



Black-White Pebbling is PSPACE-Complete

Philipp Hertel
Department of Computer Science
University of Toronto
Toronto, ON CANADA
philipp@cs.toronto.edu

Toniann Pitassi
Department of Computer Science
University of Toronto
Toronto, ON CANADA
toni@cs.toronto.edu

April 25, 2007

Abstract

The complexity of the Black-White Pebbling Game has remained an open problem for 30 years. It was devised to capture the power of non-deterministic space bounded computation. Since then it has been continuously studied and applied to problems in diverse areas of computer science including VLSI design and more recently propositional proof complexity. In 1983, determining its complexity was rated as “An Open Problem of the Month” in David Johnson’s *NP-Completeness Column*. In this paper we show that the Black-White Pebbling Game is PSPACE-complete.

1 Introduction

The Black-White Pebbling Game was introduced by Cook and Sethi in 1976 [3] in the context of determining lower bounds for space bounded Turing Machines. The problem received considerable attention throughout the next decade due to its numerous applications including VLSI design, compilers, and algebraic complexity. In 1983 determining its complexity was rated as “An Open Problem of the Month” in David Johnson’s *NP-Completeness Column* [9]. An excellent survey of pebbling results from this period can be found in Pippenger [15]. Recently, there has been a resurgence of interest in pebbling games due to their links with propositional proof complexity [1, 2, 4, 5, 13]. In this paper we prove that the Black-White Pebbling Game is PSPACE-complete.

The Black-White Pebbling Game was preceded by the Black Pebbling Game, which has also been widely studied [15]. Let $G = (V, E)$ be a directed acyclic graph with one distinguished output node, s . In the Black Pebbling Game, a player tries to place a pebble on s while minimizing the number of pebbles placed simultaneously on G . The game is split up into distinct steps, each of which takes the player from one pebbling configuration to the next. Initially, the graph contains no pebbles and each subsequent configuration follows from the previous by one of the following rules:

- At any point a black pebble can be placed on any source node v .
- At any point a black pebble can be removed from any node v .
- For any node v , if all of v ’s predecessors have pebbles on them, then a black pebble can be placed on v , or a black pebble can be slid from a predecessor u to v .

The Black Pebbling Game models deterministic space-bounded computation. Each node models a result and the placement of a black pebble on a node represents the deterministic computation of the result from previously computed results. A sequence of moves made by the player is called a *pebbling strategy*. If a strategy manages to pebble s using no more than k pebbles, then that strategy is called a k -pebbling strategy.

The Black-White Pebbling Game is a more powerful extension of the Black Pebbling Game in which white pebbles, which behave in a dual manner to the original black pebbles, can also be used. As before, the player attempts to place a black pebble on s while minimizing the number of pebbles placed simultaneously on G at any time. The Black-White Pebbling Game extends the Black Pebbling Game with the addition of the following rules:

- At any point a white pebble can be placed on any node v .
- At any point a white pebble can be removed from any source node v .
- For any node v with a white pebble on it, the pebble can be slid to an empty predecessor u if all of v ’s other predecessors are pebbled, or the white pebble can be removed if all of v ’s predecessors are pebbled.
- The game ends when s contains a black pebble and every other node is empty.

As before, the placement of each black pebble is meant to model the derivation of a deterministically-computed result, while the placement of each white pebble is meant to model a non-deterministic guess, whose verification requires all of its antecedents to be derived. Since the game ends when there is only a single black pebble on the target, the game cannot complete until all of these guesses have been verified and thereby discharged. Clearly every black pebbling strategy is a black-white pebbling strategy.

In 1978, Lingas showed that a generalization of the Black Pebbling Game, played on monotone circuits instead of DAGs, is PSPACE-complete [12]. This was a somewhat surprising result since the PSPACE-complete games of the time involved two players and it was clear how the alternation between them led to

each game’s high complexity. Lingas’s Generalized Black Pebbling Game, on the other hand, is a single player game with no obvious alternation. Its complexity stems from the necessity to repeggle some nodes many times in order to achieve the minimum pebbling number for some graphs. Lingas’s ingenious reduction exploited exactly this phenomenon to force any optimal strategy on his circuits to necessarily verify the truth of a quantified boolean formula (QBF).

In 1980, Gilbert, Lengauer, and Tarjan elaborated on the basic structure of Lingas’s construction to prove that the Black Pebbling Game on DAGs is PSPACE-complete [6]. The main difficulty in moving from monotone circuits to the more restricted class of DAGs is the creation of an OR widget using only the global bound on the number of permissible pebbles and nodes which act like AND gates. Though their exposition does not focus on it, this is a significant technical hurdle when extending Lingas’s ideas to DAGs.

Both reductions were devised to force any optimal black pebbling strategies to verify a QBF. By their nature, black pebbling strategies are very inductive and can only pebble graphs in one direction. As a result, large portions of a graph remain unpebbled while progress is being made linearly from the source nodes toward them. In contrast, white pebbles allow a much richer choice of strategies since they can be placed anywhere on the graph regardless of where pebbles were placed before, thereby breaking up the straight inductive pattern obvious in all pure black strategies. Although the black pebbling number of a graph is never more than a square of the black-white pebbling number [7], the addition of white pebbles lowers the pebbling number of many graphs [11], [16], [10]. Unfortunately, the constructions used for the previous PSPACE-completeness results are both examples of such graphs. As a result, neither can be used to differentiate between true and false QBFs in the presence of white pebbles.

In this paper, we finally resolve Johnson’s open problem by building on the construction of [6] to prove the PSPACE-completeness of the Black-White Pebbling Game. Since white pebbles can be used so unpredictably, we create graphs on which the use of even a single white pebble on anything other than a source node (where black and white pebbles are almost indistinguishable) leads to a sub-optimal pebbling. When applied to the right family of QBFs, our reduction also provides an infinite family of graphs which require exponential time to minimally black-white pebble, but can be pebbled in linear time if we use just one pebble more than the minimum. This results in a time/space tradeoff result similar to that proved in [6] for pure black pebbling.

2 Definitions and Proof Overview

Formally, the Black-White Pebbling Game takes as input a DAG \mathcal{G} with a special target node s and an integer k and asks whether there is a k -pebbling strategy for s in \mathcal{G} . We prove the following theorem.

Theorem 1: The Black-White Pebbling Game is PSPACE-complete.

It is not hard to see that black-white pebbling is in PSPACE. Given (\mathcal{G}, k) , we can easily guess a sequence of configurations that pebbles \mathcal{G} with at most k pebbles. Then by Savitch’s theorem, this implies that black-white pebbling is in (deterministic) PSPACE.

The next two sections will be devoted to showing that the Black-White Pebbling Game is PSPACE-hard. To prove this, we will reduce from QSAT. Given a QBF ψ , we will create a graph \mathcal{G} with the property that ψ is QSAT if and only if \mathcal{G} has a $4n + 3$ black-white pebbling strategy.

Following the conventions of [14] and [6], we classify pebble placements as *necessary* or *unnecessary*. The first placement of a black pebble on the target vertex is necessary. A placement of a black or white pebble on any other node v is necessary if and only if the pebble remains on v until a necessary placement occurs on a successor of v (this can occur concurrently if we are sliding a black pebble up from v to the successor). We call a pebbling strategy which contains no unnecessary placements *frugal*. Clearly, removing all unnecessary placements from a k -pebbling strategy for a graph \mathcal{G} results in a frugal k -pebbling strategy for \mathcal{G} . We can therefore limit ourselves to considering just frugal pebbings. The notion of frugality is central to proving one of our most important lemmas, Lemma 7.

Our construction is similar at a high-level to [6], where they create a graph from a QBF with the property that the formula is QSAT if and only if the graph has a small pure black pebbling strategy. The general idea behind their reduction is to have the black pebbling correspond to the exponential-time procedure that verifies that ψ is QSAT.

The construction of [6] is broken up into two main subgraphs: a linear chain of clause widgets followed by a linear chain of quantifier widgets. In all strategies which achieve the construction's minimum pebbling number, pebbles must be placed on certain special nodes in a way which corresponds to the lexicographically first truth assignment in the *QSAT* model for ψ . Since this assignment satisfies ψ 's 3CNF the player is able to successfully pebble through the clause widgets without exceeding the minimum pebbling number. The player can then begin to make progress up to the first universal quantifier widget, say widget i . In order to pebble through this widget without exceeding the pebbling number, the player must leave a pebble on a "progress node" in widget i and then repebble the special nodes for the innermost i variables, thereby placing pebbles in a way which corresponds to the lexicographically second truth assignment in the *QSAT* model. The player can then pebble up through the clause widgets again, and this time use the pebble which was previously placed on the progress node to pebble through widget i , only to have his/her progress arrested at the next universal widget, at which point the process must repeat. Minimally black pebbling the graph corresponding to a true QBF with k universal quantifier widgets therefore requires 2^k time.

Unfortunately, the graphs used in all earlier constructions are easy to pebble once white pebbles are allowed, regardless of whether or not the formula is QSAT. Thus the main obstacle in proving hardness of black-white pebbling is to determine how to modify the construction so that white pebbles will be rendered useless. We exploit an important observation to do this. In 1979, Meyer auf der Heide [7] proved a strong duality between black and white pebbles. Namely, he proved that on any graph G , for any pure black k -pebbling strategy there is a pure white k -pebbling strategy and vice versa. In order to prove this, he made a modification to the rules of the game. Pure black strategies still begin with an empty graph and end with a single black pebble on the target node, but pure white pebbling strategies now begin with a single pebble on the target node, and end with a completely empty graph. His proof amounts to showing that running a pure black k -pebbling strategy backward yields a pure white k -pebbling strategy, and vice versa. This has some implications for the original Black-White Pebbling Game, in which every strategy must end with a single black pebble on the target node. Namely, if you try to use as close to a pure white strategy as you can to black pebble the target node of some DAG G and if the maximum pebbling number k is reached in any pure black strategy of G at some time when there is no black pebble on the target node, then the black-white strategy will necessarily need to use $k + 1$ pebbles, one black pebble on the target node and k white pebbles which are simulating some optimal black pebbling in reverse. By similar reasoning, if one can build a graph which requires the player to use the maximum number of pebbles in every configuration of every optimal pure black strategy, then using a white pebble in support of a black pebbling of any intermediate node should also exceed the maximum. Our construction is designed to enforce this while maintaining the original properties found in the construction of [6].

However, we run into troubles in the case of the existentially quantified variables. The problem stems from the fact that for an existential quantifier widget, we want to be able to pebble up to that widget in either of two different ways—one corresponding to the variable being set to true, and the other way corresponding to the variable being set to false. Thus, there is an implicit OR in this argument. This difficulty was also overcome in [6], in the more limited context of black pebbling. If we were constructing monotone circuits rather than graphs (which are special cases of monotone circuits with only AND gates), then things become much easier, even when allowing the use of white pebbles, since we can use an explicit OR gate to allow for either of these two types of pebbings. This was accomplished in [8] which uses OR gates as a building block in order to prove an exponential time/space speedup theorem for Resolution. However, when OR gates are not allowed, we have to somehow simulate this implicit OR using only AND gates. Any way of doing this will necessarily involve two different pebbings, and it is quite subtle to see how to accomplish this while still prohibiting white pebbles from being used. We manage to accomplish this with another new

idea that lets us simulate this implicit OR using only AND gates.

3 The Reduction

To show that the Black-White Pebbling Game is PSPACE-hard, we reduce from QSAT. In our presentation, a QBF $\psi = Q_n x_n Q_{n-1} x_{n-1} \cdots Q_1 x_1 F$, where F is a 3CNF containing m clauses over the n quantified variables x_n, \dots, x_1 . We have inverted the numbering of the variables simply as a convenience in the proof. Given a QBF ψ , we produce a graph G whose target node s can be black-white pebbled using at most $4n + 3$ pebbles if and only if ψ is QSAT. Our construction is designed to penalize any use of white pebbles, so that the optimal strategy is all black.

The graph which we construct is composed of $n + m$ widgets, one for each quantified variable and one for each clause in F . As in [6], the quantifier widget for $Q_i x_i$ contains four vertices which represent the variable x_i , we call these nodes $x_i, x'_i, \bar{x}_i, \bar{x}'_i$. The location of pebbles on these four nodes corresponds to the truth value assigned to x_i by the current truth assignment which is being tested by the pebbling. If pebbles are on x_i and \bar{x}'_i , then the variable x_i is set to true. If pebbles are on x'_i and \bar{x}_i or if pebbles are on x'_i and \bar{x}'_i , then the variable x_i is set to false. Our construction will never allow an assignment to place pebbles on both x_i and \bar{x}_i .

The construction of the quantifier widgets relies on a subwidget we call an i -slide. An i -slide is designed to severely restrict the player's pebbling strategies. An example of a 4-slide is shown in Figure 1. Once the bottom nodes of an i -slide are all black-pebbled, an i -slide strategy, where the bottom pebbles are slid up to the top nodes in the appropriate order, is the only way to black-pebble the top nodes without using more than i pebbles.

DEFINITION 3.1: An **i -slide** is a pair of sets (V, U) together with a set of edges that satisfy the following properties. V is a set of i nodes v^1, v^2, \dots, v^i and U is a set of i nodes u^1, u^2, \dots, u^i . The edges are as follows. (1) v^j is the predecessor of all nodes v^k such that $k > j$; (2) u^j is the predecessor of all nodes u^k such that $k > j$; (3) u^j is the predecessor of all nodes v^k such that $k \leq j$; (4) u^j has at least $i - j + 1$ predecessors from outside of V or U .

Globally the construction is very much like the construction in [6]. There are a number of nodes used to encode a truth assignment, which are predecessors to nodes in both clause widgets and quantifier widgets. The clause widgets are connected linearly and can only be pebbled within the space bound of $4n + 3$ if the truth assignment encoded by the current pebbling configuration satisfies F . The quantifier widgets are also connected to each other linearly and follow the last clause widget. They slow the advance of the pebbling toward s . In order to advance through them, it will be necessary to repebble the clause widgets numerous times, once for each truth assignment required to show that ψ is QSAT. Only once the final quantifier widget is pebbled is it possible to pebble the target node s . We now describe the individual widgets and how they are connected. These descriptions are somewhat terse and are meant to be read in accompaniment to Figures 1, 3, 4, 2, and 5.

The universal widget is depicted in Figure 3. For every $i, 1 \leq i \leq n$, if widget i is a universal widget, it is composed of 4 groups of nodes, $\{\bar{x}_i, \bar{x}'_i, d_i, x_i, x'_i, y_i\}$, $G_{i-1} = \{g_{i-1}^1, \dots, g_{i-1}^{4i-1}\}$, $\{a_i, b_i\}$, and $G_i = \{g_i^1, \dots, g_i^{4i+3}\}$. These are connected as follows. y_i has $4i + 3$ source nodes $p_{x_i}^1$ through $p_{x_i}^{4i+3}$ as predecessors, x'_i has $4i + 2$ source nodes $p_{x_i}^1$ through $p_{x_i}^{4i+2}$ as predecessors, d_i has $4i + 1$ source nodes $p_{d_i}^1$ through $p_{d_i}^{4i+1}$ as predecessors, and \bar{x}'_i has $4i$ source nodes $p_{\bar{x}_i}^1$ through $p_{\bar{x}_i}^{4i}$ as predecessors. The sole predecessor of x_i is x'_i and the sole predecessor of \bar{x}_i is \bar{x}'_i . For every pair of nodes g_i^j and g_i^k of G_i , if $j < k$ then g_i^j is a predecessor of g_i^k . Similarly, for every pair of nodes g_{i-1}^j and g_{i-1}^k of G_{i-1} , if $j < k$ then g_{i-1}^j is a predecessor of g_{i-1}^k . The subgraph $(\{g_i^1, \dots, g_i^{4i-1}\}, G_{i-1})$ forms an $4i - 1$ slide. The node b_i is a successor of every node in G_{i-1} , and the node a_i is a successor of every node in $G_{i-1} \cup \{b_i\}$. Finally, \bar{x}'_i is a predecessor of every node in $\{g_i^1, \dots, g_i^{4i}\}$, \bar{x}_i is a predecessor of b_i , d_i is a predecessor of both nodes in $\{b_i, a_i\}$, x'_i is also a predecessor

of both nodes in $\{b_i, a_i\}$, x_i is a predecessor of every node in $\{g_i^1, \dots, g_i^{4i+1}\}$, a_i is a predecessor of every node in $\{g_i^1, \dots, g_i^{4i+2}\}$, and y_i is a predecessor of every node in $\{g_i^1, \dots, g_i^{4i+3}\}$.

The existential widget is depicted in Figure 4. For every i , $1 \leq i \leq n$, if widget i is an existential widget, it is composed of 4 groups of nodes, $\{\bar{x}_i, \bar{x}'_i, d_i, x_i, x'_i, y_i\}$, $G_{i-1} = \{g_{i-1}^1, \dots, g_{i-1}^{4i-1}\}$, $R_i = \{r_i^1, \dots, r_i^{4i+1}\} \cup H_i = \{h_i^1, \dots, h_i^{4i+1}\} \cup \{a_i\}$, and $G_i = \{g_i^1, \dots, g_i^{4i+3}\}$. x'_i has $4i+3$ source nodes $p_{x'_i}^1$ through $p_{x'_i}^{4i+3}$ as predecessors, y_i has $4i+2$ source nodes $p_{y_i}^1$ through $p_{y_i}^{4i+2}$ as predecessors d_i has $4i+1$ source nodes $p_{d_i}^1$ through $p_{d_i}^{4i+1}$ as predecessors, and \bar{x}'_i has $4i$ source nodes $p_{\bar{x}'_i}^1$ through $p_{\bar{x}'_i}^{4i}$ as predecessors. \bar{x}'_i also has y_i and x'_i as predecessors. The sole predecessor of x_i is x'_i and the only two predecessors of \bar{x}_i are \bar{x}'_i and y_i . For every pair of nodes g_i^j and g_i^k of G_i , if $j < k$ then g_i^j is a predecessor of g_i^k . The same is true for every pair of nodes in H_i , R_i , and G_{i-1} . Every node $g_i^j \in \{g_i^1, \dots, g_i^{4i+1}\}$ has $4i+1-j$ source nodes as predecessors. Also, a_i is a predecessor of every node in $\{g_i^1, \dots, g_i^{4i+1}\}$, \bar{x}'_i is a predecessor of every node in $\{g_i^1, \dots, g_i^{4i+2}\}$, and x_i is a predecessor of every node in $\{g_i^1, \dots, g_i^{4i+3}\}$. Also, a_i is the successor of every node in H_i , d_i is a predecessor of every node in $\{h_i^1, \dots, h_i^{4i+1}\}$, \bar{x}_i is a predecessor of every node in $\{h_i^1, \dots, h_i^{4i}\}$ and $(\{h_i^1, \dots, h_i^{4i-1}\}, R_i)$ forms a $4i-1$ slide. Finally, y_i is a predecessor of every node in R_i and (R_i, G_{i-1}) forms a $4i-1$ slide.

For all i , $1 < i < n$, G_i is part of both widget i and widget $i+1$. G_0 is special in that it connects the string of quantifier widgets to the string of clause widgets and is described below. G_n is special because every node in G_n is a predecessor of the target node s . We now describe the m clause widgets.

For each clause C_i , there is a corresponding node z_i . This node always has four predecessors, one of which is the previous clause node z_{i-1} . The other three, l_i^1 , l_i^2 , and l_i^3 , correspond to the literals which occur C_i . For example, if the first literal in the i^{th} clause is \bar{x}_j , then the node \bar{x}_j from quantifier widget j is one of the predecessors of z_i . z_1 has a special source node z_0 as a predecessor, since it has no previous clause. Finally, we add edges from z_m to all three nodes of G_0 . There are also three source nodes a_0 , b_0 , and c_0 which are connected to G_0 . a_0 and b_0 are predecessors of g_0^1 and c_0 is a predecessor of g_0^2 . Figure 1 shows both an example of a clause widget as well the connection between z_m and G_0 . This completes the construction. Figure 5 shows the outline of an entire circuit for an example QBF.

4 Proof of PSPACE Completeness

Theorem 2: The quantified Boolean formula $\psi = Q_n x_n Q_{n-1} x_{n-1} \dots Q_1 x_1 F$ is *QSAT* if and only if vertex s in the graph \mathcal{G} constructed as above can be pebbled with $4n+3$ pebbles.

DEFINITION 4.1: Let the set of all truth assignments over variables x_{i+1}, \dots, x_n be denoted by A_i . Thus each α_i in A_i is a partial assignment that sets the outermost $n-i$ variables of $Q_n x_n \dots Q_1 x_1 F$. For any assignment to α_i , define B_{α_i} to be the pebbling configuration of \mathcal{G} consisting of black pebbles on the following nodes: For each universally quantified variable x_j of ψ , $j \geq i+1$, if $\alpha_i(x_j) = 0$, then $y_j \in B_{\alpha_i}$, $x'_j \in B_{\alpha_i}$, $d_j \in B_{\alpha_i}$, and $(\bar{x}_j, \bar{x}'_j) \in B_{\alpha_i}$. Otherwise, if $\alpha_i(x_j) = 1$, then $y_j \in B_{\alpha_i}$, $\bar{x}'_j \in B_{\alpha_i}$, $a_j \in B_{\alpha_i}$ and $(x_j, x'_j) \in B_{\alpha_i}$. For each existentially quantified variable x_j of ψ , $j \geq i+1$, if $\alpha_i(x_j) = 0$, then $y_j \in B_{\alpha_i}$, $x'_j \in B_{\alpha_i}$, $d_j \in B_{\alpha_i}$, and $(\bar{x}_j, \bar{x}'_j) \in B_{\alpha_i}$. Otherwise, if $\alpha_i(x_j) = 1$, then $y_j \in B_{\alpha_i}$, $\bar{x}'_j \in B_{\alpha_i}$, $d_j \in B_{\alpha_i}$ and $(x_j, x'_j) \in B_{\alpha_i}$.

DEFINITION 4.2 (Black clamping interval) Let $t_0 \leq t_j \leq t_k \leq t_{\text{end}}$. Let S be a set of nodes. We say that $S \in [t_a, t_b]$ if all nodes from S must be black pebbled during every configuration from time t_a through time t_b . We say that $(u, v) \in [t_a, t_b]$ if either u or v is black pebbled during every configuration from time t_a to time t_b .

Lemma 3: If ψ is QSAT, then the target node s of \mathcal{G} can be pebbled with $4n+3$ pebbles.

Lemma 3 follows from the following more general lemma by setting $i = n$.

Lemma 4: For all i , $\alpha_i \in A_i$, suppose the graph \mathcal{G} is initially in configuration B_{α_i} . If ψ is QSAT, then we can black pebble G_i at some time $t > 1$ using $4n+3$ pebbles, while keeping B_{α_i} clamped (i.e., $B_{\alpha_i} \in [1, t]$).

Proof: The proof is by induction on i from 0 to n . The base case is when $i = 0$. Let α_0 be any assignment in A_0 . Suppose that $Q_n x_n \cdots Q_1 x_1 F \upharpoonright_{\alpha_0}$ is QSAT. Then some literal in every clause must be set to true. This implies that for each z_j , $1 \leq j \leq m$, at least one of l_j^1 , l_j^2 , or l_j^3 are black pebbled in B_{α_0} . We can therefore black pebble G_0 as follows. Start by putting a black pebble on z_0 . Then since at most two of z_1 's other predecessors are unpebbled, we have enough free pebbles to black pebble the rest of z_1 's predecessors. We know we can black pebble them because if some l_1^k is unpebbled, then $l_1^{k'}$ must be black pebbled in B_{α_0} . We can therefore black pebble all of z_1 's predecessors. We can then slide the pebble from z_0 to z_1 and lift the other (at most 2) pebbles which we just put down. Once z_1 is black pebbled, we can then black pebble z_2 the same way, all the way to z_m . Once z_m is black pebbled we can use the remaining two black pebbles to black pebble a_0 and b_0 , and then slide the pebble from z_m to c_0 . We can then slide the black pebble from a_0 to g_0^1 , from b_0 to g_0^2 , and from c_0 to g_0^3 . Note that this strategy uses only black pebbles. For the inductive step there are two cases depending on whether Q_i is a universal or an existential quantifier.

Case 1: Q_i is a universal quantifier. In this case, both $\Psi \upharpoonright_{\alpha_i \cup \{x_i\}}$ and $\Psi \upharpoonright_{\alpha_i \cup \{\bar{x}_i\}}$ are QSAT. We begin in configuration B_{α_i} with $4i + 3$ free pebbles. Black pebble y_i , followed by x'_i , then d_i , and then \bar{x}'_i . Then move the pebble from \bar{x}'_i to \bar{x}_i . At this point we have $4i - 1$ pebbles free and can apply the induction hypothesis to black pebble G_{i-1} . Then slide the black pebble from \bar{x}_i to b_i , then the black pebble from d_i to a_i . Remove all pebbles from widget i except for the ones on a_i , x'_i , and y_i . Then slide the black pebble from x'_i to x_i and black pebble \bar{x}'_i again. Now apply the induction hypothesis to simultaneously black pebble G_{i-1} again. Next, use the i -slide strategy to slide all of G_{i-1} 's pebbles up to g_i^1 to g_i^{4i-1} . Then slide \bar{x}'_i 's black pebble to g_i^{4i} , and then x_i 's black pebble to g_i^{4i+1} . Next slide the black pebble from a_i to g_i^{4i+2} . Finally, slide the black pebble from y_i to g_i^{4i+3} .

Case 2: Q_i is an existential quantifier. In this case, either $\Psi \upharpoonright_{\alpha_i \cup \{x_i\}}$ or $\Psi \upharpoonright_{\alpha_i \cup \{\bar{x}_i\}}$ is QSAT. As in the universal case, we begin in B_{α_i} with $4i + 3$ free pebbles. Black pebble x'_i , followed by y_i , d_i , and then \bar{x}'_i .

If $\Psi \upharpoonright_{\alpha_i \cup \{x_i\}}$ is QSAT, move the black pebble from x'_i to x_i . Then apply the induction hypothesis to black pebble G_{i-1} . Then use the i -slide strategy to move all of the pebbles from G_{i-1} to R_i . Then slide the black pebble from y_i to \bar{x}_i . Then use the i -slide strategy to move all of the pebbles from R_i to $\{h_i^1, \dots, h_i^{4i-1}\}$. After that, slide the pebble from \bar{x}_i to h_i^{4i} and then slide the pebble from d_i to h_i^{4i+1} . Then slide the pebble from h_i^{4i+1} to a_i . At this point remove all the pebbles off of the widget so that only \bar{x}'_i , x_i , and a_i remain. Use these $4i$ free pebbles to pebble the source node predecessors of g_i^1 and then slide one to g_i^1 itself. Use the pebbles left over on the source nodes to subsequently pebble each g_i^j until g_i^{4i} is pebbled. At this point slide the pebble from a_i to g_i^{4i+1} , slide the pebble from \bar{x}'_i to g_i^{4i+2} , and finish by sliding the pebble from x_i to g_i^{4i+1} .

If $\Psi \upharpoonright_{\alpha_i \cup \{\bar{x}_i\}}$ is QSAT, move the black pebble from \bar{x}'_i to \bar{x}_i . Then apply the induction hypothesis to black pebble G_{i-1} . Then use the i -slide strategy to move all of the pebbles from G_{i-1} to R_i . Then use the i -slide strategy to move all of the pebbles from R_i to $\{h_i^1, \dots, h_i^{4i-1}\}$. After that, slide the pebble from \bar{x}_i to h_i^{4i} and then slide the pebble from d_i to h_i^{4i+1} . Then slide the pebble from h_i^{4i+1} to a_i . At this point remove all the pebbles off of the widget so that only y_i , x'_i , and a_i remain. Use the $4i$ pebbles that are free to repebble \bar{x}'_i and then pick the pebble up from y_i and pick up the $4i - 1$ pebbles that remain on \bar{x}'_i source node predecessors. Slide the pebble from x'_i to x_i . At this point \bar{x}'_i , x_i , and a_i are all pebbled and we can finish by black pebbling G_i as we did in the positive case. \square

Lemma 5: Let ψ be a QBF, and let \mathcal{G} be the corresponding graph. If s has a $4n + 3$ black-white pebbling strategy in \mathcal{G} , then ψ is QSAT, and any $4n + 3$ black-white pebbling strategy requires $\Omega(2^k)$ steps, where k is the number of universal quantifiers in ψ .

We first note that s has $4n + 3$ predecessors, G_n . And each of these nodes has indegree $4n + 3$. So no node of G_n could ever contain a white pebble while s contains a black pebble, because there would not be enough free pebbles to discharge it. Therefore, in order to pebble s , G_n must first be simultaneously black pebbled. Lemma 5 therefore follows from the following more general theorem.

Lemma 6: For all $\alpha_i \in A_i$, if there exists times t', t'' such that $B_{\alpha_i} \subseteq [t', t'']$, then black pebbling G_i at t'' from B_{α_i} using no more than $4n + 3$ pebbles, requires that ψ is QSAT and requires $\Omega(2^k)$ units of time between t' and t'' , where k is the number of universal quantifiers among the i inner most quantifiers.

The following lemma will be used repeatedly. In particular, it implies that for any i -slide (V, U) , in order to pebble V using no more than i pebbles, U must first be black pebbled at some earlier time.

Lemma 7: If a node v has k predecessors and there are $4n + 3 - k$ other nodes in $[t', t'']$ and v is not white pebbled at t'' , then v can be black pebbled at most once and can never be white pebbled between t' and t'' .

Proof: If v is white pebbled, then its white pebble can only be discharged once it has contributed toward placing a black pebble beyond it. The existence of this extra black pebble means that there are at most $k - 1$ free pebbles to pebble all of v 's k predecessors. So the space bound must be exceeded to discharge the white pebble. The same argument forbids a second black pebbling. \square

Proof: [of Lemma 6] The proof is by induction on i from 0 to n . The base case is when $i = 0$. Let α_0 be any assignment in A_0 and suppose there exist times t' and t'' such that $B_{\alpha_0} \subseteq [t', t'']$. We will show that simultaneously black pebbling G_0 at t'' without ever exceeding $4n + 3$ pebbles requires that ψ is QSAT.

In order to black pebble z_j or discharge a white pebble from z_j we must either black pebble z_{j-1} or discharge a white pebble from z_{j-1} . In order to black pebble any node in G_0 , we must pebble z_m . Inductively, this means that at some point for every single z_j , it was necessary to either black pebble it or discharge a white pebble from it. But every z_j (except z_0) has 4 predecessors, $l_j^1, l_j^2, l_j^3, z_{j-1}$. Therefore, in order to pebble z_j at least one l_j^k must be black pebbled in B_{α_0} . But in this case, α_0 must satisfy clause j of F . Since every z_j must either be black pebbled or discharged, α_0 must satisfy every clause of F . Therefore $F \upharpoonright_{\alpha_0}$ is QSAT.

Induction Step: We now prove the induction step in which we will show that if we can simultaneously black pebble $G_i = \{g_i^1 \dots g_i^{4i+3}\}$ using no more than $4i + 3$ pebbles without moving any pebbles in B_{α_i} , then $\psi \upharpoonright_{\alpha_i}$ is QSAT and the pebbling must take time $\Omega(2^k)$, where k is the number of universally quantified variables among the inner most i variables of ψ .

Case 1: Q_i is a universal quantifier. We will show that in order to black pebble G_i we must necessarily pass through a number of all-black configurations, including black pebbling G_{i-1} twice, once with black pebbles on x'_i, d_i , and either \bar{x}_i or \bar{x}'_i (the false configuration), and once with black pebbles on \bar{x}'_i, a_i , and either x_i or x'_i (the true configuration).

We appeal to Lemma 7 to conclude that since y_i has $4i + 3$ source nodes as predecessors, our first action within widget i must be to black pebble y_i and it must stay in place until its last successor g_i^{4i+3} is pebbled for the final time at t_{15} , so $y_i \in [t_1, t_{15} - 1]$.

Now that y_i is clamped, we can again appeal to Lemma 7 to conclude that no node in $G_i \cup \{a_i, b_i, x'_i\}$ can be white pebbled and each can only be black pebbled once between t_1 and t_{15-1} . Since x'_i has $4i + 2$ source nodes as predecessors, our second action within widget i must be to black pebble x'_i and it must stay in place until its successor x_i is pebbled for the last time. Then a pebble must remain on x_i until all of its successors are pebbled for the last time, because we can never repebble/discharge x_i once x'_i is empty. Let t_7 be the time that a_i is pebbled and let t_{12} be the time g_i^{3i} is pebbled. Then $x'_i \in [t_2, t_7 - 1]$ and $(x_i, x'_i) \in [t_7, t_{12} - 1]$.

Our argument now divides into two sections. In order to simultaneously black pebble G_i we must black pebble g_i^{4i+3} , which requires that both a_i and $\{g_i^1, \dots, g_i^{4i}\}$ be pebbled. In the first part of the argument we prove that in order to black pebble a_i , $\psi \upharpoonright_{\alpha_i \cup \{\bar{x}_i\}}$ must be QSAT and that $\Omega(2^k)$ units of time must pass between t_0 and t_7 , where k is the number of universally quantified variables among the inner most $i - 1$ variables of ψ . In the second part of the argument, we argue that g_i^1, \dots, g_i^{4i} must also be simultaneously black pebbled in order to black pebble g_i^{4i+3} and that pebbling them without exceeding our bound necessitates that $\psi \upharpoonright_{\alpha_i \cup \{x_i\}}$ is QSAT and that $\Omega(2^k)$ units of time pass between times t_7 and $t_{14} - 1$. This will allow us to

conclude that black pebbling G_i requires that $\psi \upharpoonright_{\alpha_i}$ is QSAT and requires $\Omega(2^{k'})$ time, where $k' = k + 1$ is the number of universally quantified variables among the inner most i variables of ψ .

Since a_i can only be black pebbled once and is needed to pebble each node of G_i , $a_i \in [t_7, t_{14} - 1]$. In order to black pebble a_i at time t_7 we must pebble b_i at some time t_6 , before t_7 . Again, we know that b_i can only be black pebbled once in t_1 to t_{14} , so $b_i \in [t_6, t_7 - 1]$. Also, d_i is a predecessor of both a_i and b_i and must be pebbled at times $t_6 - 1$ and $t_7 - 1$. Since x'_i is in $[t_2, t_7]$, by Lemma 7 we can conclude that d_i cannot be white pebbled and can only be black pebbled once in this interval. Also, since it has in-degree $4i$, d_i must be black pebbled at t_3 , immediately after t_2 as in Lemma 3, so $d_i \in [t_3, t_7 - 1]$. The same argument can be made to argue that $(\bar{x}_i, \bar{x}'_i) \in [t_4, t_6 - 1]$, where t_4 is after t_3 . In order to black pebble a_i or b_i , we must first pebble G_{i-1} at some time t_5 before t_6 . This whole time the nodes x'_i , d_i , and (\bar{x}_i, \bar{x}'_i) are clamped. We can therefore apply Lemma 7 to conclude that G_{i-1} must be black pebbled at some time t_5 between t_4 and t_6 . We can now apply the induction hypothesis to conclude that black pebbling G_{i-1} requires $\psi \upharpoonright_{\alpha_i \cup \{\bar{x}_i\}}$ to be QSAT and black pebbling G_{i-1} from $B[t_4]$ requires time $\Omega(2^k)$, where k is the number of universally quantified variables among the inner most $i - 1$ variables of ψ .

We now proceed with the second phase of the argument. We know that each node in G_i cannot be white pebbled and can only be black pebbled once. So when we black pebble g_i^{4i+3} at time t_{15} , all the rest of G_i must already be black pebbled. Consider g_i^{4i+2} . In order to black pebble it at time t_{14} before t_{15} , we must first black pebble g_i^{4i+1} at time t_{13} before t_{14} . In order to black pebble g_i^{4i+1} at time t_{13} we must first black pebble g_i^{4i} at time t_{12} and in order to pebble that, we must pebble g_i^1, \dots, g_i^{4i-1} at time t_{11} . But we must also pebble \bar{x}'_i . Note that \bar{x}'_i must be empty at t_7 since y_i is clamped and a_i has $4i + 2$ other predecessors, none of which is \bar{x}'_i . Also, \bar{x}'_i must be empty again by $t_{13} - 1$, since g_i^{4i+1} has $4i + 3$ predecessors, none of which is \bar{x}'_i . We can therefore apply Lemma 7 to conclude that between t_7 and t_{13} , \bar{x}'_i cannot be white pebbled and can only be black pebbled once in that interval. We must therefore repebble \bar{x}'_i at some time t_8 after t_7 when a_i and (x_i, x'_i) are clamped and $\bar{x}'_i \in [t_8, t_{12} - 1]$. Since \bar{x}'_i is a predecessor of every node in g_i^1, \dots, g_i^{4i-1} , these nodes can only be black pebbled at some time t_{11} , with g_i^1 being pebbled first at t_{10} , after t_8 . Every node of G_{i-1} is a predecessor of g_i^1 . Since the three nodes $\{\bar{x}'_i, a_i, (x_i, x'_i)\}$ are clamped during the interval $[t_7, t_{11}]$ we can apply Lemma 7 to conclude that G_{i-1} must be black pebbled at t_9 between t_8 and t_{10} . Since $\{\bar{x}'_i, a_i, (x_i, x'_i)\}$ is the true assignment for variable x_i we can apply our induction hypothesis to conclude that $\psi \upharpoonright_{\alpha_i \cup \{x_i\}}$ must be QSAT and black pebbling G_{i-1} from $B[t_7]$ requires time $\Omega(2^k)$, where k is the number of universally quantified variables among the inner most $i - 1$ variables of ψ .

Thus we have shown that any $4n + 3$ pebbling must black pebble G_{i-1} twice between t_0 and t_{15} , once implying that $\psi \upharpoonright_{\alpha_i \cup \{\bar{x}_i\}}$ is QSAT, and once implying that $\psi \upharpoonright_{\alpha_i \cup \{x_i\}}$ is QSAT. Each time requires $\Omega(2^k)$ time, where k is the number of universally quantified variables among the inner most $i - 1$ variables of ψ . Therefore, black pebbling G_i requires time $\Omega(2^{k+1})$, and implies that $\psi \upharpoonright_{\alpha_i}$ is QSAT.

Case 2: Q_i is an existential quantifier. We will show that in order to black pebble G_i , we must necessarily pass through a number of all-black partial configurations, including simultaneously black pebbling G_{i-1} , either with black pebbles on x'_i , d_i , and either \bar{x}_i or \bar{x}'_i (the false configuration), or with black pebbles on \bar{x}'_i , d_i , and either x_i or x'_i (the true configuration).

By Lemma 7, no node in $G_i \cup \{x'_i\}$ can be white pebbled between t_0 and t_{15} , and each can be black pebbled at most once. Based on which nodes of G_i are predecessors to others, we can conclude that g_i^{4i+3} must be black pebbled last, at time t_{15} , g_i^{4i+2} must be black pebbled before that at time t_{14} and $g_i^{4i+2} \in [t_{14}, t_{15}]$, and g_i^{4i+1} must be pebbled before that at time t_{13} and $g_i^{4i+1} \in [t_{13}, t_{15}]$, g_i^1 must be pebbled before that at time t_{12} and $g_i^1 \in [t_{12}, t_{15}]$. Also, $(x_i, x'_i) \in [t_1, t_{15} - 1]$.

Now consider y_i . It has degree $4i + 2$, and it must be black pebbled at time t_2 , and can never be repebbled again. Thus it must remain black pebbled until it is used for the last time.

Clearly, both $\bar{x}'_i \in B[t_{12} - 1]$ and $a_i \in B[t_{12} - 1]$. Let t_{11} be the last time a_i is pebbled. At this time, a_i must be pebbled black. We can see this because g_i^{4i+1} cannot get its black pebble from g_i^1 through to g_i^{4i} since these can only be pebbled once. All of these must be in place when g_i^{4i+1} gets its black pebble, so it

cannot get a black pebble from either x_i or \bar{x}'_i since both of these are needed to support g_i^{4i+2} and could not be repebbled with so many black pebbles clamped in G_i . g_i^{4i+1} 's $4i+3$ rd predecessor is a_i , so it must receive its black pebble via a slide move from a_i . So a_i must be black during the interval $[t_{11}, t_{12} - 1]$.

At this point our proof splits into two cases, either a black pebble is on \bar{x}'_i at t_{11} or not. One of these cases will imply that $\psi^{\lceil \alpha_i \cup \{x_i\}} is QSAT and the other one will imply that $\psi^{\lceil \alpha_i \cup \{\bar{x}_i\}} is QSAT.$$

Suppose there is no black pebble on \bar{x}'_i at t_{11} . Then there are two subcases to consider. Subcase (i): there is no pebble at all on \bar{x}'_i or subcase (ii) there is a white pebble on \bar{x}'_i at t_{11} . First we consider subcase (i): there is no pebble at all on \bar{x}'_i . Then we must repebble \bar{x}'_i at some time t^* between t_{11} and $t_{12}-1$. We will first argue that two nodes, x'_i and y_i must be clamped during the interval $[t_2, t_{11}]$. First, because x'_i is black pebbled at t_1 , and is a predecessor of \bar{x}'_i , and can never be pebbled again (because its indegree is $4i+3$), it follows that $x'_i \in [t_1, t^* - 1]$. Secondly, since y_i is a predecessor of \bar{x}'_i (and by the above reasoning gets black pebbled only once at t_2), it follows that $y_i \in [t_2, t^* - 1]$. Thus both x'_i and y_i are clamped during the interval $[t_2, t_{11}]$.

Now we will argue that each node of H_i must be black pebbled, and can only be pebbled once. Let t_{10} be the time when h_i^{4i+1} is pebbled; let t_9 be the time when h_i^{4i} is pebbled; let t_8 be the time when h_i^{4i-1} is pebbled, and let t_7 be the time when h_i^1 is pebbled, where $t_7 < t_8 < t_9 < t_{10}$. By Lemma 7 and because x'_i and y_i are clamped, and all nodes in H_i have indegree $4i+1$, it follows that each can only be pebbled once and must be pebbled black. Thus, $h_i^{4i+3} \in [t_{10}, t_{11} - 1]$, $h_i^{4i+2} \in [t_9, t_{11} - 1]$, $h_i^{4i+1} \in [t_8, t_{11} - 1]$, and $h_i^1 \in [t_7, t_{11} - 1]$.

Next we will argue that during the interval $[t_3, t_{10} - 1]$, the three nodes d_i , x'_i and y_i are all black clamped. (We already know that x'_i and y_i are black clamped during this interval.) Because d_i has indegree $4i+1$, by Lemma 7, again we know that d_i must be black pebbled at time t_3 and can only be black pebbled once. Thus, d_i is black and clamped during the interval $[t_3, t_{10} - 1]$.

Now again we can apply Lemma 7 to R_i . Because now we know that 3 nodes are clamped during this interval, and because all nodes in R_i have degree $4i$, it follows that they can only be pebbled once between t_3 and $t_{10} - 1$ and are black. Let t_6 be the time r_i^1 is pebbled, $t_6 < t_7$.

Finally, we want to show that $(\bar{x}_i, \bar{x}'_i) \in [t_3 + 1, t_9 - 1]$ and furthermore the pebbled node is black. First, \bar{x}_i must be pebbled at time $t_7 - 1$ because it is a predecessor of h_i^1 . Furthermore we will argue that it must be black pebbled. At time $t_7 - 1$, \bar{x}'_i must be unpebbled because in order to pebble h_i^1 at time t_7 , there must be $4i+3$ pebbles already on this i^{th} widget, not including \bar{x}'_i (the $4i+1$ predecessors of h_i^1 plus the clamped nodes x'_i and y_i .) Similarly, \bar{x}'_i must unpebbled at t_9 . Now if \bar{x}_i were pebbled white rather than black at $t_7 - 1$, it would have to be discharged by t_9 ; but this cannot happen since it would have to be discharged through the unpebbled \bar{x}'_i , which would exceed our allowable space. Thus we have argued that \bar{x}_i must be pebbled black at $t_7 - 1$, and further remains black until $t_9 - 1$ since it is a predecessor of all h_i^1, \dots, h_i^{4i} .

Now to black pebble \bar{x}_i by $t_7 - 1$, \bar{x}'_i must be pebbled earlier, say at time t_4 , $t_3 < t_4 < t_7 - 1$. It is left to argue that $t_4 = t_3 + 1$. When we black pebble \bar{x}'_i at time t_4 , we have already argued that there are three nodes already clamped, x'_i , y_i and d_i . Because \bar{x}'_i has indegree $4i$, it follows that it must be black pebbled next, and can only be pebbled once. Thus $(\bar{x}_i, \bar{x}'_i) \in [t_3 + 1, t_9 - 1]$.

Now in order to black pebble r_i^1 at t_6 , every node of G_{i-1} must be pebbled at $t_6 - 1$. Again we can apply Lemma 7. Since there are 4 nodes clamped, and the degree of each node in G_i is $4i - 1$, it follows by our lemma that every node in G_i can only be pebbled once between t_4 and t_9 and must be black pebbled. Now finally we can apply the induction hypothesis to conclude that since every node in $\{(\bar{x}_i, \bar{x}'_i), x'_i, d_i, y_i\}$ is clamped while G_i is being black pebbled, $\psi^{\lceil \alpha_i \cup \{\bar{x}_i\}} is QSAT.$

The other subcase (ii) is an analogous argument to subcase (i) but for the dual case of a white pebble being discharged from a node (rather than the node being black pebbled.)

Suppose, on the other hand, that there is a black pebble on \bar{x}'_i at t_{11} . We will now show that only $B_{\alpha_i} \cup \{x'_i, d_i, y_i\}$ can be pebbled when we pebble \bar{x}'_i for the last time before t_{11} at some time t_4 . Suppose for the sake of contradiction that there is a pebble on some other node z at t_4 .

Since y_i is a predecessor of \bar{x}'_i and can only be pebbled at t_2 , $y_i \in [t_2, t_4 - 1]$. So d_i must be empty at $t_4 - 1$ because \bar{x}'_i has $4i+2$ predecessors which must be on the graph, along with z , at $t_4 - 1$, which fills up the

space bound.

In order to pebble a_i by t_{11} we must therefore pebble d_i at some time between t_4 and t_{11} . Suppose d_i is white pebbled. This pebble must be discharged by $t_{11} - 1$ because a_i has $4i + 1$ predecessors and both (x'_i, x_i) and \bar{x}'_i are clamped until t_{11} , so d_i 's pebble is needed. By frugality there must be a pebble in H_i at the time d_i is discharged. So at this time there must be pebbles on a node of H_i , one of (x_i, x'_i) , and \bar{x}_i and we must exceed the space bound. Suppose on the other hand that d_i is black pebbled between t_4 and t_{11} . This takes $4i + 1$ pebbles and there must be pebbles on (x_i, x'_i) , \bar{x}'_i and by frugality z' , where z' is between z and a_i . So we can never pebble d_i between t_4 and t_{11} . We therefore know that when \bar{x}'_i is pebbled for the last time before t_{11} , there can be no pebble on z .

By the argument which we just finished, any node of G_{i-1} can only be pebbled after t_4 . We now show that G_{i-1} must be simultaneously black pebbled in order to black pebble a_i .

We know that both $\bar{x}'_i \in [t_4, t_{11}]$ and $(x'_i, x_i) \in [t_4, t_{11}]$. Therefore by Lemma 7, any node in H_i can only be pebbled once in $[t_4, t_{11}]$ and must be black. Call the time h_i^{4i+1} is pebbled t_{10} , the time h_i^{4i} is pebbled t_9 , the time h_i^{4i-1} is pebbled t_8 , and the time h_i^1 is pebbled t_7 . So \bar{x}_i must pebbled at some time t_6 before t_7 and $\bar{x}_i \in [t_6, t_7 - 1]$. Suppose it is white pebbled. Then it must be discharged before t_{10} because its pebble is needed to pebble h_i^{4i+1} . Note that y_i must be empty at $t_7 - 1$ since our space bound is reached by h_i^1 's predecessors and the clamping of (x_i, x'_i) and \bar{x}'_i . So when \bar{x}_i is discharged, there can be no pebble on y_i . Therefore, to discharge \bar{x}_i , y_i must be pebbled again after t_7 and before t_{11} , which is impossible due to its high indegree. Suppose, on the other hand that \bar{x}_i is black pebbled at t_6 . This means that $y_i \in [t_2, t_6 - 1]$. So there are at least 3 pebbles clamped from t_4 until $t_7 - 1$. But d_i must be pebbled before $t_7 - 1$. So d_i must be pebbled before t_4 , at some time t_3 after t_2 , and $d_i \in [t_3, t_{10} - 1]$.

Thus $\{(x_i, x'_i), \bar{x}'_i, d_i, (y_i, \bar{x}_i)\} \subseteq [t_4, t_7 - 1]$, so by Lemma 7 any node of R_i can only be pebbled black and pebbled once during this interval. Let t_5 be the time r_i^1 is pebbled. The nodes $G_{i-1} \cup \{(x_i, x'_i), \bar{x}'_i, d_i, y_i\}$ must all be pebbled at $t_5 - 1$. So $\{(x_i, x'_i), \bar{x}'_i, d_i, y_i\} \in [t_4, t_5 - 1]$. So G_{i-1} must only be pebbled black and once during this interval, so we can apply the induction hypothesis for the true assignment to conclude that $\Psi \upharpoonright_{\alpha_i \cup \{x_i\}}$ is *QSAT*. \square

Corollary 8: There exists an infinite family of graphs such that any minimal space black-white pebbling of these graphs requires exponential-time, but they can be refuted in linear time with the use of 1 additional pebble.

Proof: Let \mathcal{G} be the DAG corresponding to the formula $\psi = \forall x_n \forall x_{n-1} \dots \forall x_1 (x_1 \vee \bar{x}_1 \vee x_2) \wedge (x_2 \vee \bar{x}_2 \vee x_3) \wedge \dots \wedge (x_n \vee \bar{x}_n \vee x_1)$. This formula is clearly *QSAT*, since its 3CNF part is a tautology. Also, since ψ has n universally quantified variables, by Lemma 5, the minimal $4n + 3$ pebbling strategy requires time 2^n to execute. But using just 1 additional pebble we can pebble the target node in linear time as follows.

For every i from n to 1, pebble y_i , followed by x'_i , and then d_i , and finally \bar{x}'_i . We have now placed $4n$ pebbles on the graph. Then use the remaining 4 pebbles to pebble through the clause widgets in the obvious way. Once z_m is reached, pebble G_0 . We must now pebble up through the n universal quantifier widgets in a similar way to the proof Lemma 4 except that we will use our extra pebble so that we will not have to repebble each twice. At the start we have pebbles on G_{i-1} , y_i , x'_i , d_i , and \bar{x}'_i , as well as the 4 pebbles in every widget above widget i , for a total of $4n + 3$ pebbles. Place the extra pebble on \bar{x}_i and slide it up to b_i . Then slide the pebble from d_i to a_i and lift the pebble from b_i . We now have pebbles on G_{i-1} , \bar{x}'_i , a_i , x'_i and y_i . At this point we can slide the pebble from x'_i to x_i and continue to pebble G_i without exceeding $4i + 3$ pebbles again. We can follow this procedure for every i from 1 to n , and then slide a pebble from G_n up to s . Clearly, each clause widget and each quantifier widget must only be pebbled once, so the whole procedure requires linear time in the size of \mathcal{G} . \square

5 Figures

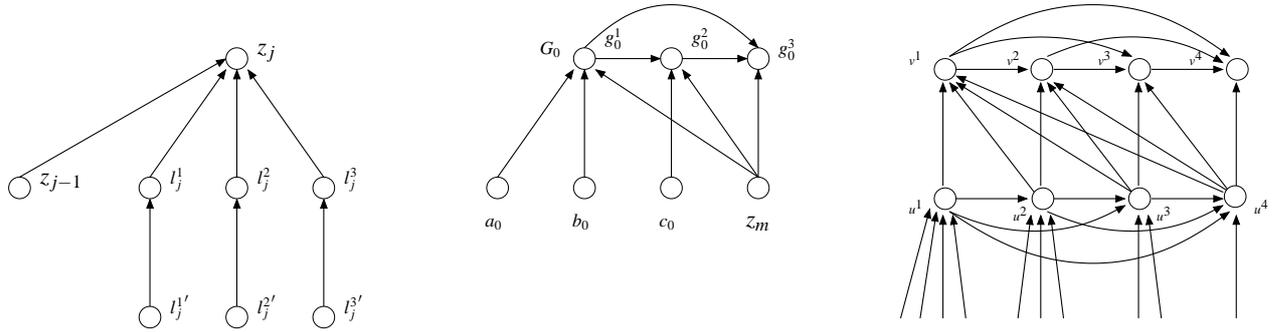


Figure 1: A clause widget for clause $z_j = (l_j^1 \vee l_j^2 \vee l_j^3)$ (left). The connection of z_m to G_0 (center). And a 4-slide ($\{v^1, v^2, v^3, v^4\}, \{u^1, u^2, u^3, u^4\}$) (right).

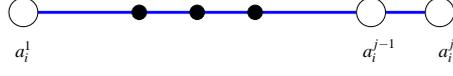
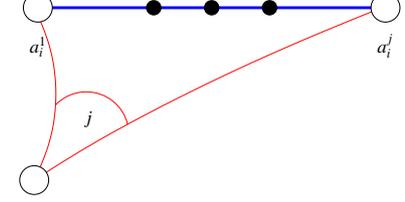
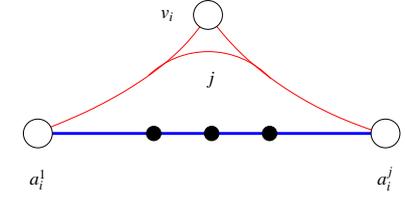
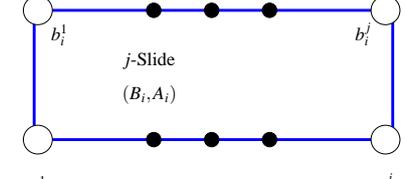
 <p style="text-align: center;">v_i</p> <p style="text-align: center;">j source nodes</p>	<p>The node v_i has j source nodes as predecessors.</p>
 <p style="text-align: center;">A_i</p> <p style="text-align: center;">a_i^1 a_i^{j-1} a_i^j</p>	<p>For every pair of nodes a_i^j and a_i^k in $A_i = \{a_i^1, \dots, a_i^m\}$, if $k < j$, then a_i^k is a predecessor of a_i^j.</p>
 <p style="text-align: center;">A_i</p> <p style="text-align: center;">a_i^1 a_i^j</p> <p style="text-align: center;">v_i</p>	<p>The node v_i is a predecessor of all j nodes in A_i.</p>
 <p style="text-align: center;">A_i</p> <p style="text-align: center;">a_i^1 a_i^j</p> <p style="text-align: center;">v_i</p>	<p>The node v_i is a successor of all j nodes in A_i.</p>
 <p style="text-align: center;">B_i</p> <p style="text-align: center;">b_i^1 b_i^j</p> <p style="text-align: center;">j-Slide (B_i, A_i)</p> <p style="text-align: center;">A_i</p> <p style="text-align: center;">a_i^1 a_i^j</p>	<p>There is a j-slide from A_i up to B_i.</p>

Figure 2: Legend explaining the components of Figures 3 and 4.

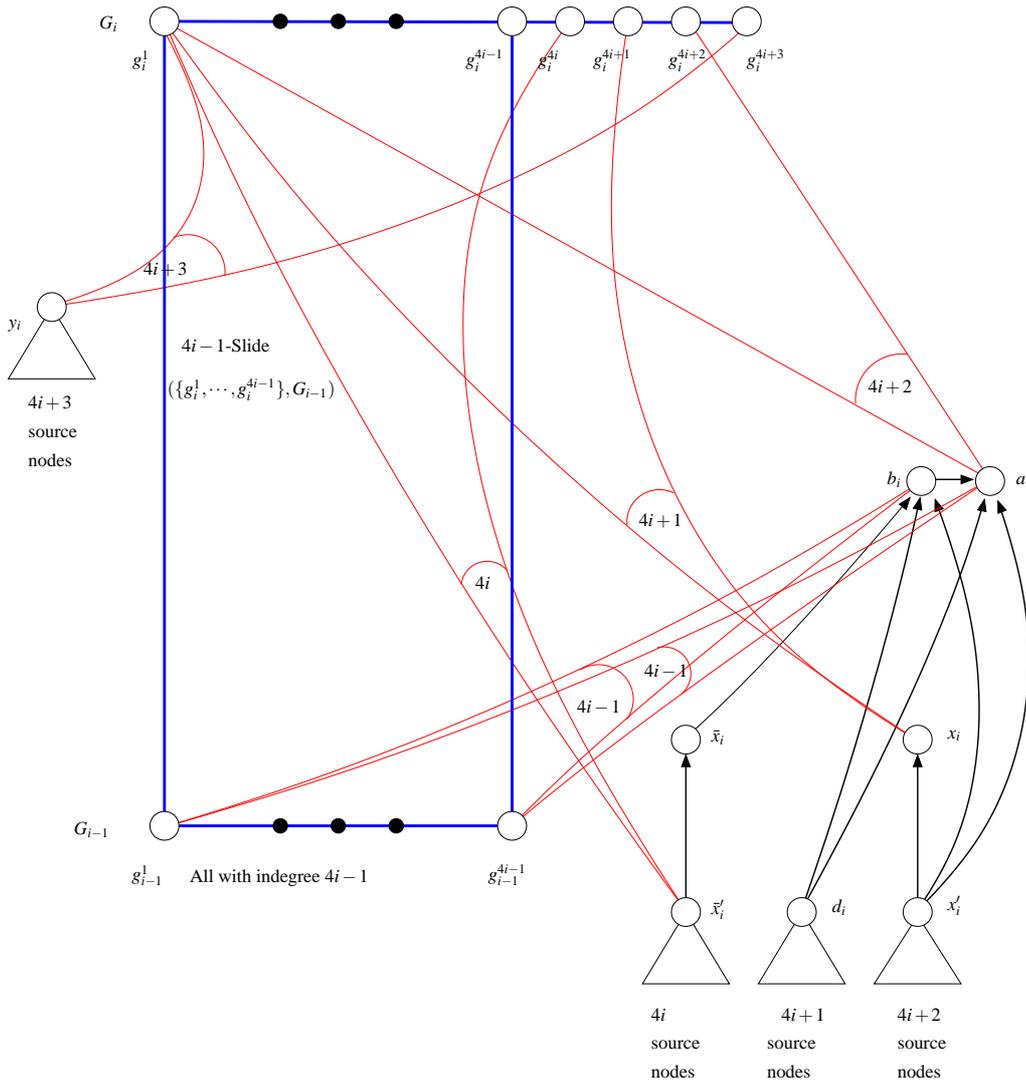


Figure 3: A universal widget.

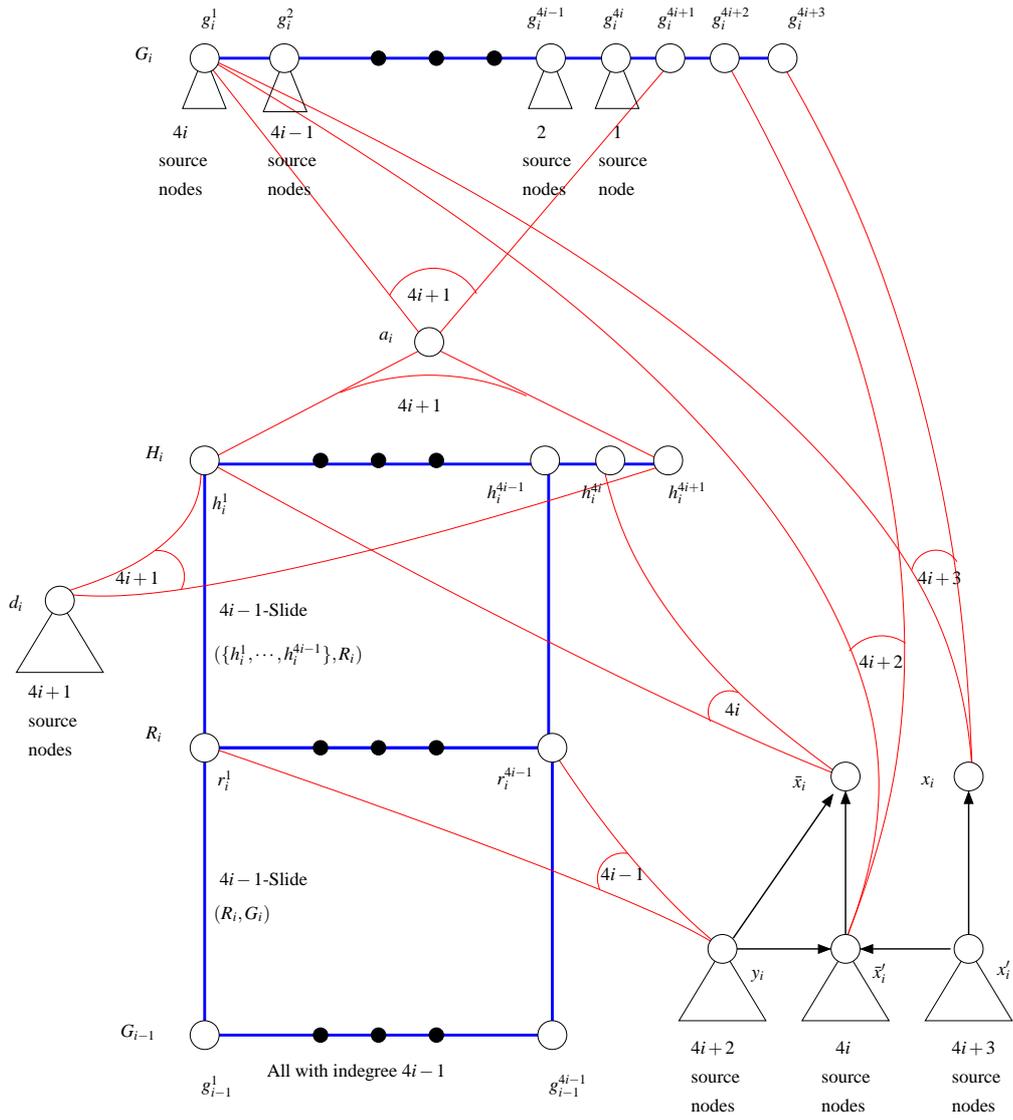


Figure 4: An existential widget.

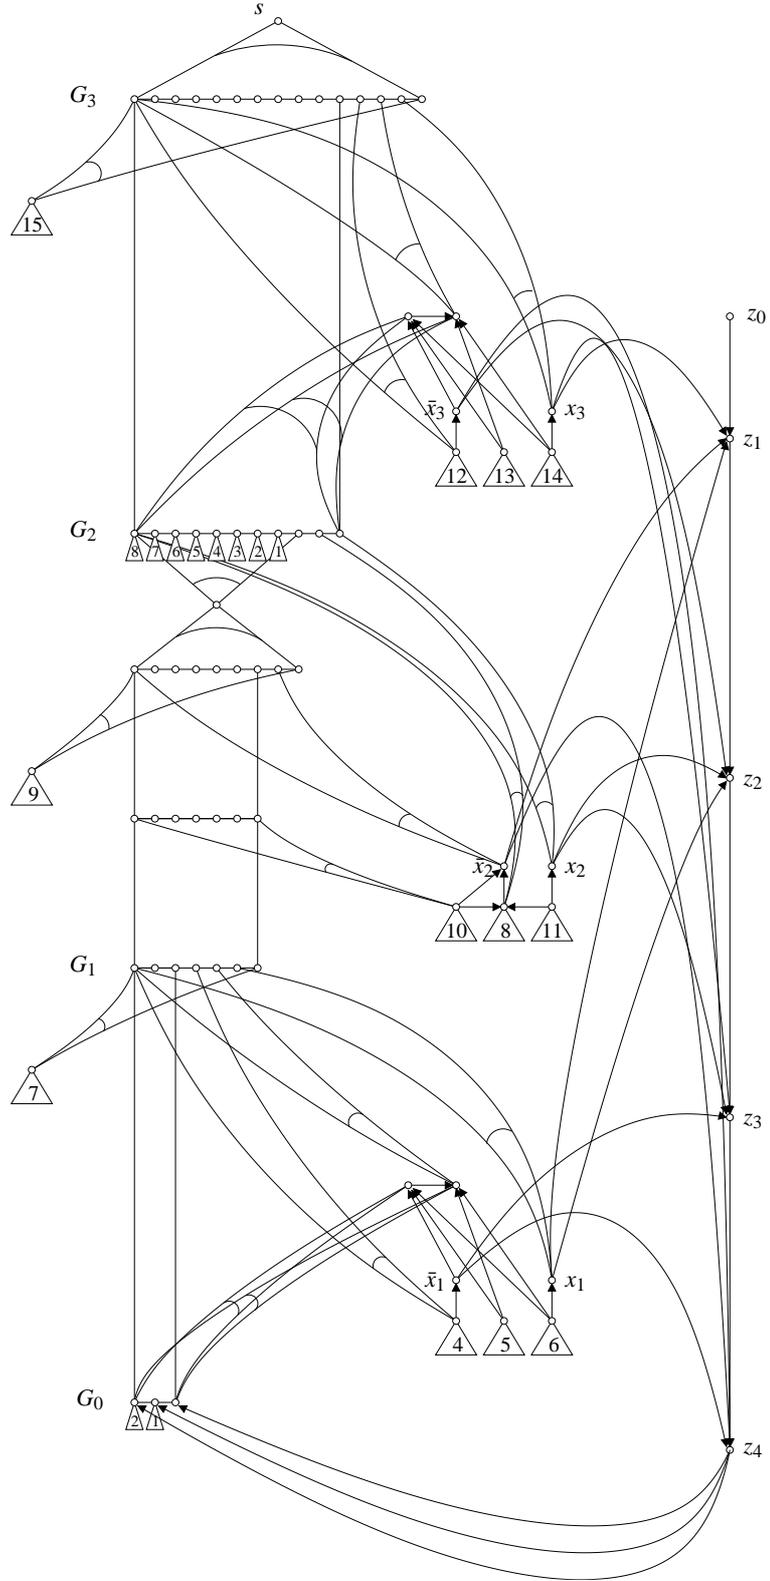


Figure 5: DAG generated for $\psi = \forall x_3 \exists x_2 \forall x_1 (x_1 \vee \bar{x}_2 \vee x_3) \wedge (x_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee x_3)$

References

- [1] Michael Alekhovich, Jan Johannsen, Toniann Pitassi, and Alasdair Urquhart. An exponential separation between regular and general resolution. *Electronic Colloquium on Computational Complexity (ECCC)*, 8(056), 2001.
- [2] E. Ben-Sasson. Size Space Tradeoffs For Resolution. *Proceedings of the 34th ACM Symposium on the Theory of Computing*, pages 457 – 464, 2002.
- [3] S. Cook and R. Sethi. Storage Requirements for Deterministic Polynomial Time Recognizable Languages. *Journal of Computer and System Sciences*, pages 25 – 37, 1976.
- [4] J. Esteban and J. Torán. Space Bounds for Resolution. *Information and Computation*, 171:84 – 97, 2001.
- [5] J. Esteban and J. Torán. A Combinatorial Characterization of Treelike Resolution Space. *Electronic Colloquium on Computational Complexity*, 44, 2003.
- [6] J. R. Gilbert, T. Lengauer, and R. E. Tarjan. The Pebbling Problem is Complete in Polynomial Space. *SIAM Journal of Computing*, Vol. 9, Issue 3:513 – 524, 1980.
- [7] F. Meyer Auf Der Heide. A Comparison of Two Variations of a Pebble Game on Graphs. *Theoretical Computer Science*, 13:315 – 322, 1981.
- [8] P. Hertel and T. Pitassi. Exponential Time/Space Speedups For Resolution. Submitted to *FOCS 2007*, 2007.
- [9] David S. Johnson. The NP-Completeness Column: An Ongoing Guide. *J. Algorithms*, 4(4):397–411, 1983.
- [10] Balasubramanian Kalyanasundaram and George Schnitger. On the power of white pebbles. In *STOC '88: Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 258–266, New York, NY, USA, 1988. ACM Press.
- [11] Maria M. Klawe. A tight bound for black and white pebbles on the pyramid. *J. ACM*, 32(1):218–228, 1985.
- [12] A. Lingas. A PSPACE-Complete Problem Related to a Pebble Game. In *Proceedings of the Fifth Colloquium on Automata, Languages and Programming*, pages 300 – 321, London, UK, 1978. Springer-Verlag.
- [13] J. Nordström. Narrow Proofs May Be Spacious: Separating Space and Width in Resolution. *Proceedings of the 38th ACM Symposium on the Theory of Computing*, 2006.
- [14] N. Pippenger. A Time-Space Trade-Off. *Journal of the ACM*, pages 509–515, 1978.
- [15] N. Pippenger. Pebbling. *Proceedings of the 5th IBM Symposium on Mathematical Foundations of Computer Science, Japan (Technical Report RC8528, IBM Watson Research Center)*, 1980.
- [16] R. Wilber. White pebbles help. In *STOC '85: Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 103–112, New York, NY, USA, 1985. ACM Press.