lexicographically minimum $(I\#b, J\#b)$ witnessing that g_i is good. We define a parity circuit D with $m + 1$ gates and $n + q + 1$ input gates as follows:

- The gates of D are those of C plus one new gate g_0 .
- The input gates of D are g_0 , the input gates of C , and the bad gates; that is, they are $g_0, g_1, \dots, g_n, g_{h_1}, \dots, g_{h_q}$. We will fix the input to g_0 to 1.
- Each interior gate g_i in D computes the parity function, whose direct inputs are given as follows:
 - If $b(i) = c(i)$, then all g_j with $j \in (I(i) \Delta J(i)) - \{i\}$.
 - If $b(i) \neq c(i)$, then all g_j with $j \in (I(i) \Delta J(i)) - \{i\}$ plus g_0 .

Note that D is topologically sorted since C is topologically sorted and if g_i is good then i is the largest in $I(i) \Delta J(i)$.

We say that $v \in \{0, 1\}^q$ is valid if the value assigned by v to each bad gate is equal to the value of the bad gate in $C\#x$; i.e., for all $t, 1 \leq t \leq q$, the t -th bit of v is equal to $C[x, h_t]$. It is obvious that there is a unique valid v .

Claim 2 v is valid if and only if for every gate $g_i, i, 1 \leq i \leq m, C[x, i] = D[1xv, i]$

Proof The implication from right to left is obvious. The other direction is proven inductively. First, note that $C[x, i] = D[1xv, i]$ holds for every bad gate g_i . Next, let g_i be a good gate and suppose that the claim holds for every direct input g_j of g_i in D . We have $f(I(i)\#b(i)) = f(J(i)\#c(i))$. So,

$$b(i) = \bigoplus_{j \in I(i)} C[x, j] \Leftrightarrow c(i) = \bigoplus_{j \in J(i)} C[x, j].$$

This implies

$$C[x, i] = C[x, j_1] \oplus \dots \oplus C[x, j_k] \oplus b(i) \oplus c(i),$$

where j_1, \dots, j_k is an enumeration of all j such that $j \neq i$ and $j \in I(i) \Delta J(i)$. By our supposition, for each $t, 1 \leq t \leq k, D[1xv, j_t] = C[x, j_t]$. Also, by definition, g_0 is among the direct inputs of g_i if and only if $b(i) \neq c(i)$, i.e., $b(i) \oplus c(i) = 1$. Thus, $C[x, i] = D[1xv, i]$. Hence, the claim holds for g_i . This proves the claim. \square

For each $v \in \{0, 1\}^q$ and $t, 1 \leq t \leq q$, we say that v is *correct* at t if, depending on the type of g_{h_t} in C , the following conditions are satisfied:

- (a) If g_{h_t} is a NOT gate in C with g_j as its direct input, g_j , then v_t is equal to $\neg(D[1xv, j])$.
- (b) If g_{h_t} is an AND gate in C with g_j and g_k as its direct inputs, then v_t is equal to $D[1xv, j] \wedge D[1xv, k]$.
- (c) If g_{h_t} is an OR gate in C with g_j and g_k as its direct inputs, then v_t is equal to $D[1xv, j] \vee D[1xv, k]$.

Claim 3 v is valid if and only if for all $t, 1 \leq t \leq q$, v is correct at t .

Proof The implication from left to right is obvious. To prove the other direction, suppose that v is not valid. Let t be the smallest i such that the i -th bit of v is not equal to the output of the gate of C on input x ; i.e., t is the smallest $i, 1 \leq i \leq t$, such that $v_i = D[1xv, j_i] \neq C[x, j_i]$. Since D is topologically sorted, by an argument similar to that in the proof of Claim 2, we have $D[1xv, k] = C[x, k]$ for all $k < j_t$. If v is correct at t , then v_t is equal to $C[x, j_t]$, a contradiction. So, v is not correct at t . \square

The above claims suggest the following algorithm to reduce C to D with the unique valid v .

Step 1: For each interior gate of C , test whether it is good, and construct H , the set of all bad gates.

Step 2: For each $v \in \{0, 1\}^q$, test whether v is valid by testing whether v is correct at all t , and if so, use the valid v to compute $D[1xv, m]$.

Claim 4 *The algorithm can be executed in $\mathcal{O}(\log^2 \ell)$ space.*

Proof Let M be a logspace machine that computes f . Note that $I \in \mathcal{T}$ is encoded as a string of length $\mathcal{O}(e \log m) = \mathcal{O}(\log^2 \ell)$. Given $I\#b$ and $J\#c$, testing whether $f(I\#b) = f(J\#c)$ can be done by simulating M on $I\#b$ and M on $J\#c$ simultaneously to compare $f(I\#b)$ and $f(J\#c)$ bit-by-bit. Since M 's output tape is certainly write-only, the comparison requires to store only the most recent output bit from each. More precisely, M on $I\#b$ and M on $J\#c$ are simulated alternatively step-by-step. If one of the simulations outputs a new bit of f , then it is suspended until the other simulation produces a new bit of f or halts without outputting a new bit. If both produce new bits, then the bits are compared and, if they are different, it must be the case that the values of f are different. So, the comparison is terminated. If only one simulation produces a new bit, then the two values of f obviously have different length, so, the values are different. So, the comparison is terminated. If both simulations halt without producing any new bits, then since the bits that have been produced so far are the same, it must be the case that they have the same value. The amount of space expended by the simulations is $\mathcal{O}(\log^2 \ell)$, the amount required to store $I\#b$ and $J\#c$, since M is logarithmic space-bounded.

In order to test whether an interior gate g_i is good, and if so, to compute $I(i), b(i), J(i), c(i)$, it suffices to test, by cycling through all possible $(I\#b, J\#c)$ in the

lexicographic increasing order, whether $(I\#b, J\#c)$ witnesses that g_i is good. By the previous discussion, the amount of space required is $\mathcal{O}(\log^2 \ell)$. There are at most e bad gates, so the amount of space required to store H , the set of all bad gates, is $\mathcal{O}(e \log m) = \mathcal{O}(\log^2 \ell)$, so, H can be computed in space $\mathcal{O}(\log^2 \ell)$.

Note that, as we are developing an $\mathcal{O}(\log^2 n)$ algorithm, there is not enough space to store the entire description of D . However, after obtaining H , each bit of the description of D is computable in $\mathcal{O}(\log^2 \ell)$ space as follows: In order to determine the direct inputs to g_i , if either $i \leq n$ or $i \in H$, then g_i is an input gate of D , and so, has no direct inputs; otherwise, $I(i), b(i), J(i), c(i)$, which are computable in $\mathcal{O}(\log^2 \ell)$ space, provide the list of direct inputs.

In order to test whether v is correct at t , since h_t is the t -th element in H and the type of g_{h_t} in C and its direct input(s) are determined from $C\#x$, it suffices to compute $D[1xv, j]$ for j such that g_j is a direct input to g_{h_t} in C . Since D is a parity circuit, the computation problem is solvable by Parity-CVP. Recall that we demand that the last gate of a circuit be the output. So, let D_j be the circuit constructed from D by making the connection of g_m identical to that of g_j . Then $D_j(1xv) = D_j[1xv, m] = D[1xv, j]$. Let N be a deterministic Turing machine that decides Parity-CVP in $\mathcal{O}(\log^2 n)$ space. Since each bit of the description of D is computable in $\mathcal{O}(\log^2 \ell)$ space, given j , each bit of the description of D_j is computable in the same amount of space. Thus, one can simulate N on $D_j\#1xv$ by keeping track of the position of N 's input head. When N needs to read the k -th bit of its input, one has only to activate the algorithm to produce D to compute the k -th bit (by recording the number of bits produced so far and the current bit), where the bits for the m -th gate are computed from those for the j -th gate. Thus, $D[1xv, j]$ is computable in $\mathcal{O}(\log^2 \ell)$ space, and therefore, whether v is valid can be tested in the same amount of space.

Once the valid v is discovered, since it is of length at most e , there is enough space to record it. Now we have only to compute $C[x, m]$ as $D[1xv, m]$ with the valid v . Again, we have only to simulate N while computing the bits of D on demand, which requires $\mathcal{O}(\log^2 \ell)$ space. Hence, the whole process can be done in $\mathcal{O}(\log^2 \ell)$ space. This proves the claim. \square

This completes the proof of the theorem.

By a straightforward generalization of the proof, we obtain the following theorem.

Theorem 2 *Let $d, e \geq 1$ and let S be a set whose density function is bounded by $2^{\mathcal{O}(\log^d n)}$. Suppose every set in P is many-one reducible to S via a function f computable in $\mathcal{O}(\log^e n)$ space. Then $\mathsf{P} \subseteq \text{DSPACE}[\log^{de+1} n]$.*

4 Conclusion

We have given a solution to Hartmanis' conjecture on sparse complete sets for P by showing that P cannot not have many-one hard sets of low density via space-efficient reductions unless $\mathsf{P} \subseteq \text{DSPACE}[\log^2 n]$. We note here that, by extending the technique in this paper, Cai and Sivakumar have recently resolved the conjecture by showing that sparse P -hard sets

exist under logspace many-one reductions if and only if $P = L$ [5]. The technique has been further extended to study the sparse P-hard set problem for more flexible reducibilities [4, 13]. A very interesting open question in this regard is whether P having sparse hard sets under logspace Turing reductions collapses P.

Acknowledgment

The author would like to thank Eric Allender, Jin-yi Cai, Lane Hemaspaandra, Ioan Macarie, D. Sivakumar, and Marius Zimand for enjoyable discussions.

References

- [1] C. Àlvarez and B. Jenner. A very hard log-space counting class. *Theoretical Computer Science*, 107:3–30, 1993.
- [2] L. Berman and J. Hartmanis. On isomorphisms and density of NP and other complete sets. *SIAM Journal on Computing*, 6(2):305–322, 1977.
- [3] G. Buntrock, C. Damm, U. Hertrampf, and C. Meinel. Structure and importance of Logspace-MOD class. *Mathematical Systems Theory*, 25:223–237, 1992.
- [4] J. Cai, A. Naik, and D. Sivakumar. On the existence of hard sparse sets under weak reductions. Technical Report 95-31, Department of Computer Science, State University of New York at Buffalo, Buffalo, NY, July 1995.
- [5] J. Cai and D. Sivakumar. The resolution of a Hartmanis conjecture. In *Proceedings of the 36th Symposium on Foundations of Computer Science*, pages 362–371. IEEE Computer Society Press, 1995.
- [6] L. Goldschlager and I. Parberry. On the construction of parallel computers from various bases of boolean functions. *Theoretical Computer Science*, 43:43–58, 1986.
- [7] J. Hartmanis. On log-tape isomorphisms of complete sets. *Theoretical Computer Science*, 7(3):273–286, 1978.
- [8] L. Hemachandra, M. Ogiwara, and S. Toda. Space-efficient recognition of sparse self-reducible languages. *Computational Complexity*, 4:262–296, 1994.
- [9] R. Ladner. The circuit value problem is log space complete for P. *SIGACT News*, 7(1):18–20, 1975.
- [10] S. Mahaney. Sparse complete sets for NP: Solution of a conjecture of Berman and Hartmanis. *Journal of Computer and System Science*, 25(2):130–143, 1982.

- [11] C. Papadimitriou and S. Zachos. Two remarks on the power of counting. In *Proceedings of the 6th GI Conference on Theoretical Computer Science*, pages 269–276. Springer-Verlag *Lecture Notes in Computer Science* #145, 1983.
- [12] W. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Science*, 4:177–192, 1970.
- [13] D. van Melkebeek. On reductions of P sets to sparse sets. Technical Report TR95-06, Department of Computer Science, University of Chicago, Chicago, IL, August 1995.