Lower Bounds for Nondeterministic Semantic Read-Once Branching Programs[*]

Stephen Cook[1], Jeff Edmonds[2], Venkatesh Medabalimi[1], and Toniann Pitassi[1]

1   University of Toronto, Computer Science Department, Toronto, ON, Canada
     sacook@cs.toronto.edu
2   York University, Computer Science Department, Toronto, ON, Canada
     jeff@cs.york.ca
3   University of Toronto, Computer Science Department, Toronto, ON, Canada
     venkatm@cs.toronto.edu
4   University of Toronto, Computer Science Department, Toronto, ON, Canada
     toni@cs.toronto.edu

──── **Abstract** ────

We prove exponential lower bounds on the size of semantic read-once 3-ary nondeterministic branching programs. Prior to our result the best that was known was for D-ary branching programs with $|D| \geq 2^{13}$.

1   Introduction

A major question in complexity theory is whether polynomial-time is the same as log-space or nondeterministic log-space. One approach to this problem is to study time/space tradeoffs for problems in P. For example, for natural problems in P, does the addition of a space restriction prevent a polynomial time solution? In the uniform setting, time-space tradeoffs for SAT were achieved in a series of papers [7, 15, 8, 9]. Fortnow-Lipton-Viglas-Van Melkebeek [9] shows that any algorithm for SAT running in space $n^{o(1)}$ requires time at least $\Omega(n^{\phi-\epsilon})$ where ϕ is the golden ratio $((\sqrt{5}+1)/2)$ and $\epsilon > 0$. Subsequent works [18, 6] improved the time lower bound to greater than $n^{1.759}$.

In the nonuniform setting, the standard model for studying time/space tradeoffs is the *branching program*. In this model, a program for computing a function $f(x_1, \ldots, x_n)$ (where the variables take values from a finite domain D) is represented by a directed acyclic graph with a unique source node called the start node. Each nonsink node is labelled by a variable and the edges out of a node correspond to the possible values of the variable. Each sink node is labelled by an output value. For Boolean functions, there is one sink node called the accept node (or 1-node), and all other sink nodes are rejecting nodes. Executing the program on an input corresponds to following a path from the start node, using the values of the input variables to determine which edges to follow. The output of the program is the value labelling the sink node reached. A D-ary branching program is deterministic if each non-sink node has exactly D edges, one for every value in D.

The length of a branching program is the number of edges in the longest path. It is clear that length of a branching program can be seen as a measure of computation time. The size of a branching program is the number of nodes in the program. For a boolean function f_n of n variables, let $BP(f_n)$ denote the minimum size of a branching program computing f_n. $BP(f_n)$ is closely related to the space complexity $S(f_n)$ of a non-uniform Turing machine computing f_n: $S(f_n) = O(\log(\max\{BP(f_n), n\}))$ and $BP(f_n) = 2^{O(\max\{S(f_n), \log n\})}$ [5, 17]. This motivates the study of branching program size lower bounds. In particular, size lower bounds on length restricted branching programs translate to time/space tradeoffs.

The state of the art time/space tradeoffs for branching programs were proven in the remarkable papers by Ajtai [1] and Beame-et-al [3]. In the first paper, Ajtai exhibited a polynomial-time computable Boolean function such that any sub-exponential size determinestic branching program requires superlinear length. This result was significantly improved and extended by Beame-et-al who showed that any sub-exponential size randomized branching program requires length $\Omega(n\frac{\log n}{\log \log n})$.

Lower bounds for nondeterministic branching programs have been more difficult to obtain. In this model, there can be several arcs (or no arcs) out of a node with the same value for the variable associated with the node. An input is accepted if there exists at least one path consistent with the input from the source to the 1-node. A nondeterministic branching program computes a function f if its accepted inputs are exactly equal to $f^{-1}(1)$. From here on, we shall restrict our attention to non-deterministic branching programs.

Length-restricted nondeterministic branching programs come in two flavors: *syntactic* and *semantic*. A length l syntactic model requires that every path in the branching program has length at most l, and similarly a read-k syntactic model requires that every path in the branching program reads every variable at most k times. In the less restricted semantic model, the requirement is only for *consistent* accepting paths from the source to the 1-node; that is, accepting paths along which no two tests $x_i = d_1$ and $x_i = d_2$, $d_1 \neq d_2$ are made. This is equivalent to requiring that for every accepting path, each variable is read at most k times. Thus for a nondeterminsitic read-k semantic branching program, the overall length of the program can be unbounded.

Note that any syntactic read-once branching program is also a semantic read-once branching program, but the the opposite direction does not hold. In fact, Jukna [11] proved that semantic read-once branching programs are exponentially more powerful than syntactic read-once branching programs, via the "Exact Perfect Matching"(EPM) problem. The input is a (Boolean) matrix A, and A is accepted if and only if every row and column of A has exactly one 1 and rest of the entries are 0's i.e if it's a permutation matrix. Jukna gave a polynomial-size semantic read-once branching program for EPM, while it was known that syntactic read-once branching programs require exponential size [14, 13].

Lower bounds for syntactic read-k (nondeterministic) branching programs have been known for some time [16, 4]. However, for *semantic* nondeterministic branching programs, even for read-once, no lower bounds are known for polynomial time computable functions for the $|D| = 2$ case. The best lower bound known prior to our work is an exponential lower bound for semantic read-once (nondeterministic) $|D|$-way branching programs, where $|D| = 2^{13}$ [10]. In fact this lower bound actually holds more generally for semantic read-k but where $|D| = 2^{3k+10}$.

Jukna obtains his result by showing that any time restricted semantic branching program of small size has a large rectangle in $f^{-1}(1)$. He uses the explicit function of computing the characteristic function of a linear code having minimum distance $m + 1$ defined over $GF(q)$. Given a parity matrix Y, the function $g(Y, x) = 1$ iff x is a codeword. Since codewords in a

linear code of minimum distance $m + 1$ can only have an m-rectangle of size 1 he argues that a time restricted branching program of length kn computing g requires a size of $2^{\Omega(n/k^2 4^k)}$. This exponential lower bound can be obtained whenever D is sufficiently large in comparison to k, specifically for $|D| = q \geq 2^{3k+10}$.

Jukna's result is an improvement over exponential lower bounds with a domain requirement of $2^{2^{ck}}$ obtained in [2]. Beame et.al [2] obtain their result by characterizing the function computed by a time restricted branching program of small size as a union of shallow decision forests where the size of the union depends on the size of the branching program. Each shallow forest is then shown to be representable by a collection of *small* number of βn-pseudo-rectangles in $f^{-1}(1)$. (Pseudo-rectangles are a generalization of what we call embedded rectangles later). This gives a representation of the branching program as a union of small (in the size 's') number of βn-pseudo-rectangles. Now, if for some function f the maximum size of a βn-pseudo-rectangle is $|D|^{(1-\psi_f(\beta))n}$ and the number of *yes*-instances $|f^{-1}(1)| \geq |D|^{(1-\eta(f))n}$ then the number of βn-pseudo-rectangles will be at least $|D|^{(\psi_f(\beta)-\eta(f))n}$. This yields an exponential lower bound on s for sufficiently large $|D|$ whenever $(\psi_f(\beta)-\eta(f))$ is bounded away from 0 by some $\epsilon > 0$. They then exhibit an explicit function with this property. Their function $QF_M : GF(q^n) \to \{0,1\}$ is based on quadratic forms using a modified Generalized Fourier Transform matrix. They show that there exists a constant $c > 0$ such that for all k and $\epsilon \in (0,1)$, if $D \geq 2^{2^{\frac{c}{\epsilon}k}}$ then a non-deterministic BP of length kn computing QF_M needs size at least $S = 2^{n \log^{1-\epsilon}|D|}$. For the specific case of $k = 1$, it can be shown that if their analysis of maximum size of βn-pseudo-rectangles in QF_M is tight, a domain size of at least $|D| \geq 2^{64}$ is needed.

Our main result is an exponential lower bound on the size of semantic read-once non-deterministic branching programs for a polynomial time decision problem f for 3-ary inputs. Similar in spirit to these previous results [10, 2] we show that a small sized semantic read once branching program is bound to have a large rectangle in $f^{-1}(1)$.

($\star$) However in addition, we show that one can always find a *balanced* rectangle in $f^{-1}(1)$ of size r^2 where r is some large constant.

A balanced rectangle is one which is reasonably close to being a square.

The particular polynomial time decision problem we use to prove the lower bound is: to decide if a polynomial over a finite field K evaluates to a value less than a certain threshold at a given input. The input is a pair (u, x) where u is the description of a degree $d-1$ polynomial over $[K]$ and $x \in [K]$, and we want to accept if and only if $u(x) < K^{1-\delta}$. We actually prove a stronger theorem: with high probability over all polynomials u, any nondeterministic semantic read-once branching program for what we shall call $Poly_u$ (along with a hyperplane constraint) requires exponential size. That is, even if the branching program knows the polynomial u, for a typical u it cannot efficiently do polynomial evaluation. The main properties of polynomials over finite fields we are using are polynomial interpolation, and Lemma 7, which might be interpreted to mean something like: the spread of values of a typical random polynomial of degree d over a field K is roughly close to being uniform over K, provided K is sufficiently large.

Continuing with the above observation ($\star$) that we can find a balanced rectangle in $f^{-1}(1)$ for a function with a small semantic read once branching program, since the number of balanced rectangles of a certain size $d = r^2$ is *small* and since each one of them can be a rectangle in $f^{-1}(1)$ for a relatively *small* number of degree d polynomials over K as a consequence of polynomial interpolation, we argue that there must be a polynomial with no balanced rectangle of this size in $f^{-1}(1)$ and hence the branching program computing it

should be *large.* A key idea of this argument is that for a balanced rectangle the sum of the lengths of the rectangle can be at most a *small* fraction of its area.

By a simple padding argument, we can modify our problem $Poly_u$ and actually achieve the lower bound for domain size $2 + \epsilon$ for arbitrarily small $\epsilon > 0$. In this model, we can define the problem to have $N = n + M$ variables, $M = \Theta(N)$ of them with domain size 3 and the rest, with domain size 2, do not affect the output. In section 5, we show why it might be harder to prove lower bounds for semantic read-once branching programs when $|D| = 2$ by showing how these branching programs can altogether evade having an exponential number of states in many purported choices of bottleneck layer by giving polynomial upper bounds.

2   Definitions

Throughout this article, D denotes a finite set. For finite set N, D^N is the set of maps from N to D. An element of N is called a *variable index* or simply an *index.* We normally take N to be $[n]$ for some integer n, and write D^N for $D^{[n]}$. If $A \subseteq N$, a point $\sigma \in D^A$ is a *partial input* on A. For a partial input σ, $fixed(\sigma)$ denotes the index set A on which it is defined and $unfixed(\sigma)$ denote the set $N - A$. If σ and π are partial inputs with $fixed(\sigma) \cap fixed(\pi) = \emptyset$, then $\sigma\pi$ denote the partial input on $fixed(\sigma) \cup fixed(\pi)$ that agrees with σ on $fixed(\sigma)$ and with π on $fixed(\pi)$.

For $x \in D^N$ and $A \subseteq N$, the *projection* x_A of x onto A is the partial input on A that agrees with x. For $S \subseteq D^N$, $S_A = \{x_A \mid x \in S\}$.

2.1   Nondeterministic Semantic Read-Once Branching Programs

Let $f : D^N \to \{0, 1\}$ be a boolean function whose input is given in $|D|$-ary. Let the input variables be $x_1, \ldots, x_n$ where $x_i \in D$ for all $i \leq n$. A $|D|$-way nondeterministic branching program (for f) is an acyclic directed graph G with a distinguished source node q_{start} and a distinguished sink node (the accept node) q_{accept}. We refer to nodes as *states.* Each non-sink state is labeled with some input variable x_i, and each edge directed out of a state is labeled with some value $b \in D$ for x_i. For each $Z \in D^N$, the branching program accepts Z if and only if there exists at least one (directed) path starting at the q_{start} and leading to the accepting state q_{accept}, and such that all labels along this path are consistent with Z. The size of a branching program is the number states (i.e. nodes) in the graph.

A branching program is *semantic read-k* if for every path from q_{start} to q_{accept} that is consistent with some input, each variable occurs at most k times along the path. For the read-once case, a semantic branching program allows variables to be read more than once, but each accepting path may only query each variable at most once.

2.2   Polynomial Evaluation Problem

Our hard computational problem is the polynomial evaluation problem, $Poly$, with parameters K, d, δ, where $0 < \delta < 1$. The input is a pair (u, x) where $u \in [K]^d$ specifies a degree $d - 1$ polynomial over the field $[K]$ (K a prime power), and $x \in [K]$ specifies a field value. $Poly(u, x) = 1$ if and only if the polynomial specified by u on input x evaluates to a number less than $K^{1-\delta}$. (We compare two field elements by comparing them using the natural ordering on ternary strings.)

We will work with $|D|$-ary branching programs (with $|D|$ prime), and let $K = |D|^n$. The input will be given as a vector in $D^{(d+1)n}$. The first dn coordinates specify u and the last n coordinates specify x. Thus the total input length is $(d + 1)n$. In the remainder of the

paper, $|D| = 3$, and thus the parameters of $Poly$ are d, δ, n. Both d and δ will be fixed constants. Let $Poly_u$ denote the polynomial evaluation problem with parameters d, δ, n where the polynomial u is fixed.

The actual lower bounds we show will be for a *sensitive* function f_u obtained from $Poly_u$ as follows. Let $a \in GF(q)$ where $q = |D|$ is a prime number. Let $h : D^n \to \{0, 1\}$ be the characteristic function of the hyperplane at a:

$$h_a(x) = 1 \quad \text{iff} \quad x_1 + x_2 + ... + x_n = a \mod q \,.$$

Fix an element $a(u)$ for which h_a accepts the largest number of vectors accepted by $Poly_u$ and define the function

$$f_u(x) = Poly_u(x) \wedge h_{a(u)}(x) \,.$$

We call f_u sensitive because it has the property that changing the value of exactly one variable in a yes input always gives an input vector that is a no instance. As a result any two accepted inputs differ in the value of at least two variables. Similarly for the polynomial evaluation problem $Poly$, where the coefficient vector u is part of the input, we define $f(u, x) = Poly(u, x) \wedge h_{a(u)}(x)$, which is sensitive in x.

2.3   Rectangles and Embedded Rectangles

We use the same definitions and conventions as in [3]. A product $U \times V$ is called a (combinatorial) rectangle. If $A \subseteq N$ is an index subset and $Y \subseteq D^A$ and $Z \subseteq D^{N-A}$, then the product set $Y \times Z$ is naturally identified with the subset $R = \{\sigma\rho \mid \sigma \in Y, \ \rho \in Z\}$ of D^N, and a set of this form is called a *rectangle* in D^N.

An *embedded rectangle* R in D^N is a triple $(\pi_{red}, \pi_{white}, C)$ where π_{red}, π_{white} are disjoint subsets of N, and $C \subseteq D^N$ satisfies: (i) The projection $C_{N-\pi_{red}-\pi_{white}}$ consists of a single partial input w, (ii) if $\tau_1 \in C_{\pi_{red}}$ and $\tau_2 \in C_{\pi_{white}}$, then the point $\tau_1\tau_2 w \in C$. C is called the *body* of R. The sets π_{red}, π_{white} are called the *feet* of the rectangle; the sets $C_{\pi_{red}}$ and $C_{\pi_{white}}$ are the *legs*, and w is the *spine*. We can also specify an embedded rectangle by its feet, legs and spine: $(\pi_{red}, \pi_{white}, A, B, w)$ where π_{red}, π_{white} are the feet, $A = C_{\pi_{red}}, B = C_{\pi_{white}}$ are the legs, and w is the spine.

We will sometimes refer to A as the *red* side of the rectangle and to B as the *white* side of the rectangle. The *size* of the rectangle is $|A| \cdot |B|$, and the dimension of the rectangle is m_r-by-m_w where $m_r = |\pi_{red}|$ and $m_w = |\pi_{white}|$.

3   Lower Bound for $|D| = 3$

▶ **Theorem 1.** *There exists constants d, δ such that for sufficiently large n, for a random u, with probability greater than $1/4$, any 3-ary nondeterministic semantic read-once branching program for f_u requires size at least $2^{\Omega(n)}$.*

▶ **Corollary 2.** *There exists constants d, δ such that for sufficiently large n, any 3-ary nondeterministic semantic read-once branching program for $f(u, x)$ with parameters d, δ, n requires size at least $2^{\Omega(n)}$.*

Overview of Proof.   Call a degree $d - 1$ polynomial "good" if the fraction of accepting instances is roughly what you would expect from a random function; that is, the fraction of yes instances is at least $\frac{1}{2}K^{-\delta}$. Lemma 7 shows that at least half of all degree $d - 1$ polynomials are good.

The main lemma (Lemma 3) shows that for all good polynomials $Poly_u$ to their corresponding sensitive function f_u, we can associate with every size $s = 2^{o(n)}$ branching program $\mathcal{P}$ computing f_u, an m_r-by-m_w embedded rectangle $R_{\mathcal{P}}$ of size r^2, where r will be a large constant, and m_r and m_w will be roughly equal, and will each be a constant fraction of n. For simplicity of calculations for now, assume that $m_r = m_w = m$. The rectangle will have the property that $\mathcal{P}$ accepts every input in $R_{\mathcal{P}}$; in other words, $R_{\mathcal{P}}$ is a 1-rectangle of $\mathcal{P}$. Choosing $d = r^2$, each rectangle of size r^2 can be a 1-rectangle for very few degree $d - 1$ polynomials – at most a $|D|^{-n\delta r^2}$ fraction of all degree $d - 1$ polynomials. (This is Lemma 6.) Secondly, the total number of such rectangles is fairly small – of size roughly $|D|^{O(rm)}$ (Lemma 5). The key point is that the *number* of rectangles is roughly $|D|^{2rm}$ – the exponent grows linearly in r. (More precisely it grows linearly in the sum of the lengths of the sides of the rectangle, $|A| + |B|$). But on the other hand, the probability that a degree $d = r^2$ polynomial takes on values less than $K^{1-\delta}$ within the rectangle is roughly $|D|^{-mr^2}$ – that is, the exponent grows *quadratically* with r. (More precisely it grows linearly in the size of the rectangle $|A|.|B|$). Because $|D|^{-n\delta r^2}|D|^{O(rn)}$ is less than $1/4$, this implies that many good degree $d - 1$ polynomials have no size r^2 1-rectangle, thus proving the theorem.

Note that we set our parameters so that the area of the rectangle $R_{\mathcal{P}}$ is at least the degree d of the polynomial u. (Thus $r^2 \geq d$.) A crucial point in the above argument is that the sum of the lengths of the sides of $R_{\mathcal{P}}$ must be at most a fraction of its area. This requires that the rectangle is reasonably close to being square. We put extra effort into making sure that the rectangle is square (without compromising too much of its size in order to make it square). This enables us to achieve domain size 3; a somewhat simpler argument achieves domain size 5.

▶ **Lemma 3** (Main Lemma). *Let $f : D^n \to \{0, 1\}$ be any sensitive boolean function such that the density of 1's is at least $\frac{1}{2|D|}K^{-\delta}$. Suppose that the following inequalities are satisfied for our parameters:*

(1) $m_w = 4m_r = \gamma n$;

(2) $|D|^{m_r} \leq |D|^{m_w}(1/2 - 2\gamma)^{m_w}$;

(3) $r \leq \frac{1}{4|D|s}(1/2 - \gamma)^{m_r}|D|^{m_r - \delta n}$.

Then if $\mathcal{P}$ is a $|D|$-way nondeterministic semantic read-once branching program of size s for f then there is an m_r-by-m_w embedded rectangle $R = (\pi_{red}, \pi_{white}, A, B, w)$ such that every input in R is accepted by $\mathcal{P}$, and where $|A|, |B| = r$.

Proof. Let f be a sensitive function such that the density of 1's is at least $\frac{1}{2|D|}K^{-\delta}$. Suppose there is a size s nondeterministic semantic read-once branching program, $\mathcal{P}$ for f. Let S_0 be the set of inputs that are accepted by $\mathcal{P}$; since $\mathcal{P}$ is assumed to be correct for all inputs of f, we have $|S_0| \geq \frac{1}{2|D|}K^{-\delta}|D|^n$. For each accepted instance $I \in S_0$, fix one accepting path, p_I, in the branching program. Since the function is sensitive each of the n variables must be read along any accepting path. For if some variable is not read along a computation path then changing the value of that variable alone would produce an accepting instance. However, this can't be the case for a sensitive function since any two accepted inputs will have to differ in at least two positions. So each of the n variables must be read along this path exactly once and thus each accepting instance I has an associated permutation π_I of the n variables associated with its accepting path p_I. Designate state q_I as the state in p_I which occurs just after the first half of the variables in π_I. Now define q to be the most common designated state (over all accepting inputs $I \in S_0$), and let $S_1 \subseteq S_0$ denote the corresponding set of inputs whose designated state is q. Thus for each input I in S_1, there is an accepting path

p_I that passes through state q. Because $\mathcal{P}$ has size s, it follows that

$$|S_1| \geq |S_0|/s \geq \frac{1}{2|D|s} K^{-\delta} |D|^n = \frac{1}{2|D|s} |D|^{n-\delta n} \tag{1}$$

We now want to pick two subsets of coordinates $\pi_{red} \subseteq N$ and $\pi_{white} \subseteq N$, of size m_r and m_w respectively, and a set $S^* \subseteq S_1$ of inputs with the property that for every input $I \in S^*$, and associated accepting path p_I, not only does it pass through state q, but every coordinate in π_{red} is read before state q, and every coordinate in π_{white} read at or after state q.

We will first pick π_{red} greedily. For each $I \in S_1$, at least $n/2$ of the n coordinates in p_I occur in π_I before reaching state q, and thus there is some coordinate i such that for at least half of the inputs $I \in S_1$, i occurs in π_I before reaching state q. After choosing the first coordinate, there are at least $|S_1|/2$ inputs remaining. Continue greedily until we pick m_r coordinates, π_{red}, always choosing the most popular coordinate that occurs in π_I before reaching state q. By averaging, when the i^{th} coordinate, $i \leq m_r < \gamma n$ is chosen, the fraction of inputs that remain is at least $\frac{(n/2-i)}{(n-i)} \geq \frac{(n/2-\gamma n)}{(n-\gamma n)} \geq \frac{(n/2-\gamma n)}{n} = (1/2 - \gamma)$. Let $S_2 \subseteq S_1$ denote the set of inputs such that all coordinates in π_{red} are read before reaching q. It follows that

$$|S_2| \geq (1/2 - \gamma)^{m_r} |S_1| \tag{2}$$

By assumption (3) in the statement of the Lemma, we have

$$r \leq \frac{1}{4|D|s} (1/2 - \gamma)^{m_r} |D|^{m_r - \delta n} \tag{3}$$

Then from (1), (2), and (3) we have

$$|S_2| \geq 2r |D|^{n-m_r} \tag{4}$$

For each $w \in D^{n-\pi_{red}}$, the average number of extensions of w in S_2 is $2r$. We want to prune S_2 such that every $w \in D^{n-\pi_{red}}$ has at least r extensions. To do this, define $S_3 \subseteq S_2$, where we remove all inputs $(w, *)$ from S_2 such that w has less than r extensions in S_2. Since we delete at most $r|D|^{n-m_r}$ elements from S_2, and $|S_2| \geq 2r|D|^{n-m_r}$, it follows that

$$|S_3| \geq r |D|^{n-m_r} \tag{5}$$

Next we will choose m_w coordinates, π_{white} in the same greedy fashion, and let S_4 denote the set of all inputs in S_3 such that all coordinates in π_{white} are read after reaching q. Again by averaging,

$$|S_4| \geq (1/2 - 2\gamma)^{m_w} |S_3| \tag{6}$$

We will express S_4 as the disjoint union of sets R_w: choose a value w for the coordinates outside of $\pi_{red} \cup \pi_{white}$. The corresponding set $R_w \subseteq S_4$ consists of all inputs (α, w, β) such that α is an assignment to the variables in π_{red}, β is an assignment to the variables in π_{white}, and $(\alpha, w, \beta) \in S_4$.

▶ **Lemma 4.** *For each w: (i) R_w is an embedded rectangle and (ii) as long as R_w is not empty, the size of its red leg is at least r.*

Proof. We will first show that R_w is an embedded rectangle. Let $S_{red} \subseteq D^{\pi_{red}}$ be the projection of R_w onto the coordinates of π_{red} and let $S_{white} \in D^{\pi_{white}}$ be the projection of R_w onto the coordinates of π_{white}. Setting $A = S_{red}$, $B = S_{white}$ and $w = w$, we claim

that R_w is equal to the embedded rectangle defined by $(\pi_{red}, \pi_{white}, A, B, w)$. To see this, consider $x, x' \in A$ and $y, y' \in B$ such that $xyw \in R_w$, and $x'y'w \in R_w$. Let I be the input corresponding to xyw and let p_I be the corresponding path going through state q. Note that in p_I the x-variables are all read prior to reaching q, and the y-variables are read after reaching q, and there is some split of the w variables into w_1, w_2 where the w_1 variables are read prior to q and the w_2 variables are read after q. Similarly, let I' be the input corresponding to $x'y'w$ and let $p_{I'}$ be the corresponding path. There is now a possibly different split of w into w_1', w_2', so x', w_1' are read before q and y', w_2' are read after q. We claim that $xy'w \in R_w$: consider the path (x, w_1) (the first half of p_I) and (y', w_2') (the second half of $p_{I'}$). This path must be consistent since w_1 and w_2' are consistent and x, y' are on disjoint variables. Thus there is an input consistent with this path; it is an accepting path going through q and consistent with w; the variables in π_{red} are all read before q, and the variables in π_{white} are all read after q. Thus it is in R_w. An analogous argument shows that $x'yw \in R_w$. Thus R_w is an embedded rectangle.

Secondly we will show (ii) for each $R_w \subseteq S_4$, the size of the red leg is at least r. (That is, $|A| \geq r$.) Consider a nonempty rectangle R_w with red leg A, white leg B and spine w. Recall that the inputs in S_3 consist of a partial input $w+ \in D^{N - \pi_{red}}$ together with a set $A \subseteq D^{\pi_{red}}$ such that $|A| \geq r$. We obtain S_4 from S_3 by selecting m_w coordinates from $N - \pi_{red}$, one at a time, choosing each coordinate greedily, where a coordinate is chosen if it is read *after* state q in the most inputs. Consider a block of inputs $(A, w+) \in S_3$. If some input $(\alpha, w+) \in (A, w+)$ survives, then all coordinates in π_{white} that were chosen must all be read after state q on input $(\alpha, w+)$. But this means that for every input $(\alpha', w+) \in (A, w+)$, all coordinates in π_{white} are also read after q. (Otherwise, some coordinate would be read twice along this accepting input, violating the read-once condition.) Thus, either the entire block $(A, w+)$ is in S_4, or the entire block is removed from S_4.

Now let $R_w = (\pi_{red}, \pi_{white}, A, B, w) \subseteq S_4$ be a nonempty rectangle, $w \in D^{N - \pi_{red} - \pi_{white}}$. R_w is obtained by taking the union of (nonempty) blocks $(A', w+) \in S_4$, $w+ \in D^{N - \pi_{white}}$. Since as we argued above, for each such block, $|A'| \geq r$, it follows that $|A| \geq r$ as well. ◄

Let r_{avg} denote the average size of the white leg of the rectangle over all rectangles $R_w \subseteq S_4$. It is easy to see that $|D|^{n - m_w} r_{avg} \geq r|D|^{n - m_r}(1/2 - 2\gamma)^{m_w}$. It follows that $r_{avg} \geq r$ if $|D|^{m_w - m_r}(1/2 - 2\gamma)^{m_w} \geq 1$. The latter inequality follows from condition (2). Thus, by condition (2) assumed in the hypothesis of Lemma 3, we can pick some setting w^* to the remaining $n - m_r - m_w$ uncoloured coordinates (the coordinates that are not in π_{red} or π_{white}) such that the white leg of the rectangle R_{w^*} has size at least $r_{avg} = r$. Let S^* equal R_{w^*}. By construction, both the red leg of $S^* = R_{w^*}$ and the white leg of R_{w^*} have size at least r. Prune S^* so that each leg has size exactly r, thus completing the proof of the lemma. ◄

▶ **Lemma 5.** *Let $\mathcal{R}$ be the set of all m_r-by-m_w embedded rectangles over D^N such that $|A| = |B| = r$, where $m_w = \gamma n$ and $m_r = m_w/4$. Then $|\mathcal{R}| \leq (e/\gamma)^{\frac{5}{4} m_w} |D|^{\frac{5}{4} r m_w + m_w/\gamma}$.*

Proof. The number of choices for π_{red}, the coordinates of A, is $\binom{n}{m_r}$. Given π_{red}, we choose r vectors from the $|D|^{m_r}$ possible values for the elements of A. Thus the total number of possible sets A is at most $\binom{n}{m_r}|D|^{r m_r}$. Similarly the number of choices for the set B is at most $\binom{n}{m_w}|D|^{r m_w}$. The number of choices for $w \in D^{N - \pi_{red} - \pi_{white}}$ is $|D|^{n - m_r - m_w}$. Thus $|\mathcal{R}|$ is at most $\binom{n}{m_r}\binom{n}{m_w}|D|^{r m_r}|D|^{r m_w}|D|^{n - \frac{5}{4} m_w}$. Using the inequality $\binom{n}{k} \leq (\frac{en}{k})^k$ we conclude the number of choices for $|\mathcal{R}|$ is at most $(en/m_w)^{m_w}(4en/m_w)^{\frac{1}{4} m_w}|D|^{n - \frac{5}{4} m_w}|D|^{\frac{5}{4} r m_w} \leq (e/\gamma)^{\frac{5}{4} m_w}|D|^{\frac{5}{4} r m_w + m_w/\gamma}$ ◄

▶ **Lemma 6.** *Define the predicate $Good(R, u)$ to be true if for every input x in the rectangle R, the polynomial u on input x is less than $K^{1-\delta}$ (i.e. $Poly_u(x)$ is true). Then for all embedded rectangles R of size d, $Pr_u[Good(R, u)] \leq p$ where $p = |D|^{-\delta nd}$.*

Proof. Assume $Good(R, u)$. Suppose that $|R| = d$ and let $B' \in [K^{1-\delta}]^d$ specify a vector of d accepting values. Let $Good_{B'}(R, u)$ to be the event that for all $x \in R$, $Poly_u(x) = B'(x)$. Then $Pr_u[Good(R, u)] = K^{(1-\delta)d} \cdot Pr_u[Good_{B'}(R, u)]$.

To bound $Pr_u[Good_{B'}(R, u)]$, suppose that it is true that $\forall x \in R$, $F_u(x) = \sum_{i<d} u_i x^i = B'(x)$. Note that this fixes the output of the degree $d-1$ polynomial for d values of x. By interpolation, this uniquely determines the polynomial, u'. Thus, $Pr_u[Good_{B'}(R, u)] = Pr_u[u = u'] = K^{-d} = |D|^{-nd}$. Overall, $Pr_u[Good(R, u)] \leq K^{(1-\delta)d}|D|^{-nd} = |D|^{n(1-\delta)d}|D|^{-nd} = |D|^{-\delta nd}$. This completes the proof of Lemma 6.   ◀

▶ **Lemma 7.** *For a random u, for fixed parameters d, δ the probability that $Poly_u(x)$ does not accept a $(1 \pm o(1))K^{-\delta}$ fraction of all the inputs is at most $o(1)$. (Here both $o(1)$ are $K^{-(1-\delta)/3}$.)*

Proof. Randomly choose the coefficients $u \in [K]^d$ of the $d-1$ degree polynomial. For each instance $x \in [K]$ (and value $b \in [K]$), let $A_{\langle x,b \rangle}$ denote the event that the output of this polynomial on input x is b. Let a_x denote the event that this value is less than $K^{1-\delta}$ so that x is a yes instance. Let $Y = \sum_{x \in K} a_x$ denote the number of yes instances for the chosen u. Note $p = Pr_u[a_x] = K^{1-\delta}/K = K^{-\delta}$ because just choosing the constant coefficient u_0 of the polynomial randomly makes the polynomial's output on x uniformly random in $[K]$. Hence, by linearity of expectation $\overline{Y} = Exp[Y] = K \cdot Pr_u[a_x] = K^{1-\delta}$. We show that the $A_{\langle x,b \rangle}$ events for different x are d-wise independent as follows. Consider any subset $\{x_1, x_2, \ldots, x_d\} \subset [K]$ of the instances. Knowing the value of the polynomial at each of these instances, by interpolation, uniquely determines the coefficients u of the polynomial. Hence, if all you know about u is the values on $d-1$ of these instances, then the value on the remaining is still uniformly random within $[K]$. Formally stated, $Pr_u[A_{\langle x_d,b_d \rangle} \mid A_{\langle x_1,b_1 \rangle}, \ldots, A_{\langle x_{d-1},b_{d-1} \rangle}] = Pr_u[A_{\langle x_d,b_d \rangle}]$. Not fully knowing the value of the first $d-1$ of the instances, but only that their value is small, give you even less information. Hence, $Pr_u[a_{x_d} \mid a_{x_1}, \ldots, a_{x_{d-1}}] = Pr_u[a_{x_d}]$. It follows that $Pr_u[a_{x_1} \wedge \ldots \wedge a_{x_d}] = Pr_u[a_{x_1}] \cdot \ldots \cdot Pr_u[a_{x_d}]$. Because the a_x events are d-wise independent, it follows that the d^{th} order standard deviation of their sum Y is the same as it would be if they were completely independent events. We, however, only need to consider the variance. More formally, for each x, let a'_x be an independent event with probability $K^{-\delta}$ of success and $Y' = \sum_{x \in K} a'_x$. The variance is $Var[Y] = Exp_u[(Y - \overline{Y})^2] = Exp_u[(\sum_x a_x - \overline{Y})^2]$. The non-linear part of this is $Exp_u[(\sum_x a_x)^2] = \sum_{x,x'} Exp_u[a_x \cdot a_{x'}]$, which we know from pair-wise independence is $\sum_{x,x'} Exp_u[a_x] \cdot Exp_u[a_{x'}] = \sum_{x,x'} Exp_u[a'_x] \cdot Exp_u[a'_{x'}]$. The same computation for the a'_x, gives that $\sigma^2 = Var[Y] = Var[Y'] = K \cdot p(1-p) \approx K \dot{K}^{-\delta} = K^{1-\delta} = \overline{Y}$. By Chebycheff's inequality, $\forall \eta > 0$ we have $Pr_u(|Y - \overline{Y}| \geq \eta\sigma) < \frac{1}{\eta^2}$. Setting $\eta = \overline{Y}^{\frac{1}{6}}$, gives $Pr_u(Y \notin (1 \pm \overline{Y}^{-\frac{1}{3}})\overline{Y}) \leq \overline{Y}^{-\frac{1}{3}}$.   ◀

We are now ready to complete the proof of the theorem. Call a polynomial u "good" if $Poly_u$ accepts at least a $\frac{1}{2}K^{-\delta}$ fraction of all inputs. By Lemma 7, we know that at least half of all u's are good. For each good u, the corresponding sensitive function f_u has density at least $\frac{1}{2|D|}K^{-\delta}$. Since f_u is sensitive and has sufficient density Lemma 3 tells us that any small branching program for f_u implies that there exists an m_r-by-m_w embedded rectangle size r^2 that is accepted (assuming that conditions (1), (2), and (3) are satisfied).

On the other hand, by union bound Lemmas 5 and 6 together tell us that at most a $p|\mathcal{R}|$ fraction of degree $d - 1$ polynomials u have such m_r-by-m_w embedded rectangles of size r^2

that are accepted. Suppose we can choose a setting of the parameters so that $p|\mathcal{R}| < 1/4$. If follows that for at least $\frac{1}{4}^{th}$ of all good polynomials the corresponding sensitive functions f_u do not have such m_r-by-m_w embedded rectangles of size r^2 that are accepted since the hyperplane constraint $h_{a(u)}(x)$ can only shrink an accepting rectangle. Then by Lemma 3 this implies that at least as many f_u cannot have small branching programs, and thus the theorem is proven.

It is left to show that we can set the parameters such that $p|\mathcal{R}| < 1/4$, and properties (1), (2), and (3) of Lemma 3 are satisfied. We will set the parameters as follows: $|D| = 3$, $m_w = 4m_r = \gamma n$, $\gamma = .01$, $\delta = \gamma/300$, $r = 3000$, and $d = r^2$. To achieve $p|\mathcal{R}| < 1/4$, we require $|D|^{\delta m_w r^2/\gamma - m_w/\gamma - \frac{5}{4}rm_w} > 4(e/\gamma)^{\frac{5}{4}m_w}$. Using $|D| = 3$ and factoring out m_w, it is sufficient if we have $3^{\delta r^2/\gamma - 1/\gamma - \frac{5}{4}r} > 4(e/\gamma)^{\frac{5}{4}}$. With our choice of parameters, this is satisfied for $r \geq 3000$.

For Lemma 3, we also require assumptions (2) and (3). First for (2): $|D|^{m_r} \leq |D|^{m_w}(1/2 - 2\gamma)^{m_w}$. For $|D| = 3$ and $m_w = 4m_r$, this is satisfied. For (3) we require: $r \leq \frac{1}{4|D|}(1/2 - \gamma)^{m_r}|D|^{m_r - \delta n} = \frac{1}{4|D|}(1/2 - \gamma)^{m_r}|D|^{m_r(1-4\delta/\gamma)}$. For $|D| = 3$, $\gamma = .01$, $\delta = \gamma/300$, we have $(1/2 - \gamma)|D|^{(1-4\delta/\gamma)} \geq 1.44$ and thus it suffices to show $r \leq \frac{1}{12}(1.44)^{m_r}/s$. This holds for our choice $r = 3000$ when $s \leq 2^{cm_r} = 2^{cn/(4\gamma)}$ for some $c > 0$ and sufficiently large n. Note that $|D| > 2$ helps us in ensuring assumptions (2) and (3) hold.

4   Conclusion

We have proved an exponential lower bound on the size of non-deterministic semantic read once branching programs computing a polynomial time computable function $f : D^n \to \{0,1\}$ when $D = \{0,1,2\}$ with just three elements. Our contribution is that we bring down the size of the domain required to achieve this. Prior to our result the best that was known was for D-ary branching programs with $|D| \geq 2^{13}$. The explicit function f for which we show the lower bound is the decision problem of determining whether a certain degree d polynomial over a finite field K evaluates to a value less than a certain threshold at a given input (along with a hyperplane constraint). This result brings down the focus to the first non-boolean case, $|D| = 3$ vs the boolean case, $|D| = 2$, since, interestingly the case where D is boolean $\{0,1\}$ still remains open and no non-trivial lower bounds are known for binary non-deterministic semantic read once branching programs [12]. In the next section we explore the Booelan case.

5   Semantic Branching Programs with $|D| = 2$ can evade Large Bottleneck Rectangles

In this section we show how binary non-deterministic semantic read once branching programs can behave differently by evading lower bounds in certain bottleneck layers by having a small number of states in those layers irrespective of what function $f : \{0,1\}^n \to \{0,1\}$ they are computing. The example upper bounds we give demonstrate why it is likely harder to prove lower bounds for semantic read once branching programs with domain size $|D| = 2$.

When the domain size is $|D| > 2$, the technique is to prove that the set of yes instances handled by any one state of the branching program contains a rectangle and then identify a computational problem that has no large rectangle of yes instances. Hence, the rectangle for each state must be small. Because there are exponentially many yes instances and each must be handled by at least one state at a selected *bottle neck* level of the branching program, there must be an exponential number of states at that level. We show here that for domain

size $|D| = 2$, the set of yes instances handled by one state can be quite arbitrary and quite large. This does not mean that the total number of branching program states can necessarily be small. But it does mean that at the one level of the branching program that the prover is hoping to use for a bottleneck, the number of states might be quite small.

A lower bound that attempts to prove that a selected bottleneck level of the branching program must have many states, must start by selecting which level of the branching program will be the level in question. It might do this by specifying how many or which variables have been read so far. Given *any boolean computational problem* with input from $\{0,1\}^n$ and a criteria for choosing the bottle neck level chosen from a wide (but not exhaustive) range of possible choices, we now show how to fool such a lower bound, by giving a branching program that gives a polynomial upper bound on the number states at the selected layer.

The branching program is constructed as follows. For each yes instance $A \in \{0,1\}^n$, we form an accepting branching program path $\langle C_1(A), q(A), C_2(A) \rangle$ where $q(A)$ denotes the state A passes through at the bottleneck level, $C_1(A)$ the path before this level and $C_2(A)$ that after. Note that to get a counter example to a lower bound technique using some bottleneck layer, we don't need to give a full poly-size branching program. We only need the number of states $q(A)$ to be small. It can have exponential number of states before and after this layer. Hence, we will have all of these paths $C_1(A)$ for different yes instances A be completely disjoint from each other. Similarly for $C_2(A)$. These paths only come together and interact at the special layer of states $q(A)$. In order to make the properties of this level more arbitrary, let $A_1, A_2 \subset [n]$ be any partition of the input variables into two parts. Let $C_1(A)$ read all the ones in A_1 and all the zeros in A_2. Let $C_2(A)$ read all the zeros in A_1 and all the ones in A_2. Let $q(A) = \langle u, v \rangle$ be the state, where $u \in [n]$ is the number of ones in A_1 and $v \in [n]$ is the number of the zeros in A_2. Hence only n^2 states are needed in the layer. Note that because we have allowed the computational problem to be arbitrary, other than partitioning its yes instances based on their hamming weights, the sets of instances handled by a state $q(A)$ is completely arbitrary.

Note that as long as A_1, A_2 are comparable in size, for most of the inputs, the incoming path $C_1(A)$ and $C_2(A)$ are of comparable length. However, the purported bottleneck layer for which we give the above upperbound is not identical to the one we use for our lower bound for $|D| \geq 3$ in the sense that the bottleneck states like $q(A)$ do not appear exactly midway through the accepting paths at length $n/2$ on all the paths as is required in Lemma 3. Nevertheless, the upper bound is interesting because for most inputs A the incoming and outgoing paths through a state in the layer are of comparable length.

We will now in two ways, prove that this branching program solves the given computational problem. We will start with a *communication game interpretation*. Think about the algorithm as a game between two players C_1 and C_2 and Charlie who they don't trust. Charlie shows C_1 the ones in the first part A_1 of the input and the zeros in the second part A_2. Assuming he trusts Charlie, this lets C_1 know the entire input. Hence, he can answer *any* question about the input. The only way that Charlie can cheat is to not show all of the entries. In order to verify that he is not lying, C_1 sends to C_2 the number of ones in A_1 and the number of zeros in A_2. C_2 can then check that they have both been shown all of what they were supposed to see.

Now lets consider the *branching program interpretation*. Clearly, the branching program described accepts all yes instances of the given problem, because it has a separate accepting path $\langle C_1(A), q(A), C_2(A) \rangle$ for each yes instance A. What remains is to prove 1) the branching program is semantic read once and 2) that no no instances are accepted. We do this by showing for every pair $\langle A, B \rangle$ of different yes instances that pass through the same bottleneck

states $q(A) = q(B) = \langle u, v \rangle = q$, that the cross path $\langle C_1(A), q, C_2(B) \rangle$ is inconsistent in that it reads some variable twice with different values. Hence, by the definition of *semantic*, it does not matter that this path is not read once and because it is inconsistent, it cannot be accepting a no instance. Because A and B are different, either A_1 and B_1 are different and/or A_2 and B_2 are different. Assume A_1 and B_1 are different. Because A_1 and B_1 have the same number u of ones, there is an element that is one in A_1 and zero in B_1. Hence $C_1(A)$ and $C_2(B)$ both read it. This element is read twice in $\langle C_1(A), q, C_2(B) \rangle$ with different values and so is inconsistent.

So for $D = 2$, presence of a small number of states in a supposed bottleneck layer of a branching program need not imply that there exists a balanced embedded rectangle of accepting instances. In particular, our lower bound finds a rectangle within the set of yes instances handled by narrowing the set down to a subset within which for many variables it is fixed whether it is read before or after the state. However in this upper bound, whether a variable is read before or after the state $q(A)$ is completely determined by whether its value is 0 or 1. Hence, fixing this fixes its value. If this is done for every variable, the set of inputs left in the eventual rectangle identified by this lower bound method is narrowed down to a singleton.

───── **References** ─────

1 M. Ajtai. A non-linear time lower bound for boolean branching programs. In *Proceedings 40th FOCS*, pages 60–70, 1999.

2 P. Beame, T.S. Jayram, and M. Saks. Time-space tradeoffs for branching programs. *J. Comput. Syst. Sci*, 63(4):542–572, 2001.

3 P. Beame, M. Saks, X. Sun, and E. Vee. Time-space trade-off lower bounds for randomized computation of decision problems. *Journal of the ACM*, 50(2):154–195, 2003.

4 Allan Borodin, A Razborov, and Roman Smolensky. On lower bounds for read-k-times branching programs. *Computational Complexity*, 3(1):1–18, 1993.

5 Alan Cobham. The recognition problem for the set of perfect squares. In *Switching and Automata Theory, 1966., IEEE Conference Record of Seventh Annual Symposium on*, pages 78–87. IEEE, 1966.

6 Scott Diehl and Dieter Van Melkebeek. Time-space lower bounds for the polynomial-time hierarchy on randomized machines. *SIAM Journal on Computing*, 36(3):563–594, 2006.

7 L. Fortnow. Nondeterministic polynomial time versus nondeterministic logarithmic space: Time space tradeoffs for satifiability. In *Proceedings 12th Conference on Computational Complexity*, pages 52–60, 1997.

8 L. Fortnow and D. Van Melkebeek. Time-space tradeoffs for nondeterministic computation. In *Proceedings 15th Conference on Computational Complexity*, pages 2–13, 2000.

9 Lance Fortnow, Richard Lipton, Dieter Van Melkebeek, and Anastasios Viglas. Time-space lower bounds for satisfiability. *Journal of the ACM (JACM)*, 52(6):835–865, 2005.

10 S. Jukna. A nondeterministic space-time tradeoff for linear codes. *Information Processing Letters*, 109(5):286–289, 2009.

11 Stasys Jukna. A note on read-k times branching programs. *Informatique théorique et applications*, 29(1):75–83, 1995.

12 Stasys Jukna. *Boolean function complexity: advances and frontiers*, volume 27. Springer Science & Business Media, 2012.

13 Stasys P Jukna. The effect of null-chains on the complexity of contact schemes. In *Fundamentals of Computation Theory*, pages 246–256. Springer, 1989.

14 Matthias Krause, Christoph Meinel, and Stephan Waack. Separating the eraser turing machine classes le, nle, co-nle and pe. In *Mathematical Foundations of Computer Science 1988*, pages 405–413. Springer, 1988.

15 R. Lipton and A. Viglas. Time-space tradeoffs for sat. In *Proceedings 40th FOCS*, pages 459–464, 1999.

16 EA Okolnishnikova. On lower bounds for branching programs. *Siberian Advances in Mathematics*, 3(1):152–166, 1993.

17 Pavel Pudlak and Stanislav Zak. Space complexity of computations. *Preprint Univ. of Prague*, 1983.

18 Ryan Williams. Better time-space lower bounds for sat and related problems. In *Computational Complexity, 2005. Proceedings. Twentieth Annual IEEE Conference on*, pages 40–49. IEEE, 2005.