

Understanding Space in Resolution: Optimal Lower Bounds and Exponential Trade-offs

Eli Ben-Sasson*

Computer Science Department
Technion — Israel Institute of Technology
Haifa, 32000, Israel
eli@cs.technion.ac.il

Jakob Nordström†

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology‡
Cambridge, MA 02139, USA
jakobn@mit.edu

November 18, 2008

Abstract

For current state-of-the-art satisfiability algorithms based on the DPLL procedure and clause learning, the two main bottlenecks are the amounts of time and memory used. Understanding time and memory consumption, and how they are related to one another, is therefore a question of considerable practical importance. In the field of proof complexity, these resources correspond to the length and space of resolution proofs for formulas in conjunctive normal form (CNF). There has been a long line of research investigating these proof complexity measures, but while strong results have been established for length, our understanding of space and how it relates to length has remained quite poor.

The key technical contribution of this paper is the following, somewhat surprising, theorem: Any CNF formula F can be transformed by simple substitution into a new formula F' such that if F has the right properties, F' can be proven in resolution in essentially the same length as F but the minimal space needed for F' is lower-bounded by the number of variables that have to be mentioned simultaneously in any proof for F . As immediate corollaries of this, we get simpler proofs for previously known optimal lower bounds for space and for the recent space-length separation in [Ben-Sasson and Nordström 2008].

Moreover, applying our theorem to so-called pebbling formulas defined in terms of pebble games on directed acyclic graphs, and then studying black-white pebblings on these graphs, we obtain a host of space-length trade-off results for space in the range from constant to $O(n/\log \log n)$, most of them superpolynomial or even exponential.

*Research supported by the Israeli Science Foundation and by the US-Israel Binational Science Foundation.

†Research supported in part by the Ericsson Research Foundation, the Foundation Olle Engkvist Byggmästare, and the Foundation Blanceflor Boncompagni-Ludovisi, née Bildt.

‡Part of this work performed while at the Royal Institute of Technology (KTH) and while visiting the Technion.

1 Introduction

1.1 Previous Work

Resolution length and space The *resolution* proof system, introduced by Blake [Bla37] in 1937 is the single-most studied proof system in propositional proof complexity. The interest in resolution is due to its lying at the very base of the important bounded-depth Frege hierarchy of propositional proof systems and because the *proof* complexity of resolution is tightly connected to the *computational* complexity of the prominent family of *SAT solvers* based on the DPLL algorithm of [DLL62, DP60, Rob65].

The interest in resolution has led to an extensive study of the complexity of refutations in this system. The most important proof complexity measure is the *length* of refutations and the most important question regarding this measure has been (and still is) to establish techniques for proving lower bounds on length. Over the past half century, starting with the seminal superpolynomial lower bound for *regular* resolution by Tseitin in 1968 [Tse68], several techniques for proving superpolynomial lower bounds on this complexity measure have been discovered. Notable examples include [Hak85, Urq87, CS88, Pud97, BKPS02, BSW01, Raz03, Raz04]. We refer to the surveys [Tor04, Seg07] for more information on this topic.

The study of resolution *space* complexity was initiated more recently—about ten years ago—by Esteban and Torán [ET01, Tor99]. Intuitively, the space of a refutation is the maximal amount of memory needed while verifying it, and the space of refuting the CNF formula F is defined as the minimal space of any resolution refutation of F . Over the past decade, a number of upper and lower bounds for refutation space in resolution have been presented in, for example, [ABSRW02, BSG03, EGM04, ET03].

There are two main ways to measure the amount of memory needed to verify a refutation and these measures are known as *clause space* and *variable space*. The former measure is defined as the number of different clauses in the memory, regardless of the amount of memory each clause requires. The latter is the number of literals kept in memory, i.e., it is the sum of the sizes of the clauses kept in memory. While variable space is more clearly related to the actual amount of memory required to verify a proof—the actual memory is at most $\log n$ times the variable space—clause space has attracted most of the attention. The reason for this seems to be that clause space has interesting connections to refutation length and *width*, which is the size of a largest clause in the refutation. Esteban and Torán [ET01] proved that clause space is at most logarithmic in the minimal length of a tree-like refutation of a formula, which implies that clause space is bounded by the number of variables appearing in the formula, and Atserias and Dalmau [AD03] proved that space is lower bounded by width.

The question of the relation between clause space and length of general resolution proofs was raised by the first author in [BS02] and has also been discussed in, for instance, [ET03, Seg07, Tor04]. A pair of works of the second author and Håstad [Nor06, NH08b] have shown that, in contrast to the case of tree-like resolution, length and clause space of general resolution proofs are not strongly related. By this we mean that the existence of a short proof does not necessarily imply the existence of a proof that can be carried out in small clause space. In our recent joint work [BSN08] we showed that the separation of clause space and length can be “maximally” large. More precisely, the main result in our paper is an explicit construction of k -CNF formulas of size n (for arbitrarily large n) that have refutations of size $O(n)$ but require clause space $\Omega(n/\log n)$. We say this separation is “maximal” because these bounds are tight up to constant factors.

Length-space trade-offs The focus of this paper is the fundamental question of the *trade-off* between length and space in resolution. Informally, this question asks how much time one can save when verifying a refutation by allowing more working memory during the verification process. Notice that the above-mentioned lower bounds on length and on space do not deal with this question, but rather state absolute lower bounds on each individual complexity measure. Consider for instance the maximal separation of length and space described in the previous paragraph. This separation is maximal since by combining

results from [ET01, HPV77] we know that any formula refutable in time $O(n)$ can also be refuted in space $O(n/\log n)$. But can this linear-length refutation be carried out in space, say, $100 \cdot n/\log n$? As we show later in this paper, the answer is no in general. Sometimes short refutations require large space, and small space implies long proofs. Analogous time-space trade-offs are well-known in computational complexity (see, e.g., [CS80, CS82, LT82]) and one of the main results of this work is to show how such classical time-space results can be “lifted” to give length-space trade-offs for resolution.

The question of length-space trade-offs in resolution was first studied by the first author in [BS02] and more recently by Hertel and Pitassi in [HP07] and by the second author in [Nor07]. These works have a number of limitations that are overcome in the current paper. The results of [BS02] are limited to the very restricted case of tree-like resolution. The paper [HP07] deals with variable space only and in addition require formulas with rapidly growing width, and [Nor07] uses a somewhat artificial construction of formulas “glued together” from two different unsatisfiable subformulas over disjoint variable sets. Moreover, both trade-off results for general resolution apply only for a very carefully selected ratio of space-to-formula-size and display a sharp and abrupt decay of proof length when space is increased even by small amounts. For instance, the refutation length of the formulas of [HP07] drops exponentially once the variable space is increased to 3 literals above the bare minimal variable space required.

1.2 Our contribution

This paper contains two main results regarding resolution length and space, and one auxiliary result about “classical” time-space trade-offs. Our first result is a new method to obtain clause space lower bounds from lower bounds on a space measure related to variable space. The second result, which builds upon the first, is a technique to convert time-space trade-offs from the “classical” computational setting to resolution.

The Substitution Space Theorem To describe our first result we define the *variable support size* of a refutation as the maximal number of distinct variables appearing simultaneously in memory during the refutation. Thus, in particular, variable support size is a lower bound on variable space. We present a general method to transform lower bounds on the variable support size for F to clause space lower bounds on a formula F' obtained from F as follows. Suppose F mentions variables x_1, \dots, x_n . To produce F' all we do is substitute each variable x_i with the exclusive-or¹ (xor) of two copies of x_i , denoted $x_i^{(1)}, x_i^{(2)}$ and expand the resulting “clauses” (which became disjunctions of xors after substitution) to obtain a CNF formula in the standard way. Our first main theorem can now be stated (informally) as follows.

Theorem 1.1 (Substitution Space Theorem (Informal)). *For any CNF formula F over the set of variables $\{x_1, \dots, x_n\}$, let F' denote the formula with the exclusive-or $x_i^{(1)} \oplus x_i^{(2)}$ substituted for x_i , written in CNF in the canonical way.*

Then any refutation π of F in bounded width can be transformed into a refutation π' of F' such that the length and variable space of π' is at most a constant times the length and variable space of π , respectively.

In the other direction, any refutation π' of the substitution formula F' can be translated back into a refutation π of F such that the length of π is upper-bounded by the length of π' and the variable support size of π is at most the clause space of π' .

The most surprising aspect of this theorem, which is also the hardest to prove, is that one can convert support size lower bounds for F to clause space lower bounds for F' . This reduces the problem of proving lower bounds on clause space to the easier task of proving lower bounds on variable support size.

¹There is nothing magical about the exclusive-or of two variables. Substituting each variable with any function whose value is never dictated by only one variable will lead to essentially the same Substitution Space Theorem.

The proof of the Substitution Space Theorem is outlined in Section 4. We believe it is of independent interest; to wit, in subsequent work (not included in this report) we generalize it to understand the connection between length and space of the stronger proof system known as k -DNF resolution. Let us briefly describe the main ideas in the proof of the clause space–variable support size connection. A refutation π' of F' is a sequence of *clause configurations* where the t th configuration is a set of clauses over variables $x_1^{(1)}, x_1^{(2)}, \dots, x_n^{(1)}, x_n^{(2)}$ corresponding to the content of the memory at time t in the proof. We start by “projecting” each memory configuration down on a set of clauses over the original variables x_1, \dots, x_n . Next, we argue that the sequence of projected sets is (almost) a resolution refutation of F , which we call π . Finally, we show that the variable support of each projected set in π is a lower bound on the clause space of its projecting clause configuration in π' .

The Substitution Space Theorem is inspired by our recent work [BSN08] and indeed our main theorem there is a special case of this new theorem. Let us highlight the important novel aspects of this more general theorem. First and foremost, our previous statement applied only to a very special kind of formulas known as *pebbling contradictions* whereas the Substitution Space Theorem can be applied convert *any* CNF formula requiring large variable support size into a new and closely related CNF formula requiring large clause space. Second, the proof of the Substitution Space Theorem is much cleaner and simpler than the previous one. There is no longer any need to assume the existence of any “underlying directed acyclic graphs” and construct intricate intermediate resolution-like pebble games on these DAGs. Third, the Substitution Space Theorem gives length-preserving reductions from π to π' and vice versa, whereas it was unclear how to derive similar reductions from our previous work. And length-preserving reductions are crucial for our length-space trade-offs described below.

We end the discussion of the Substitution Space Theorem by pointing out that the space bounds obtained from the Substitution Space Theorem apply to both clause and variable space. This is because the lower bound on space of π' is in terms of clause space. Thus, it implies a similar lower bound on the variable space of π' because variable space is always at least as large as clause space. In the other direction, the upper bound on the space of π' is in terms of the larger of the two space measures, variable space, and hence applies also to clause space. The “tightness of bounds” of the Substitution Space Theorem plays a pivotal role in our second main result, namely, the length-space trade-offs described next.

Trade-offs in resolution Our second main result is a new method to “lift” classical time-space trade-off results to the proof complexity world and obtain a host of “robust” length-space trade-offs for resolution. By “robust” we mean that the trade-off is not significantly affected by small changes to either space or time and displays a rather slow and gradual decrease in one parameter (say, length) as the other (say, space) is increased. Prior to this work such “robust” trade-offs were known only for tree-like resolution [BS02].

All trade-off results reported here follow the same proof strategy, which is described in loose terms next (the full details appear in Appendix C). We start with a computational time-space trade-off which is typically stated as a result about the *pebbling price* of a directed acyclic graph. The use of pebbling in the context of space lower bounds is by now standard and we refer the reader to [Pip80] for a survey of pebbling results and to [Nor08] for a discussion of pebbling and resolution. (Relevant formal definitions appear in Appendix A). The pebbling trade-off results we need are of the following nature.

“There exists (arbitrarily large) directed acyclic graphs G over n vertices and bounded indegree that (i) can be pebbled with p pebbles in time t , but (ii) any pebbling strategy of G using $s < p$ pebbles requires time $f(s)$, where f monotonically decreases in s .”

One should think of t as linear in n and of $f(s)$ as being much larger than t for small values of s (We will discuss later how “large” can $f(s)$ get to be.)

With such a pebbling trade-off in hand, we construct from G a CNF formula F , known as a *pebbling contradiction* (see Definition A.8) and promptly substitute each variable by (say) the exclusive-or of two

copies of the variable, as described above. Our hope is that the resulting formula, denoted F' , will display a length-space trade-off similar in spirit to the pebbling trade-off of the underlying graph. More to the point, the upper bound of t on the time required to pebble G using p pebbles should imply that F' can be refuted in length $\approx t$ and variable space $\approx p$ (consequently, the upper bound on clause space is also $\approx p$). And the Substitution Space Theorem says that a refutation π' of F' in time t' and clause space s implies a refutation π of F in time $\approx t'$ and variable space $\approx s$. Finally, by a close reading of the construction in [BS02], we deduce that any refutation of length t' and variable support size s yields a pebbling strategy for G of time t' and space s , which implies $t' > f(s)$.

Unfortunately, things are not that simple. We know how to convert a pebbling strategy into a short and space-efficient refutation only if the pebbling strategy is a so-called *black pebbling* (which corresponds to deterministic space). On the other hand, the result of [BS02] converts the proof π into a *black-white pebbling* strategy (which corresponds to nondeterministic space). To complicate matters further, it is known that black white pebbling can be asymptotically more efficient than black pebbling [KS88, Wil85].

Thus, to obtain our trade-off results we need a strong form of “dual” pebbling trade-offs, where the upper bound (i) is stated in terms of *black* pebbling while the matching lower bound (ii) applies to the stronger model of *black-white* pebbling. Appealing to the Substitution Space Theorem, we can show that any such strong pebbling trade-off translates into a length-space trade-off for resolution.

Using this method of proof we present a number of robust size-space trade-offs for resolution. Before giving a few examples we explain why the need arises for different trade-offs (as opposed to just one global statement). In a nutshell, this is a mirror-picture of the state of size-space trade-offs for pebbling graphs upon which we rely. For instance, suppose G can be pebbled in constant space. Then counting arguments show that G can be pebbled in polynomial time and constant space simultaneously. Thus, if we want to present a nontrivial size-space trade-off for a formula that can be refuted in constant space we cannot hope to get this trade-off to be superpolynomial. Similarly, if G can be pebbled in, say, polylogarithmic space, we cannot obtain exponential time-space trade-offs. We are interested in deriving robust trade-offs for a large range of space complexity parameters and thus we must rely on diverse size-space trade-off results which each come from a different family of graphs. We end this section by describing a couple of trade-off results (many more appear in Appendix E).

Our strong pebbling trade-offs come from three sources. First, we prove a new strong trade-off result for a family of graphs introduced by Carlson and Savage in [CS80, CS82]. Carlson and Savage prove time-space trade-offs for these graphs in the black pebbling model, but to get a strong dual trade-off we need to modify their construction a bit and above all apply different ideas to prove lower bounds in the more challenging black-white pebbling setting. (Details appear in Appendix D.2.) One of the results derived from this is the rather striking statement that superpolynomial length-space trade-offs can occur for *arbitrarily slowly* growing non-constant space. (The formal statement appears as Theorem E.2.)

Theorem 1.2 (Superpolynomial trade-offs for super-constant space (Informal)). *For any arbitrarily slowly growing function $\omega(1) = s(n) = O(n^{1/7})$ and any $\epsilon > 0$ there exists a family of k -CNF formulas $\{F_n\}_{n=1}^\infty$ of size $O(n)$ refutable in length $O(n)$, refutable in space $s(n)$, but not simultaneously so. On the contrary, these formulas are refutable in length $O(n)$ and variable space $O\left((n/s^2(n))^{1/3}\right)$ simultaneously, but any refutation of F_n in clause space $O\left((n/s^2(n))^{1/3-\epsilon}\right)$ must have superpolynomial length.*

Three remarks should be made. First, notice that the trade-off applies to both clause and variable space. This is because the upper bounds are stated in terms of the larger of these two measures (variable space) while the lower bounds are in terms of the smaller one (clause space). This optimality of bound-type is inherited from the Substitution Space Theorem. Second, observe the “robust” nature of the trade-off, which is displayed by the long range of space complexity (from $\omega(1)$ up to $\approx n^{1/3}$) which requires superpolynomial

length. Finally, we remark that the lower bound on length reaches up till very close to where our upper bound kicks in.

A second source of trade-off results for resolution comes from studying the graphs appearing in the study of “classical” time-space trade-offs but deriving *strictly better upper bounds on their refutation complexity than what can provably be obtained for black pebbling*. To do this, we cannot use the machinery developed in this paper as a black box, but need to prove upper bounds in resolution directly. Our quadratic length-space trade-off for constant space in Theorem E.1, the statement of which is omitted here due to space constraints, is of this type.

Our third and final source of trade-off results comes from the seminal work of Lengauer and Tarjan [LT82], in which they showed strong pebbling trade-offs for variety of graphs. For instance, we can obtain the following very strong trade-off in this way.

Theorem 1.3 (Exponential trade-offs for nearly-linear space (Informal)). *There exists constants $K < K'$ and $\epsilon > 0$ and a family of k -CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $O(n)$ that are refutable in length $O(n)$, in space $K \cdot n / \log n$, but not simultaneously so. On the contrary, any refutation π of F_n in space $\leq K' \cdot n / \log n$ must be of length $\exp(n^\epsilon)$.*

1.3 Organization of the Rest of This Paper

After a few basic definitions in Section 2 (more definitions appear in Appendix A) we state our first main result, the Substitution Space Theorem, in Section 3, along with two immediate corollaries that follow from it. We sketch the proof of the Substitution Space Theorem in Section 4 and complete it (due to space limitations) in Appendix B. Our second main result, namely, the method for converting strong pebbling trade-offs into length-space trade-offs for resolution, is described in Appendix C. In Appendix D, we derive our new pebbling trade-off and survey some previously known ones. These results are needed for the robust length-space trade-offs that conclude our paper in Appendix E.

2 Preliminaries

In this section we present a few definitions regarding resolution that are crucial to what follows. Due to space limitations we moved some essential but commonly known definitions regarding resolution and pebbling to Appendix A.

2.1 The Resolution Proof System

When we want to study length and space simultaneously in resolution, we have to be slightly careful with the definitions so that we will be able to capture length-space trade-offs. Just listing the clauses used in a resolution refutation does not tell us *how* the refutation was performed, and essentially the same refutation can be carried out in vastly different time depending on the space constraints (as is shown in this paper). Following the exposition in [ET01], a resolution refutation can be seen as a Turing machine computation, with a special read-only input tape from which the axioms can be downloaded and a working memory where all derivation steps are made. Then the length of a proof is essentially the time of the computation and space measures memory consumption. The formal definitions follow.

Definition 2.1 (Resolution ([ABSRW02])). A *clause configuration* \mathbb{C} is a set of clauses. A sequence of clause configurations $\{\mathbb{C}_0, \dots, \mathbb{C}_\tau\}$ is a *resolution derivation* from a CNF formula F if $\mathbb{C}_0 = \emptyset$ and for all $t \in [\tau]$, \mathbb{C}_t is obtained from \mathbb{C}_{t-1} by one of the following rules:

Axiom Download $\mathbb{C}_t = \mathbb{C}_{t-1} \cup \{C\}$ for some $C \in F$ (an *axiom*).

Erasure $\mathbb{C}_t = \mathbb{C}_{t-1} \setminus \{C\}$ for some $C \in \mathbb{C}_{t-1}$.

Inference $\mathbb{C}_t = \mathbb{C}_{t-1} \cup \{D\}$ for some D inferred by resolution from $C_1, C_2 \in \mathbb{C}_{t-1}$.

A resolution derivation $\pi : F \vdash A$ of a clause A from a formula F is a derivation $\{\mathbb{C}_0, \dots, \mathbb{C}_\tau\}$ such that $\mathbb{C}_\tau = \{A\}$. A *resolution refutation* of F is a derivation of the empty clause 0, i.e., the clause with no literals, from F .

We will be interested in studying length and space in resolution, which are formalized as proof complexity measures in the next definition. Also, it will be convenient to define what width in resolution is.

Definition 2.2 (Length, width and space). The *length* $L(\pi)$ of a resolution derivation π is the total number of axiom downloads and inferences made in π , i.e., the total number of clauses counted with repetitions.

The *width* $W(C)$ of a clause C is the number of literals in it, the width $W(F)$ of a formula F is the size of a widest clause in F , and the width $W(\pi)$ of a derivation π is defined in the same way.

The *clause space* $Sp(\mathbb{C})$ of a clause configuration \mathbb{C} is $|\mathbb{C}|$, i.e., the number of clauses in \mathbb{C} , and the *variable space* $VarSp(\mathbb{C})$ is $\sum_{C \in \mathbb{C}} |C|$, i.e., the total number of literals in \mathbb{C} counted with repetitions.²

Taking the minimum over all refutations of a formula F , we define $L(F \vdash 0) = \min_{\pi: F \vdash 0} \{L(\pi)\}$ as the length of refuting F , $W(F \vdash 0) = \min_{\pi: F \vdash 0} \{W(\pi)\}$ as the width of refuting F , and $Sp(F \vdash 0) = \min_{\pi: F \vdash 0} \{Sp(\pi)\}$ and $VarSp(F \vdash 0) = \min_{\pi: F \vdash 0} \{VarSp(\pi)\}$ as the clause space and variable space, respectively, of refuting F in resolution.

Note that this definition of length exactly captures the minimum length as the number of lines in a listing of the refutation (just construct a refutation that only does downloads and inferences until it gets to 0, and only then erase all the other clauses). For tree-like resolution, we obtain the standard length measure by insisting that every clause be used at most once before being erased. In general, Definition 2.2 unifies previous definitions for various subsystems of resolution and gives us the possibility to measure length and space simultaneously in a meaningful way.

We also need to define a new measure which is related to, but weaker than, variable space.

Definition 2.3 (Variable support size). Let us say that the *variable support size*, or just *support size*, of a clause set \mathbb{C} is $SuppSize(\mathbb{C}) = |Vars(\mathbb{C})|$, i.e., the number of variables mentioned in \mathbb{C} . We define the support size of a resolution derivation $\pi = \{\mathbb{C}_0, \dots, \mathbb{C}_\tau\}$ to be $SuppSize(\pi) = \max_{t \in [\tau]} \{SuppSize(\mathbb{C}_t)\}$ and the minimal support size of refuting F is then $SuppSize(F \vdash 0) = \min_{\pi: F \vdash 0} \{SuppSize(\pi)\}$.

The difference between variable space and variable support size is that the variables space counts the number of variable occurrences in \mathbb{C} *with repetitions*, but for variable support size we only count each variable once no matter how often it occurs. It follows that the support size of refuting a formula is always at most linear in the formula size, while the refutation variable space could potentially be quadratic in the formula size in the worst case. (It should be noted, though, that no such formulas are known to exist, and to the best of our knowledge it is even an open problem to prove superlinear lower bounds on variable space.)

2.2 Substitution Formulas

Throughout this paper, we will let f_d denote any (non-constant) Boolean function $f_d : \{0, 1\}^d \mapsto \{0, 1\}$ of arity d . We use the shorthand $\vec{x} = (x_1, \dots, x_d)$, so that $f_d(\vec{x})$ is just an equivalent way of writing $f_d(x_1, \dots, x_d)$. Every function $f_d(x_1, \dots, x_d)$ is equivalent to a CNF formula over x_1, \dots, x_d with at

²Note that if one wanted to be really precise, space (as well as formula size) should probably measure the number of *bits* rather than the number of literals. However, counting literals makes matters substantially cleaner, and the difference is at most a logarithmic factor. Therefore, counting literals seems to be the established way of measuring formula size and variable space.

most 2^d clauses. Fix a canonical way to represent functions as CNF formulas and let $Cl[f_d(\vec{x})]$ denote the canonical set of clauses representing f_d . Similarly, let $Cl[\neg f_d(\vec{x})]$ denote the clauses in the canonical representation of the negation of f . The following definition extends the notion of substitution to a CNF formula F . For notational convenience, we assume that F only has variables x, y, z , et cetera, without subscripts, so that $x_1, \dots, x_d, y_1, \dots, y_d, z_1, \dots, z_d, \dots$ are new variables not occurring in F .

Definition 2.4 (Substitution formula). For a positive literal x and a non-constant Boolean function f_d , we define the f_d -substitution of x to be $x[f_d] = Cl[f_d(\vec{x})]$, i.e., the canonical representation of $f_d(x_1, \dots, x_d)$ as a CNF formula. For a negative literal \bar{y} , the f_d -substitution is $\bar{y}[f_d] = Cl[\neg f_d(\vec{y})]$. The f_d -substitution of a clause $C = a_1 \vee \dots \vee a_k$ is the CNF formula

$$C[f_d] = \bigwedge_{C_1 \in a_1[f_d]} \dots \bigwedge_{C_k \in a_k[f_d]} (C_1 \vee \dots \vee C_k) \quad (1)$$

and the f_d -substitution of a CNF formula F is $F[f_d] = \bigwedge_{C \in F} C[f_d]$.

In Appendix A we list several useful properties of substituted formulas that appear later on in our proofs.

3 The Substitution Space Theorem

Let us now state formally our main technical result. It is phrased in terms of substitution using functions f having the property that no single variable x_i determines the value of $f(x_1, \dots, x_d)$.

Definition 3.1 (Non-authoritarian function). We will call a Boolean function f over d variables x_1, \dots, x_d *non-authoritarian* if for any variable x_i and any truth value $\alpha(x_i) = \nu_i$ assigned to x_i , α can be extended to a truth value assignment α' satisfying $f(x_1, \dots, x_d)$ and another truth value assignment α'' falsifying $f(x_1, \dots, x_d)$. If f does not satisfy this requirement, then we will call the function *authoritarian*.

Examples of non-authoritarian functions include exclusive-or and threshold functions over d variables for which the threshold lies above 1 and below d .

Loosely put, the Substitution Space Theorem says that if a formula F can be refuted in resolution in small length and width simultaneously, then so can the substitution formula $F[f_d]$. There is an analogous result in the other direction as well in the sense that we can translate any refutation π_f of $F[f_d]$ into a refutation π of the original formula F where the length of π is almost upper-bounded by the length of π_f (this will be made precise below). So far this is nothing very unexpected, but what is more interesting is that if f_d is non-authoritarian, then the clause space of π_f is an upper bound on the number of variables mentioned simultaneously in π . Thus, the theorem says that we can convert lower bounds on variable support size into lower bounds on clause space by making substitutions using non-authoritarian functions.

Theorem 3.2 (Substitution Space Theorem). *Let F be any unsatisfiable CNF formula and f_d be any non-constant Boolean function of arity d . Then it holds that the substitution formula $F[f_d]$ can be refuted in width*

$$W(F[f_d] \vdash 0) = O(d \cdot W(F \vdash 0))$$

and length

$$L(F[f_d] \vdash 0) \leq \min_{\pi: F \vdash 0} \{L(\pi) \cdot \exp(O(d \cdot W(\pi)))\} .$$

In the other direction, any refutation $\pi_f : F[f_d] \vdash 0$ of the substitution formula can be transformed into a refutation $\pi : F \vdash 0$ of the original formula such that the number of axiom downloads in π is at most the number of axiom downloads in π_f . If in addition f_d is non-authoritarian, it holds that $Sp(\pi_f) > SuppSize(\pi)$, i.e., the clause space of refuting the substitution formula $F[f_d]$ is lower-bounded by the variable support size of refuting the original formula F .

Note that if F is refutable simultaneously in linear length and constant width, then the bound in Theorem 3.2 on $L(F[f_d] \vdash 0)$ becomes linear in $L(F \vdash 0)$. It would be interesting to know if the bound in terms of number of axiom downloads could in fact be strengthened to a bound in terms of length, but we do not know if this is the case or not. Luckily enough, however, the bound in terms of axiom downloads turns out to be exactly what we need for our applications.

Although this might not be immediately obvious, Theorem 3.2 is remarkably powerful as a tool for understanding space in resolution. It will take some more work before we can present our main applications of this theorem, which are the strong time-space trade-off results discussed in Appendix E. Let us note, however, for starters, that without any extra work we immediately get lower bounds on space.

Esteban and Torán [ET01] proved that the clause space of refuting F is upper-bounded by the formula size. In the papers [ABSRW02, BSG03, ET01] it was shown, using quite elaborate arguments, that there are polynomial-size k -CNF formulas with lower bounds on clause space matching this upper bound up to constant factors. Using Theorem 3.2 we can get a different proof of this fact.

Corollary 3.3 ([ABSRW02, BSG03, ET01]). *There are k -CNF formula families $\{F\}_{n=1}^{\infty}$ with $\Theta(n)$ clauses over $\Theta(n)$ variables such that $Sp(F_n \vdash 0) = \Theta(n)$.*

Proof. Just pick any formula family for which it is shown that any refutation of F_n must at some point in the refutation mention $\Omega(n)$ variables at the same time (e.g., from [BSW01]), and then apply Theorem 3.2. \square

It should be noted, though, that when we apply Theorem 3.2 the formulas in [ABSRW02, BSG03, ET01] are changed. We remark that there is another, and even more elegant way to derive Corollary 3.3 from [BSW01] without changing the formulas, namely by using the lower bound on clause space in terms of width in [AD03].

For our next corollary, however, there is no other, simpler way known to prove the same result. Instead, our proof in this paper actually improves the constants in the result.

Corollary 3.4 ([BSN08]). *There are families $\{F_n\}_{n=1}^{\infty}$ of k -CNF formulas of size $O(n)$ refutable in linear length $L(F_n \vdash 0) = O(n)$ and constant width $W(F_n \vdash 0) = O(1)$ such that the minimum clause space required is $Sp(F_n \vdash 0) = \Omega(n/\log n)$.*

Proof. In [BS02], the first author showed that there are formulas refutable simultaneously in linear length and constant width, but for which any refutation must at some point mention $\Omega(n/\log n)$ distinct variables at the same time (although the result was stated in slightly different terms). Corollary 3.4 follows immediately from this by applying Theorem 3.2. \square

In fact, the ideas in [BS02], which provide a way of translating back and forth between resolution and pebbling, are also what allows us to prove strong trade-off results for resolution. We will return to this in Appendix C where we formalize this resolution-pebbling correspondence.

4 Outline of Proof of the Substitution Space Theorem

We divide the proof of Theorem 3.2 into three parts in Theorems 4.1, 4.4, and 4.5 below.

Theorem 4.1. *For any CNF formula F and any non-constant Boolean function f_d , it holds that*

$$W(F[f_d] \vdash 0) = O(d \cdot W(F \vdash 0))$$

and

$$L(F[f_d] \vdash 0) \leq \min_{\pi: F \vdash 0} \{L(\pi) \cdot \exp(O(d \cdot W(\pi)))\} .$$

These upper bounds on refutation width and length for $F[f_d]$ are not hard to show. The proof proceeds along the following lines. Given a resolution refutation π of F , we construct a refutation $\pi_f : F[f_d] \vdash 0$ mimicking the derivation steps in π . When π downloads an axiom C , we download the $\exp(O(d \cdot W(C)))$ axiom clauses in $C[f_d]$. When π resolves $C_1 \vee x$ and $C_2 \vee \bar{x}$ to derive $C_1 \vee C_2$, we use the fact that resolution is implicationally complete to derive $(C_1 \vee C_2)[f_d]$ from $(C_1 \vee x)[f_d]$ and $(C_2 \vee \bar{x})[f_d]$ in at most $\exp(O(d \cdot W(C_1 \vee C_2)))$ steps. We return to the details of the proof in Section B.1.

It is more challenging, however, to prove that we can get lower bounds on clause space for $F[f_d]$ from lower bounds on support size for F . The idea is to look at refutations of $F[f_d]$ and “project” them down on refutations of F . To do this, we first define a special kind of “precise implication.”

Definition 4.2 (Precise implication). Let F be a CNF formula and f_d a non-constant Boolean function, and suppose that \mathbb{D} is a set of clauses derived from $F[f_d]$ and that P and N are (disjoint) subset of variables of F . If

$$\mathbb{D} \models \bigvee_{x \in P} f_d(\vec{x}) \vee \bigvee_{y \in N} \neg f_d(\vec{y}) \quad (2a)$$

but for all strict subsets $\mathbb{D}' \subsetneq \mathbb{D}$, $P' \subsetneq P$, and $N' \subsetneq N$ it holds that

$$\mathbb{D}' \not\models \bigvee_{x \in P} f_d(\vec{x}) \vee \bigvee_{y \in N} \neg f_d(\vec{y}) \quad , \quad (2b)$$

$$\mathbb{D} \not\models \bigvee_{x \in P'} f_d(\vec{x}) \vee \bigvee_{y \in N} \neg f_d(\vec{y}) \quad , \quad \text{and} \quad (2c)$$

$$\mathbb{D} \not\models \bigvee_{x \in P} f_d(\vec{x}) \vee \bigvee_{y \in N'} \neg f_d(\vec{y}) \quad , \quad (2d)$$

we say that the clause set \mathbb{D} implies $\bigvee_{x \in P} f_d(\vec{x}) \vee \bigvee_{y \in N} \neg f_d(\vec{y})$ *precisely* and write

$$\mathbb{D} \triangleright \bigvee_{x \in P} f_d(\vec{x}) \vee \bigvee_{y \in N} \neg f_d(\vec{y}) \quad . \quad (3)$$

Note that $P = N = \emptyset$ in Definition 4.2 corresponds to \mathbb{D} being unsatisfiable.

Let us also use the convention that any clause C can be written $C = C^+ \vee C^-$, where $C^+ = \bigvee_{x \in \text{Lit}(C)} x$ is the disjunction of the positive literals in C and $C^- = \bigvee_{\bar{y} \in \text{Lit}(C)} \bar{y}$ is the disjunction of the negative literals.

Definition 4.3 (Projected clauses). Let F be a CNF formula and f_d a non-constant Boolean function, and suppose that \mathbb{D} is a set of clauses derived from $F[f_d]$. Then we say that \mathbb{D} *projects* the clause $C = C^+ \vee C^-$ on F —or, perhaps more correctly, on $\text{Vars}(F)$ —if there is a subset $\mathbb{D}_C \subseteq \mathbb{D}$ such that

$$\mathbb{D}_C \triangleright \bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C^-} \neg f_d(\vec{y}) \quad (4)$$

and we write $\text{proj}_F(\mathbb{D}) = \{C \mid \exists \mathbb{D}_C \subseteq \mathbb{D} \text{ s.t. } \mathbb{D}_C \triangleright \bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C^-} \neg f_d(\vec{y})\}$ to denote the set of all clauses that \mathbb{D} projects on F .

Given that we now know how to translate clauses derived from $F[f_d]$ into clauses over $\text{Vars}(F)$, the next step is to show that this translation preserves resolution refutations.

Theorem 4.4. *Suppose that $\pi_f = \{\mathbb{D}_0, \dots, \mathbb{D}_\tau\}$ is a resolution refutation of $F[f_d]$ for some arbitrary unsatisfiable CNF formula F and some arbitrary non-constant function f_d . Then the sets of projected clauses $\{\text{proj}_F(\mathbb{D}_0), \dots, \text{proj}_F(\mathbb{D}_\tau)\}$ form the “backbone” of a resolution refutation π of F in the sense that:*

- $proj_F(\mathbb{D}_0) = \emptyset$.
- $proj_F(\mathbb{D}_\tau) = \{0\}$.
- All transitions from $proj_F(\mathbb{D}_{t-1})$ to $proj_F(\mathbb{D}_t)$ for $t \in [\tau]$ can be accomplished by axiom downloads from F , inferences, erasures, and possibly weakening steps in such a way that the variable support size in π during these intermediate derivation steps never exceeds $\max_{\mathbb{D} \in \pi_f} \{SuppSize(proj_F(\mathbb{D}))\}$.
- The only time π performs a download of some axiom C in F is when π_f downloads some axiom $D \in C[f_d]$ in $F[f_d]$.

Using standard techniques we can get rid of the weakening moves in a postprocessing step, but allowing them in the statement of Theorem 4.4 makes the proof much cleaner. Accepting Theorem 4.4 on faith for the moment, the final missing link in the proof of the Substitution Space Theorem is the following lower bound.

Theorem 4.5. *Suppose that \mathbb{D} is a set of clauses derived from $F[f_d]$ for some arbitrary unsatisfiable CNF formula F and some non-authoritarian function f_d . Then $Sp(\mathbb{D}) = |\mathbb{D}| > SuppSize(proj_F(\mathbb{D}))$.*

Combining Theorems 4.1, 4.4, and 4.5 (proven in Appendix B), the Substitution Space Theorem follows.

5 Directions for Further Research

We end by briefly mentioning a few open questions related to our reported work that we find most interesting.

Open Question 1. *Are there polynomial-size k -CNF formulas which require variable refutation space $VarSp(F \vdash 0) = \Omega((\text{size of } F)^2)$?*

The answer has been conjectured by [ABSRW02] to be “yes”, but as far as we are aware, there are no stronger lower bounds on variable space known than those that follow trivially from corresponding linear lower bounds on clause space. Thus, a first step would be to show superlinear lower bounds on variable space.

Two other questions concern the possible trade-offs at the extremal points of the space interval, where we can only get polynomial trade-offs for constant space and no trade-offs at all for linear space.

Open Question 2. *Are there superpolynomial trade-offs for formulas refutable in constant space?*

Open Question 3. *Are there formulas with trade-offs in the range space $>$ formula size? Or can every refutation be carried out in at most linear space?*

We find the Open Question 3 especially intriguing. Note that all bounds on clause space proven so far, including the trade-offs in the current paper, are in the regime where the space is less than formula size (which is quite natural, since by [ET01] we know the size of the formula is an upper bound on the minimal clause space needed). It is unclear to what extent such lower bounds on space are relevant to state-of-the-art SAT solvers, however, since such algorithms will presumably use at least a linear amount of memory to store the formula to begin with. For this reason, it seems to be a highly interesting problem to determine what can be said if we allow extra clause space above linear. Are there formulas exhibiting trade-offs in this superlinear regime, or is it always possible to carry out a minimal-length refutation in, say, at most a constant factor times the linear upper bound on the space required for any formula?

We point out that pebbling formulas cannot help answer these questions as these formulas are always refutable in linear time and linear space simultaneously by construction, and since constant pebbling space implies polynomial pebbling time.

Acknowledgements

The authors want to thank David Carlson, John Gilbert, Thomas Lengauer, Nicholas Pippenger, John Savage, and Robert Tarjan for helpful correspondence regarding their papers on pebbling.

References

- [ABSRW02] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *SIAM Journal on Computing*, 31(4):1184–1211, 2002.
- [AC03] Noga Alon and Michael Capalbo. Smaller explicit superconcentrators. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '03)*, pages 340–346, 2003.
- [AD03] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. In *Proceedings of the 18th IEEE Annual Conference on Computational Complexity (CCC '03)*, pages 239–247, July 2003.
- [AJPU02] Michael Alekhnovich, Jan Johannsen, Toniann Pitassi, and Alasdair Urquhart. An exponential separation between regular and general resolution. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 448–456, May 2002.
- [BEGJ00] Maria Luisa Bonet, Juan Luis Esteban, Nicola Galesi, and Jan Johannsen. On the relative complexity of resolution refinements and cutting planes proof systems. *SIAM Journal on Computing*, 30(5):1462–1484, 2000.
- [BKPS02] Paul Beame, Richard Karp, Toniann Pitassi, and Michael Saks. The efficiency of resolution and Davis-Putnam procedures. *SIAM Journal on Computing*, 31(4):1048–1075, 2002.
- [Bla37] Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.
- [BOP03] Josh Buresh-Oppenheimer and Toniann Pitassi. The complexity of resolution refinements. In *Proceedings of the 18th IEEE Symposium on Logic in Computer Science (LICS '03)*, pages 138–147, June 2003.
- [BS02] Eli Ben-Sasson. Size space tradeoffs for resolution. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 457–464, May 2002.
- [BSG03] Eli Ben-Sasson and Nicola Galesi. Space complexity of random formulae in resolution. *Random Structures and Algorithms*, 23(1):92–109, August 2003.
- [BSIW04] Eli Ben-Sasson, Russell Impagliazzo, and Avi Wigderson. Near optimal separation of treelike and general resolution. *Combinatorica*, 24(4):585–603, September 2004.
- [BSN08] Eli Ben-Sasson and Jakob Nordström. Short proofs may be spacious: An optimal separation of space and length in resolution. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 709–718, October 2008.
- [BSW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001.

- [Coo74] Stephen A. Cook. An observation on time-storage trade off. *Journal of Computer and System Sciences*, 9:308–316, 1974.
- [CS76] Stephen A. Cook and Ravi Sethi. Storage requirements for deterministic polynomial time recognizable languages. *Journal of Computer and System Sciences*, 13(1):25–37, 1976.
- [CS80] David A. Carlson and John E. Savage. Graph pebbling with many free pebbles can be difficult. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC '80)*, pages 326–332, 1980.
- [CS82] David A. Carlson and John E. Savage. Extreme time-space tradeoffs for graphs with small space requirements. *Information Processing Letters*, 14(5):223–227, 1982.
- [CS88] Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988.
- [DLL62] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem proving. *Communications of the ACM*, 5(7):394–397, July 1962.
- [DP60] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.
- [EGM04] Juan Luis Esteban, Nicola Galesi, and Jochen Messner. On the complexity of resolution with bounded conjunctions. *Theoretical Computer Science*, 321(2-3):347–370, August 2004.
- [ET01] Juan Luis Esteban and Jacobo Torán. Space bounds for resolution. *Information and Computation*, 171(1):84–97, 2001.
- [ET03] Juan Luis Esteban and Jacobo Torán. A combinatorial characterization of treelike resolution space. *Information Processing Letters*, 87(6):295–300, 2003.
- [GG81] Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981.
- [GLT80] John R. Gilbert, Thomas Lengauer, and Robert Endre Tarjan. The pebbling problem is complete in polynomial space. *SIAM Journal on Computing*, 9(3):513–524, August 1980.
- [GT78] John R. Gilbert and Robert Endre Tarjan. Variations of a pebble game on graphs. Technical Report STAN-CS-78-661, Stanford University, 1978. Available at <http://infolab.stanford.edu/TR/CS-TR-78-661.html>.
- [Hak85] Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985.
- [HP07] Philipp Hertel and Toniann Pitassi. Exponential time/space speedups for resolution and the PSPACE-completeness of black-white pebbling. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS '07)*, pages 137–149, October 2007.
- [HPV77] John Hopcroft, Wolfgang Paul, and Leslie Valiant. On time versus space. *Journal of the ACM*, 24(2):332–337, April 1977.
- [Kla85] Maria M. Klawe. A tight bound for black and white pebbles on the pyramid. *Journal of the ACM*, 32(1):218–228, January 1985.

REFERENCES

- [KS88] Balasubramanian Kalyanasundaram and George Schnitger. On the power of white pebbles. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC '88)*, pages 258–266, 1988.
- [LT80] Thomas Lengauer and Robert Endre Tarjan. The space complexity of pebble games on trees. *Information Processing Letters*, 10(4/5):184–188, July 1980.
- [LT82] Thomas Lengauer and Robert Endre Tarjan. Asymptotically tight bounds on time-space trade-offs in a pebble game. *Journal of the ACM*, 29(4):1087–1130, October 1982.
- [NH08a] Jakob Nordström and Johan Håstad. Towards an optimal separation of space and length in resolution. Technical Report TR08-026, Electronic Colloquium on Computational Complexity (ECCC), February 2008.
- [NH08b] Jakob Nordström and Johan Håstad. Towards an optimal separation of space and length in resolution (Extended abstract). In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC '08)*, pages 701–710, May 2008.
- [Nor06] Jakob Nordström. Narrow proofs may be spacious: Separating space and width in resolution (Extended abstract). In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*, pages 507–516, May 2006.
- [Nor07] Jakob Nordström. A simplified way of proving trade-off results for resolution. Technical Report TR07-114, Electronic Colloquium on Computational Complexity (ECCC), September 2007.
- [Nor08] Jakob Nordström. *Short Proofs May Be Spacious: Understanding Space in Resolution*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, May 2008. Available at <http://www.csc.kth.se/~jakobn/research/>.
- [Pip77] Nicholas Pippenger. Superconcentrators. *SIAM Journal on Computing*, 6(2):298–304, 1977.
- [Pip80] Nicholas Pippenger. Pebbling. Technical Report RC8258, IBM Watson Research Center, 1980. Appeared in Proceedings of the 5th IBM Symposium on Mathematical Foundations of Computer Science, Japan.
- [PTC77] Wolfgang J. Paul, Robert Endre Tarjan, and James R. Celoni. Space bounds for a game on graphs. *Mathematical Systems Theory*, 10:239–251, 1977.
- [Pud97] Pavel Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *Journal of Symbolic Logic*, 62(3):981–998, September 1997.
- [Raz03] Alexander A. Razborov. Resolution lower bounds for the weak functional pigeonhole principle. *Theoretical Computer Science*, 1(303):233–243, June 2003.
- [Raz04] Ran Raz. Resolution lower bounds for the weak pigeonhole principle. *Journal of the ACM*, 51(2):115–138, 2004.
- [RM99] Ran Raz and Pierre McKenzie. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3):403–435, March 1999.
- [Rob65] John Alan Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, January 1965.

- [Seg07] Nathan Segerlind. The complexity of propositional proofs. *Bulletin of Symbolic Logic*, 13(4):482–537, December 2007.
- [Tor99] Jacobo Torán. Lower bounds for space in resolution. In *Proceedings of the 13th International Workshop on Computer Science Logic (CSL '99)*, volume 1683 of *Lecture Notes in Computer Science*, pages 362–373. Springer, 1999.
- [Tor04] Jacobo Torán. Space and width in propositional resolution. *Bulletin of the European Association for Theoretical Computer Science*, 83:86–104, June 2004.
- [Tse68] Grigori Tseitin. On the complexity of derivation in propositional calculus. In A. O. Silenko, editor, *Structures in Constructive Mathematics and Mathematical Logic, Part II*, pages 115–125. Consultants Bureau, New York-London, 1968.
- [Urq87] Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, January 1987.
- [Val76] Leslie G. Valiant. Graph-theoretic properties in computational complexity. *Journal of Computer and System Sciences*, 13(3):278–285, 1976.
- [Wil85] Robert Wilber. White pebbles help. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing (STOC '85)*, pages 103–112, 1985.

A Additional Preliminary Definitions

A.1 The Resolution Proof System

A *literal* is either a propositional logic variable or its negation, denoted x and \bar{x} , respectively, or sometimes x^1 and x^0 . We define $\overline{\bar{x}} = x$. Two literals a and b are *strictly distinct* if $a \neq b$ and $a \neq \bar{b}$, i.e., if they refer to distinct variables.

A *clause* $C = a_1 \vee \dots \vee a_k$ is a set of literals. Without loss of generality, all clauses C are assumed to be nontrivial in the sense that all literals in C are pairwise strictly distinct (otherwise C is trivially true). We say that C is a *subclause* of D if $C \subseteq D$. A clause containing at most k literals is called a *k-clause*.

A *CNF formula* $F = C_1 \wedge \dots \wedge C_m$ is a set of clauses. A *k-CNF formula* is a CNF formula consisting of k -clauses. We define the *size* $S(F)$ of the formula F to be the total number of literals in F counted with repetitions. More often, we will be interested in the number of clauses $|F|$ of F .

In this paper, when nothing else is stated it is assumed that A, B, C, D denote clauses, \mathbb{C}, \mathbb{D} sets of clauses, x, y propositional variables, a, b, c literals, α, β truth value assignments and ν a truth value 0 or 1. We write

$$\alpha^{x=\nu}(y) = \begin{cases} \alpha(y) & \text{if } y \neq x, \\ \nu & \text{if } y = x, \end{cases} \quad (5)$$

to denote the truth value assignment that agrees with α everywhere except possibly at x , to which it assigns the value ν . We let $Vars(C)$ denote the set of variables and $Lit(C)$ the set of literals in a clause C .³ This notation is extended to sets of clauses by taking unions. Also, we employ the standard notation $[n] = \{1, 2, \dots, n\}$.

In its simplest form, a *resolution derivation* $\pi : F \vdash A$ of a clause A from a CNF formula F can be viewed as a sequence of clauses $\pi = \{D_1, \dots, D_\tau\}$ such that $D_\tau = A$ and each line $D_i, i \in [\tau]$, either is

³Although the notation $Lit(C)$ is slightly redundant given the definition of a clause as a set of literals, we include it for clarity.

one of the clauses in F (an *axiom*) or is derived from clauses D_j, D_k in π with $j, k < i$ by the *resolution rule*

$$\frac{B \vee x \quad C \vee \bar{x}}{B \vee C} . \quad (6)$$

We refer to (6) as *resolution on the variable x* and to $B \vee C$ as the *resolvent* of $B \vee x$ and $C \vee \bar{x}$ on x .

A.2 Some Auxiliary Technical Results for Resolution

We start off by an easy observation.

Observation A.1. *Any unsatisfiable CNF formula F over n variables can be refuted in length at most $2^{n+1} - 1$, clause space at most $O(n)$, and variable space at most $O(n^2)$ simultaneously.*

Proof sketch. Build a search tree where all vertices on level i query the i th variable and where we go to the left, say, if the variable is false under a given truth value assignment α and to the right if the variable is true. As soon as some axiom in F is falsified by the partial assignment defined by the path to a vertex, we make that vertex into a leaf labelled by that clause. This tree has size at most $2^{n+1} - 1$, and if we turn it upside down we can obtain a legal tree-like refutation of F , possibly using weakening. This refutation can be carried out in clause space linear in the tree depth and variable space upper-bounded by the clause space times the number of distinct variables. We refer to, for instance, [BS02, ET01] for more details. \square

For technical reasons, it is sometimes convenient to add a rule for *weakening*, saying that we can always derive a weaker clause $C' \supseteq C$ from C . It is easy to show that any weakening steps can always be eliminated from a resolution refutation without changing anything essential. Let us state this more formally since we will need the precise formulation later on in this paper. The proof is an easy induction over the refutation and we omit the details.

Proposition A.2. *Any resolution refutation $\pi : F \vdash 0$ using the weakening rule can be transformed into a refutation $\pi' : F \vdash 0$ without weakening in at most the same length, width, clause space and variable space, and performing at most the same number of axiom downloads, inferences and erasures as π .*

Another tool that we will use to to simplify some of the proofs is the concept of *restrictions*.

Definition A.3 (Restriction). A *partial assignment* or *restriction* ρ is a partial function $\rho : X \mapsto \{0, 1\}$, where X is a set of Boolean variables. We identify ρ with the set of literals $\{a_1, \dots, a_m\}$ set to true by ρ . The ρ -*restriction* of a clause C is defined to be

$$C \upharpoonright_\rho = \begin{cases} 1 & \text{(i.e., the trivially true clause) if } \text{Lit}(C) \cap \rho \neq \emptyset, \\ C \setminus \{\bar{a} \mid a \in \rho\} & \text{otherwise.} \end{cases}$$

This definition is extended to set of clauses by taking unions.

We write $\rho(-C)$ to denote the minimal restriction fixing C to false, i.e., $\rho(-C) = \{\bar{a} \mid a \in C\}$.

Proposition A.4. *If π is a resolution refutation of F and ρ is a restriction on $\text{Vars}(F)$, then $\pi \upharpoonright_\rho$ can be transformed into a resolution refutation of $F \upharpoonright_\rho$ in at most the same length, width, clause space and variable space as π .*

In fact, $\pi \upharpoonright_\rho$ is a refutation of $F \upharpoonright_\rho$ (removing all trivially true clauses), but possibly using weakening. The proof of this is again an easy induction over the resolution refutation π .

In a resolution refutation of a formula F , there is nothing in Definition 2.1 that rules out that completely unnecessary derivation steps are made on the way, such as axioms being downloaded and them immediately erased again, or entire subderivations being made to no use. In our constructions it will be important that we can rule out some redundancies and enforce the following requirements for any resolution refutation:

- Every clause in memory is used in an inference step before being erased.
- Every clause is erased from memory immediately after having been used for the last time.

We say that a resolution refutation that meets these requirements is *frugal*. The formal definition, which is a mildly modified version of that in [BS02], follows.

Definition A.5 (Frugal refutation). Let $\pi = \{\mathbb{C}_0 = \emptyset, \mathbb{C}_1, \dots, \mathbb{C}_\tau = \{0\}\}$ be a resolution refutation of some CNF formula F . The *essential clauses* in π are defined by backward induction:

- If \mathbb{C}_t is the first configuration containing 0, then 0 is essential at time t .
- If $D \in \mathbb{C}_t$ is essential and is inferred at time t from $C_1, C_2 \in \mathbb{C}_{t-1}$ by resolution, then C_1 and C_2 are essential at time $t - 1$.
- If D is essential at time t and $D \in \mathbb{C}_{t-1}$, then D is essential at time $t - 1$.

Essential clause configurations are defined by forward induction over π . The configuration $\mathbb{C}_t \in \pi$ is essential if all clauses $D \in \mathbb{C}_t$ are essential at time t , if \mathbb{C}_t is obtained by inference from a configuration \mathbb{C}_{t-1} containing only essential clauses at time $t - 1$, or if \mathbb{C}_t is obtained from an essential configuration \mathbb{C}_{t-1} by an erasure step.

Finally, $\pi = \{\mathbb{C}_0, \dots, \mathbb{C}_\tau\}$ is a *frugal refutation* if all configurations $\mathbb{C}_t \in \pi$ are essential.

Without loss of generality, we can always assume that resolution refutations are frugal.

Lemma A.6. Any resolution refutation $\pi : F \vdash 0$ can be converted into a frugal refutation $\pi' : F \vdash 0$ without increasing the length, width, clause space or variable space. Furthermore, the axiom downloads, inferences and erasures performed in π' are a subset of those in π .

Proof. The construction of π' is by backward induction over π . Set $s = \min\{t : 0 \in \mathbb{C}_t\}$ and $\mathbb{C}'_s = \{0\}$. Assume that $\mathbb{C}'_s, \mathbb{C}'_{s-1}, \dots, \mathbb{C}'_{t+2}, \mathbb{C}'_{t+1}$ have been constructed and consider \mathbb{C}_t and the transition $\mathbb{C}_t \rightsquigarrow \mathbb{C}_{t+1}$.

Axiom Download $\mathbb{C}_{t+1} = \mathbb{C}_t \cup \{C\}$: Set $\mathbb{C}'_t = \mathbb{C}'_{t+1} \setminus \{C\}$. (If C is not essential we get $\mathbb{C}'_t = \mathbb{C}'_{t+1}$.)

Erasure $\mathbb{C}_{t+1} = \mathbb{C}_t \setminus \{D\}$: Ignore, i.e., set $\mathbb{C}'_t = \mathbb{C}'_{t+1}$.

Inference $\mathbb{C}_{t+1} = \mathbb{C}_t \cup \{D\}$ inferred from $C_1, C_2 \in \mathbb{C}_t$: If $D \notin \mathbb{C}'_{t+1}$, ignore the step and set $\mathbb{C}'_t = \mathbb{C}'_{t+1}$. Otherwise (using fractional time steps for notational convenience) insert the configurations $\mathbb{C}'_t = \mathbb{C}'_{t+1} \cup \{C_1, C_2\} \setminus \{D\}$, $\mathbb{C}'_{t+\frac{1}{3}} = \mathbb{C}'_{t+1} \cup \{C_1, C_2\}$, $\mathbb{C}'_{t+\frac{2}{3}} = \mathbb{C}'_{t+1} \cup \{C_2\}$.

Finally go through π' and eliminate any consecutive duplicate clause configurations.

It is straightforward to check that π' is a legal resolution refutation. Let us verify that π' is frugal. By backward induction, each \mathbb{C}'_t for integral time steps t contains only essential clauses. By forward induction, if $\mathbb{C}'_{t+1} = \mathbb{C}'_t \cup \{C\}$ is obtained by axiom download, all clauses in \mathbb{C}'_{t+1} are essential. Erasures in π are ignored. For inference steps, \mathbb{C}'_t contains only essential clauses by induction, $\mathbb{C}'_{t+\frac{1}{3}}$ is essential by inference, and $\mathbb{C}'_{t+\frac{2}{3}}$ and \mathbb{C}'_{t+1} are essential since they are derived by erasure from essential configurations. Finally, it is clear that π' performs a subset of the derivation steps in π and that the length, width, and space does not increase. \square

A.3 Pebble Games

Pebble games were devised for studying programming languages and compiler construction, but have found a variety of applications in computational complexity theory. In connection with resolution, pebble games have been employed both to analyze resolution derivations with respect to how much memory they consume (using the original definition of space in [ET01]) and to construct CNF formulas which are hard for different variants of resolution in various respects (see for example [AJPU02, BSIW04, BEGJ00, BOP03] and the sequence of papers [Nor06, NH08b, BSN08] leading up to this work). An excellent survey of pebbling up to ca 1980 is [Pip80].

The black pebbling price of a DAG G captures the memory space, i.e., the number of registers, required to perform the deterministic computation described by G . The space of a non-deterministic computation is measured by the black-white pebbling price of G . We say that vertices of G with indegree 0 are *sources* and that vertices with outdegree 0 are *sinks* (or *targets*). In the following, unless otherwise stated we will assume that all DAGs under discussion have a unique sink and this sink will always be denoted z . The next definition is adapted from [CS76], though we use the established pebbling terminology introduced by [HPV77].

Definition A.7 (Black-white pebble game). Suppose that G is a DAG with sources S and a unique sink z . The *black-white pebble game* on G is the following one-player game. At any point in the game, there are black and white pebbles placed on some vertices of G , at most one pebble per vertex. A *pebble configuration* is a pair of subsets $\mathbb{P} = (B, W)$ of $V(G)$, comprising the black-pebbled vertices B and white-pebbled vertices W . The rules of the game are as follows:

1. If all immediate predecessors of an empty vertex v have pebbles on them, a black pebble may be placed on v . In particular, a black pebble can always be placed on any vertex in S .
2. A black pebble may be removed from any vertex at any time.
3. A white pebble may be placed on any empty vertex at any time.
4. If all immediate predecessors of a white-pebbled vertex v have pebbles on them, the white pebble on v may be removed. In particular, a white pebble can always be removed from a source vertex.

A *black-white pebbling* from (B_0, W_0) to (B_τ, W_τ) in G is a sequence of pebble configurations $\mathcal{P} = \{\mathbb{P}_0, \dots, \mathbb{P}_\tau\}$ such that $\mathbb{P}_0 = (B_0, W_0)$, $\mathbb{P}_\tau = (B_\tau, W_\tau)$, and for all $t \in [\tau]$, \mathbb{P}_t follows from \mathbb{P}_{t-1} by one of the rules above. A (*complete*) *pebbling of G* , also called a *pebbling strategy for G* , is a pebbling such that $(B_0, W_0) = (\emptyset, \emptyset)$ and $(B_\tau, W_\tau) = (\{z\}, \emptyset)$.

The *time* of a pebbling $\mathcal{P} = \{\mathbb{P}_0, \dots, \mathbb{P}_\tau\}$ is simply $\text{time}(\mathcal{P}) = \tau$ and the *space* is $\text{space}(\mathcal{P}) = \max_{0 \leq t \leq \tau} \{|B_t \cup W_t|\}$. The *black-white pebbling price* (also known as the *pebbling measure* or *pebbling number*) of G , denoted $\text{BW-Peb}(G)$, is the minimum space of any complete pebbling of G .

A *black pebbling* is a pebbling using black pebbles only, i.e., having $W_t = \emptyset$ for all t . The (*black*) *pebbling price* of G , denoted $\text{Peb}(G)$, is the minimum space of any complete black pebbling of G .

For any DAG G over n vertices with bounded indegree, the black pebbling price (and thus also the black-white pebbling price) is at most $O(n/\log n)$ [HPV77], where the hidden constant depends on the indegree. A number of exact or asymptotically tight bounds on different graph families have been proven in the whole range from constant to $\Theta(n/\log n)$, for instance in [GT78, Kla85, LT80, PTC77]. As to time, obviously any DAG G over n vertices can be pebbled in time $2n - 1$, and for all graphs we will study this is also a lower bound, so studying the time measure in isolation is not that exciting. A very interesting question, however, is how time and space are related in a single pebbling of G if one wants to optimize both measures simultaneously. We will return to this question in Appendix D.

A.4 Pebbling Contradictions

A *pebbling contradiction* defined on a DAG G is a CNF formula that encodes the pebble game on G by postulating the sources to be true and the target to be false, and specifying that truth propagates through the graph according to the pebbling rules. These formulas have previously been studied in, for instance, [BSW01, BEGJ00, RM99].

Definition A.8 (Pebbling contradiction). Suppose that G is a DAG with sources S and a unique sink z . Identify every vertex $v \in V(G)$ with a propositional logic variable v . The *pebbling contradiction* over G , denoted Peb_G , is the conjunction of the following clauses:

- for all $s \in S$, a unit clause s (*source axioms*),
- For all non-source vertices v with immediate predecessors u_1, \dots, u_ℓ , the clause $\bar{u}_1 \vee \dots \vee \bar{u}_\ell \vee v$ (*pebbling axioms*),
- for the sink z , the unit clause \bar{z} (*target or sink axiom*).

If G has n vertices and maximal indegree ℓ , the formula Peb_G is an unsatisfiable $(1+\ell)$ -CNF formula with n clauses over n variables.

A.5 Substitution Formulas

The following observation is rather immediate, but nevertheless it might be helpful to state it explicitly.

Observation A.9. Suppose for any non-constant Boolean function f_d that $C \in Cl[f_d(\vec{x})]$ and that ρ is any partial truth value assignment such that $\rho(C) = 0$. Then for all $D \in Cl[\neg f_d(\vec{x})]$ it holds that $\rho(D) = 1$.

Proof. If $\rho(C) = 0$ this means that $\rho(f_d) = 0$. Then clearly $\rho(\neg f_d) = 1$, so, in particular, ρ must fix all clauses $D \in Cl[\neg f_d(\vec{x})]$ to true. \square

We have the following easy observation, the proof of which is presented for completeness.

Observation A.10. For any non-constant Boolean function $f_d : \{0, 1\}^d \mapsto \{0, 1\}$, it holds that $F[f_d]$ is unsatisfiable if and only if F is unsatisfiable.

Proof. Suppose that F is satisfiable and let α be a truth value assignment such that $\alpha(F) = 1$. Then we can satisfy $F[f_d]$ by choosing an assignment α' for $Vars(F[f_d])$ in such a way that $f_d(\alpha'(x_1), \dots, \alpha'(x_d)) = \alpha(x)$. For if $C \in F$ is satisfied by some literal a_i set to true by α , then α' will satisfy all clauses $C_i \in a_i[f_d]$ and thus also the whole CNF formula $C[f_d]$ in (1).

Conversely, suppose F is unsatisfiable and consider any truth value assignment α' for $F[f_d]$. Then α' defines a truth value assignment α for F in the natural way by setting $\alpha(x) = f_d(\alpha'(x_1), \dots, \alpha'(x_d))$, and we know that there is some clause $C \in F$ that is not satisfied by α . That is, for every literal $a_i \in C = a_1 \vee \dots \vee a_k$ it holds that $\alpha(a_i) = 0$. But then α' does not satisfy $a_i[f_d]$, so there is some clause $C'_i \in a_i[f_d]$ such that $\alpha'(C'_i) = 0$. This shows that α' falsifies the disjunction $C'_1 \vee \dots \vee C'_k \in C[f_d]$, and consequently $F[f_d]$ must also be unsatisfiable. \square

B Completing the proof of the Substitution Theorem

In this section we prove Theorems 4.1, 4.4, and 4.5, from which the Substitution Space Theorem 3.2 follows.

For convenience of notation, let us define the disjunction $\mathbb{C} \vee \mathbb{D}$ of two clause sets \mathbb{C} and \mathbb{D} to be the clause set

$$\mathbb{C} \vee \mathbb{D} = \{C \vee D \mid C \in \mathbb{C}, D \in \mathbb{D}\} . \quad (7)$$

This notation extends to more than two clause sets in the natural way. Rewriting (1) in Definition 2.4 using this notation, we have that

$$(D \vee a)[f_d] = D[f_d] \vee a[f_d] = \bigwedge_{C_1 \in D[f_d]} \bigwedge_{C_2 \in a[f_d]} (C_1 \vee C_2) . \quad (8)$$

B.1 Proof of Theorem 4.1

Given $\pi : F \vdash 0$, we construct $\pi_f : D[f_d] \vdash 0$ by maintaining the invariant that if we have \mathbb{C} in memory for π , then we have $\mathbb{C}[f_d]$ in memory for π_f . We get the following case analysis.

Axiom download If π downloads C , we download all of $C[f_d]$, i.e., less than $2^{d \cdot W(C)}$ clauses which all have width at most $d \cdot W(C)$.

Erasure If π erases C , we erase all of $C[f_d]$ in less than $2^{d \cdot W(C)}$ erasure steps.

Inference This is the only interesting case. Suppose that π infers $C_1 \vee C_2$ from $C_1 \vee x$ and $C_2 \vee \bar{x}$. Then by induction we have $(C_1 \vee x)[f_d]$ and $(C_2 \vee \bar{x})[f_d]$ in memory in π_f . It is a straightforward extension of Observation A.10 that if $\mathbb{C} \models D$, then $\mathbb{C}[f_d] \models D[f_d]$, so in particular it holds that $(C_1 \vee x)[f_d]$ and $(C_2 \vee \bar{x})[f_d]$ imply $(C_1 \vee C_2)[f_d]$. By the implicational completeness of resolution, these clauses can all be derived.

A not necessarily tight upper bound for the width of this derivation in π_f is $d \cdot (W(C_1 \vee x) + W(C_2 \vee \bar{x}) + W(C_1 \vee C_2)) = O(d \cdot W(\pi))$, as claimed.

To bound the length, note that $(C_1 \vee C_2)[f_d]$ contains less than $2^{d \cdot W(C_1 \vee C_2)}$ clauses. For every clause $D \in (C_1 \vee C_2)[f_d]$, consider the minimal restriction $\rho(\neg D)$ falsifying D . Since

$$(C_1 \vee x)[f_d] \wedge (C_2 \vee \bar{x})[f_d] \models D \quad (9)$$

we have that

$$(C_1 \vee x)[f_d] \upharpoonright_{\rho(\neg D)} \wedge (C_2 \vee \bar{x})[f_d] \upharpoonright_{\rho(\neg D)} \models 0 . \quad (10)$$

The number of variables is at most $d \cdot (W(C_1 \vee C_2) + 1) = N$, and by Observation A.1 there is a refutation of $(C_1 \vee x)[f_d] \upharpoonright_{\rho(\neg D)} \wedge (C_2 \vee \bar{x})[f_d] \upharpoonright_{\rho(\neg D)}$ in length at most $2^{N+1} - 1$. Looking at this refutation and removing the restriction $\rho(\neg D)$, it is straightforward to verify that we get a derivation of D from $(C_1 \vee x)[f_d] \wedge (C_2 \vee \bar{x})[f_d]$ in the same length (see, for instance, the inductive proof in [BSW01]). We can repeat this for every clause $D \in (C_1 \vee C_2)[f_d]$ to derive all of the less than $2^{d \cdot (W(C_1 \vee C_2))}$ clauses in this set in total length at most

$$2^{d \cdot (W(C_1 \vee C_2))} \cdot 2^{d \cdot (W(C_1 \vee C_2) + 2)} \leq 2^{3d \cdot W(\pi)} = 2^{O(d \cdot W(\pi))} . \quad (11)$$

Taken together, we see that we get a refutation π_f in length at most $L(\pi) \cdot 2^{O(d \cdot W(\pi))}$ and width at most $O(d \cdot W(\pi))$. Theorem 4.1 follows.

B.2 Proof of Theorem 4.4

Let us use the convention that \mathbb{D} and D denote clause sets and clauses derived from $F[f_d]$ while \mathbb{C} and C denote clause sets and clauses derived from F .

Let us also overload the notation and write $\mathbb{D} \models C$, $\mathbb{D} \not\models C$, and $\mathbb{D} \triangleright C$ for $C = C^+ \vee C^-$ when the corresponding implications hold or do not hold for \mathbb{D} with respect to $\bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{y \in C^-} \neg f_d(\vec{y})$. Note that it will always be clear when we use the notation in this overloaded sense since \mathbb{D} and C are defined over different sets of variables, so the non-overloaded interpretation would not be very meaningful.

Recall from Definition 4.3 that $\text{proj}_F(\mathbb{D}) = \{C \mid \exists \mathbb{D}_C \subseteq \mathbb{D} \text{ s.t. } \mathbb{D}_C \triangleright \bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{y \in C^-} \neg f_d(\vec{y})\}$ is the set of clauses projected by \mathbb{D} . In the spirit of the notational convention just introduced, we will let \mathbb{C}_t be a shorthand for $\text{proj}_F(\mathbb{D}_t)$.

Suppose now that $\pi_f = \{\mathbb{D}_0, \dots, \mathbb{D}_\tau\}$ is a resolution refutation of $F[f_d]$ for some arbitrary unsatisfiable CNF formula F and some arbitrary non-constant function f_d .

The first two bullets in Theorem 4.4 are immediate. For $\mathbb{D}_0 = \emptyset$ we have $\mathbb{C}_0 = \text{proj}_F(\mathbb{D}_0) = \emptyset$, and it is easy to verify that $\mathbb{D}_\tau = \{0\}$ yields $\mathbb{C}_\tau = \text{proj}_F(\mathbb{D}_\tau) = \{0\}$. We note, however, that the empty clause will have appeared in $\mathbb{C}_t = \text{proj}_F(\mathbb{D}_t)$ earlier, namely for the first t such that \mathbb{D}_t is contradictory.

Perhaps the trickiest part is to show that all transitions from $\mathbb{C}_{t-1} = \text{proj}_F(\mathbb{D}_{t-1})$ to $\mathbb{C}_t = \text{proj}_F(\mathbb{D}_t)$ can be performed in such a way that the variable support size in our refutation under construction $\pi : F \vdash 0$ never exceeds $\max\{\text{SuppSize}(\mathbb{C}_{t-1}), \text{SuppSize}(\mathbb{C}_t)\}$ during the intermediate derivation steps needed in π . The proof is by a case analysis of the derivation steps. Before plunging into the proof, let us make a simple but useful observation.

Observation B.1. *Using the above notation, if $\mathbb{D}_t \models C$ then $C = C^+ \vee C^-$ is derivable from $\mathbb{C}_t = \text{proj}_F(\mathbb{D}_t)$ by weakening.*

Proof. Pick $\mathbb{D}' \subseteq \mathbb{D}_t$, $C_1^+ \subseteq C^+$, and $C_2^- \subseteq C^-$ minimal so that $\mathbb{D}' \models C_1^+ \vee C_2^-$ still holds. Then by definition $\mathbb{D}' \triangleright C_1^+ \vee C_2^-$ so $C_1^+ \vee C_2^- \in \mathbb{C}_t$ and $C \supseteq C_1^+ \vee C_2^-$ can be derived from \mathbb{C}_t by weakening as claimed. \square

Consider now the rule applied in π_f at time t to get from \mathbb{D}_{t-1} to \mathbb{D}_t . We analyze the three possible cases—inference, erasure and axiom download—in this order.

B.2.1 Inference

Since $\mathbb{D}_t \supseteq \mathbb{D}_{t-1}$, it is immediate from Definition 4.3 that no clauses in \mathbb{C}_{t-1} can disappear at time t , i.e., $\mathbb{C}_{t-1} \setminus \mathbb{C}_t = \emptyset$. There can appear new clauses in \mathbb{C}_t , but by Observation B.1 all such clauses are derivable by weakening from \mathbb{C}_{t-1} . During such weakening moves the variable support size increases monotonically and is bounded from above by $\text{SuppSize}(\mathbb{C}_t)$.

B.2.2 Erasure

Since $\mathbb{D}_{t-1} \subseteq \mathbb{D}_t$, it is immediate from Definition 4.3 that no new clauses can appear at time t . Any clauses in $\mathbb{C}_{t-1} \setminus \mathbb{C}_t$ can simply be erased, which decreases the variable support size monotonically.

B.2.3 Axiom download

This is the only place in the case analysis where we need to do some work. Suppose that $\mathbb{D}_t = \mathbb{D}_{t-1} \cup \{D\}$ for some axiom clause $D \in A[f_d]$, where A in turn is an axiom of F . If $C \in \mathbb{C}_t \setminus \mathbb{C}_{t-1}$ is a new projected clause, D must be involved in projecting it so there is some subset $\mathbb{D} \subseteq \mathbb{D}_{t-1}$ such that

$$\mathbb{D} \cup \{D\} \triangleright C . \quad (12)$$

Also note that if $\mathbb{D}_{t-1} \models C$ we are done since C can be derived from \mathbb{C}_{t-1} by weakening, so we can assume that

$$\mathbb{D}_{t-1} \not\models C . \quad (13)$$

We want to show that all such clauses C can be derived from $\mathbb{C}_{t-1} = \text{proj}_F(\mathbb{D}_{t-1})$ by downloading $A \in F$, making inferences, and then possibly erasing A , and that this can be done without the variable support size exceeding $\max\{\text{SuppSize}(\mathbb{C}_{t-1}), \text{SuppSize}(\mathbb{C}_t)\}$. The key to our proof is the next lemma.

Lemma B.2. *Suppose that \mathbb{D} derived from $D \in F[f_d]$, $D \in A[f_d]$, and C a clause over $\text{Vars}(F)$ are such that $\mathbb{D} \cup \{D\} \triangleright C$ but $\mathbb{D} \not\models C$. Then if $A = a_1 \vee \dots \vee a_k$, for every $a_i \in A \setminus C$ there is a clause subset $\mathbb{D}^i \subseteq \mathbb{D}$ and a subclause $C^i \subseteq C$ such that $\mathbb{D}^i \triangleright C^i \vee \bar{a}_i$. That is, all clauses $C \vee \bar{a}_i$ for $a_i \in A \setminus C$ can be derived from $\mathbb{C} = \text{proj}_F(\mathbb{D})$ by weakening.*

Proof. Consider any truth value assignment α such that $\alpha(\mathbb{D}) = 1$ but $\alpha(\bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C^-} \neg f_d(\vec{y})) = 0$. Such an assignment exists since $\mathbb{D} \not\models C$ by assumption. Also, since by assumption $\mathbb{D} \cup \{D\} \triangleright C$ we must have $\alpha(D) = 0$. If $A = a_1 \vee \dots \vee a_k$, we can write $D \in A[f_d]$ on the form $D = D_1 \vee \dots \vee D_k$ for $D_i \in a_i[f_d]$ (compare with (8)). Fix any $a \in A$ and suppose for the moment that $a = x$ is a positive literal. Then $\alpha(D_i) = 0$ implies that $\alpha(f_d(\vec{x})) = 0$. By Observation A.9, this means that $\alpha(\neg f_d(\vec{x})) = 1$. Since exactly the same argument holds if $a = \bar{y}$ is a negative literal, we conclude that

$$\mathbb{D} \models \bigvee_{x \in (C \vee a_i)^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in (C \vee a_i)^-} \neg f_d(\vec{y}) \quad (14)$$

or, rewriting (14) using our overloaded notation, that

$$\mathbb{D} \models C \vee \bar{a}_i . \quad (15)$$

If $a_i \in C$, the clause $C \vee \bar{a}_i$ is trivially true and thus uninteresting, but otherwise we pick $\mathbb{D}^i \subseteq \mathbb{D}$ and $C^i \subseteq C$ minimal such that (15) still holds (and notice that since $\mathbb{D} \not\models C$, the literal \bar{a}_i cannot be dropped from the implication). Then by Definition 4.3 we have $\mathbb{D}^i \triangleright C^i \vee \bar{a}_i$ as claimed. \square

We remark that Lemma B.2 can be seen to imply that $\text{Vars}(A) \subseteq \text{Vars}(\mathbb{C}_t) = \text{Vars}(\text{proj}_F(\mathbb{D}_t))$. For $x \in \text{Vars}(A) \cap \text{Vars}(C)$ this is of course trivially true, but for $x \in \text{Vars}(A) \setminus \text{Vars}(C)$ Lemma B.2 tells us that already at time $t-1$, there is a clause in $\mathbb{C}_{t-1} = \text{proj}_F(\mathbb{D}_{t-1})$ containing x , namely the clause $C^i \vee \bar{a}_i$ found in the proof above. Since $\mathbb{D}_t \supseteq \mathbb{D}_{t-1}$, this clause does not disappear at time t . This means that if we download $A \in F$ in our refutation $\pi : F \vdash 0$ under construction, we have $\text{SuppSize}(\mathbb{C}_{t-1} \cup \{A\}) \leq \text{SuppSize}(\mathbb{C}_t)$.

Thus, we can download $A \in F$, and then possibly erase this clause again at the end of our intermediate resolution derivation to get from \mathbb{C}_{t-1} to \mathbb{C}_t , without the variable support size ever exceeding $\max\{\text{SuppSize}(\mathbb{C}_{t-1}), \text{SuppSize}(\mathbb{C}_t)\}$. Let us now argue that all new clauses $C \in \mathbb{C}_t \setminus \mathbb{C}_{t-1}$ can be derived from $\mathbb{C}_{t-1} \cup \{A\}$.

If $A \setminus C = \emptyset$, then the weakening rule applied on A is enough. Suppose therefore that this is not the case and let $A' = A \setminus C = \bigvee_{a \in \text{Lit}(A) \setminus \text{Lit}(C)} a$. Appealing to Lemma B.2 we know that for every $a \in A$ there is a $C_a \subseteq C$ such that $C_a \vee \bar{a} \in \mathbb{C}_{t-1}$. Note that by assumption (13) this means that if $x \in \text{Vars}(A) \cap \text{Vars}(C)$, then x occurs with the same sign in A and C , since otherwise we would get the contradiction $\mathbb{D} \models C \vee \bar{a} = C$. Summing up, \mathbb{C}_{t-1} contains $C_a \vee \bar{a}$ for some $C_a \subseteq C$ for all $a \in \text{Lit}(A) \setminus \text{Lit}(C)$ and in addition we know that $\text{Lit}(A) \cap \{\bar{a} \mid a \in \text{Lit}(C)\} = \emptyset$. Let us write $A' = a_1 \vee \dots \vee a_m$ and do the following weakening

derivation steps from $\mathbb{C}_{t-1} \cup \{A\}$:

$$\begin{aligned}
 A &\rightsquigarrow C \vee A' \\
 C_{a_1} \vee \bar{a}_1 &\rightsquigarrow C \vee \bar{a}_1 \\
 C_{a_2} \vee \bar{a}_2 &\rightsquigarrow C \vee \bar{a}_2 \\
 &\vdots \\
 C_{a_m} \vee \bar{a}_m &\rightsquigarrow C \vee \bar{a}_m
 \end{aligned} \tag{16}$$

Then resolve $C \vee A'$ in turn with all clauses $C \vee \bar{a}_1, C \vee \bar{a}_2, \dots, C_{a_m} \vee \bar{a}_m$, finally yielding the clause C .

In this way all clauses $C \in \mathbb{C}_t \setminus \mathbb{C}_{t-1}$ can be derived one by one, and we note that we never mention any variables outside of $\text{Vars}(\mathbb{C}_{t-1} \cup \{A\}) \subseteq \text{Vars}(\mathbb{C}_t)$ in these derivations.

B.2.4 Wrapping up the Proof of Theorem 4.4

We have proven that no matter what derivation step is made in the transition $\mathbb{D}_{t-1} \rightsquigarrow \mathbb{D}_t$, we can perform the corresponding transition $\mathbb{C}_{t-1} \rightsquigarrow \mathbb{C}_t$ for our projected clause sets without the variable support size going above $\max\{\text{SuppSize}(\mathbb{C}_{t-1}), \text{SuppSize}(\mathbb{C}_t)\}$. Also, the only time we need to download an axiom $A \in F$ in our projected refutation π of F is when π_f downloads some axiom $D \in A[f_d]$. This completes the proof of Theorem 4.4.

B.3 Proof of Theorem 4.5

Recall the convention that x, y, z refer to variables in F while $x_1, \dots, x_d, y_1, \dots, y_d, z_1, \dots, z_d$ refer to variables in $F[f_d]$. We start with an intuitively plausible lemma saying that for all variables x appearing in some clause projected by \mathbb{D} , the clause set \mathbb{D} itself must contain at least one of the variables x_1, \dots, x_d .

Lemma B.3. *Suppose that \mathbb{D} is a set of clauses derived from $F[f_d]$ and that $C \in \text{proj}_F(\mathbb{D})$. Then for all variables $x \in \text{Vars}(C)$ it holds that $\{x_1, \dots, x_d\} \cap \text{Vars}(\mathbb{D}) \neq \emptyset$.*

Proof. Fix any $\mathbb{D}' \subseteq \mathbb{D}$ such that $\mathbb{D}' \triangleright C$. By definition, for all $z \in \text{Vars}(C)$ we have $\mathbb{D}' \not\equiv C \setminus \{z, \bar{z}\}$. Suppose that z appears as a positive literal in C (the case of a negative literal is completely analogous). This means that there is an assignment α such that $\alpha(\mathbb{D}') = 1$ but $\alpha(\bigvee_{x \in C^+ \setminus \{z\}} f_d(\vec{x}) \vee \bigvee_{y \in C^-} \neg f_d(\vec{y})) = 0$. Since $\mathbb{D}' \triangleright C$, it must hold that $\alpha(f_d(\vec{z})) = 1$. Modify α into α' by changing the assignments to z_1, \dots, z_d in such a way that $\alpha'(f_d(\vec{z})) = 0$. Then $\alpha'(\bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{y \in C^-} \neg f_d(\vec{y})) = 0$, so we must have $\alpha'(\mathbb{D}') = 0$. Since we only changed the assignments to (a subset of) the variables z_1, \dots, z_d , the clause set $\mathbb{D}' \subseteq \mathbb{D}$ must mention at least one of these variables. \square

With Lemma B.3 in hand, we are ready to prove Theorem 4.5. Note that everything said so far in Section 4 and Appendix B (in particular, all of the proofs) applies to any non-constant Boolean function. In the proof of Theorem 4.5, however, it will be essential that we are dealing with non-authoritarian functions, i.e., functions f_d having the property that no single variable x_i can fix the value of $f_d(x_1, \dots, x_d)$.

Suppose that \mathbb{D} is a set of clauses derived from $F[f_d]$ and write $V^* = \text{Vars}(\text{proj}_F(\mathbb{D}))$ to denote the set of all variables in $\text{Vars}(F)$ appearing in any clause projected by \mathbb{D} . We want to prove that $Sp(\mathbb{D}) = |\mathbb{D}| > |V^*|$ provided that f_d is non-authoritarian.

To this end, consider the bipartite graph with the clauses in \mathbb{D} labelling the vertices on the left-hand side and variables in V^* labelling the vertices on the right-hand side. We draw an edge between $D \in \mathbb{D}$ and $x \in V^*$ if $\text{Vars}(D) \cap \{x_1, \dots, x_d\} \neq \emptyset$. By Lemma B.3 it holds that $\text{Vars}(\mathbb{D}) \cap \{x_1, \dots, x_d\} \neq \emptyset$ for all variables $x \in V^*$, so in particular every variable $x \in V^*$ is the neighbour of at least one clause $D \in \mathbb{D}$. Let

us write $N(D)$ to denote the neighbours of a left-hand vertex D and extend this notation to sets of vertices by taking unions.

Fix $\mathbb{D}_1 \subseteq \mathbb{D}$ to be any largest subset such that $|\mathbb{D}_1| > N(\mathbb{D}_1)$. If $\mathbb{D}_1 = \mathbb{D}$ we are done (remember that $N(\mathbb{D}) = V^*$), so suppose $\mathbb{D}_1 \neq \mathbb{D}$. We show that this assumption leads to a contradiction.

Let $\mathbb{D}_2 = \mathbb{D} \setminus \mathbb{D}_1 \neq \emptyset$ and define the vertex sets $V_1^* = N(\mathbb{D}_1)$ and $V_2^* = V^* \setminus V_1^*$. Note that we must have $V_2^* \subseteq N(\mathbb{D}_2)$ since $N(\mathbb{D}) = N(\mathbb{D}_1) \cup N(\mathbb{D}_2) = V^*$. By the maximality of \mathbb{D}_1 it must hold for all $\mathbb{D}' \subseteq \mathbb{D}_2$ that $|\mathbb{D}'| \leq |N(\mathbb{D}') \setminus V_1^*|$, because otherwise $\mathbb{D}'' = \mathbb{D}_1 \cup \mathbb{D}'$ would be a larger set with $|\mathbb{D}''| > |N(\mathbb{D}'')|$. But this means that by Hall's marriage theorem, there is a matching M of \mathbb{D}_2 into $N(\mathbb{D}_2) \setminus V_1^* = V_2^*$. Consider any clause $C \in \text{proj}_F(\mathbb{D})$ such that $\text{Vars}(C) \cap V_2^* \neq \emptyset$ and let $\mathbb{D}' \subseteq \mathbb{D}$ be any clause set such that

$$\mathbb{D}' \triangleright \bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C^-} \neg f_d(\vec{y}) \quad (17)$$

(the existence of which is guaranteed by Definition 4.3). We claim that we can construct an assignment α that makes \mathbb{D}' true but $\bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C^-} \neg f_d(\vec{y})$ false. This is clearly a contradiction, so if we can prove this claim it follows that our assumption $\mathbb{D}_1 \neq \mathbb{D}$ is false and that it instead must hold that $\mathbb{D}_1 = \mathbb{D}$ and thus $|N(\mathbb{D})| = |V^*| < |\mathbb{D}|$, which proves the theorem.

To establish the claim, let $\mathbb{D}'_i = \mathbb{D}' \cap \mathbb{D}_i$ for $i = 1, 2$ and let $C_i = C_i^+ \vee C_i^-$ for

$$C_i^+ = \bigvee_{\substack{x \in C \\ x \in V_i^*}} x \quad \text{and} \quad C_i^- = \bigvee_{\substack{\bar{y} \in C \\ y \in V_i^*}} \bar{y} \quad (18)$$

and $i = 1, 2$. We construct the assignment α satisfying \mathbb{D}' but falsifying $\bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C^-} \neg f_d(\vec{y})$ in three steps:

1. Since $C_1^+ \vee C_1^- = C_1 \not\subseteq C$ by construction (recall that we chose our clause C in such a way that $\text{Vars}(C) \cap V_2^* \neq \emptyset$), the minimality condition in Definition 4.3 yields that

$$\mathbb{D}'_1 \not\models \bigvee_{x \in C_1^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C_1^-} \neg f_d(\vec{y}) \quad (19)$$

and hence we can find a truth value assignment α_1 that sets \mathbb{D}'_1 to true, all $f_d(x_1, \dots, x_d)$, $x \in C_1^+$, to false, and all $f_d(y_1, \dots, y_d)$, $\bar{y} \in C_1^-$, to true. Note that α_1 need only assign values to $\{z_1, \dots, z_d \mid z \in \text{Vars}(C_1)\}$.

2. For \mathbb{D}'_2 , we use the matching M into V_2^* found above to pick a distinct variable $x(D) \in \text{Vars}(F)$ for every $D \in \mathbb{D}'_2$ and then a variable $x(D)_i \in \text{Vars}(F[f_d])$ appearing in D , the existence of which is guaranteed by the edge between D and $x(D)$. Let α_2 be the assignment that sets all these variables $x(D)_i$ to the values that fix all $D \in \mathbb{D}'_2$ to true. We stress that α_2 assigns a value to at most one variable $x(D)_i$ for every $x(D) \in \text{Vars}(F)$.
3. But since f_d is non-authoritarian, this means that we can extend α_2 to an assignment to all variables $x(D)_1, \dots, x(D)_d$ that still satisfies \mathbb{D}'_2 but sets all $f_d(x_1, \dots, x_d)$, $x \in C_2^+$, to false and all $f_d(y_1, \dots, y_d)$, $\bar{y} \in C_2^-$, to true.

Hence, $\alpha = \alpha_1 \cup \alpha_2$ is an assignment such that $\alpha(\mathbb{D}') = 1$ but $\alpha(\bigvee_{x \in C^+} f_d(\vec{x}) \vee \bigvee_{\bar{y} \in C^-} \neg f_d(\vec{y})) = 0$, which proves the claim. This concludes the proof of Theorem 4.5.

Since Theorems 4.1, 4.4, and 4.5 have now all been established, the proof of Theorem 3.2 is finished.

C Reductions Between Resolution and Pebbling

It is not hard to see how a black pebbling \mathcal{P} of a DAG G can be used to construct a resolution refutation of the pebbling contradiction Peb_G in Definition A.8 in length and space upper-bounded by $time(\mathcal{P})$ and $space(\mathcal{P})$, respectively. It is straightforward to show that this translation from pebblings to refutations works even if we do an f_d -substitution in the pebbling contradiction. We present a proof of this fact in Section C.1.

Using our new results in Section 3, we can prove the more surprising fact that there is also a fairly tight reduction in the other direction: provided that the function f_d is non-authoritarian, any resolution refutation of $Peb_G[f_d]$ translates into a black-white pebbling of G with the same time-space properties (adjusting for constant factors depending on the function f_d and the maximal indegree of G). This new reduction is given in Section C.2.

Finally, in Section C.3 we appeal to both of these reductions to prove a meta-theorem saying that for DAGs G having the right time-space trade-off properties, we can prove that pebbling contradictions defined over such DAGs inherit the same trade-off properties. This will allow us, after having studied pebbling time-space trade-offs in Appendix D, to prove a wealth of strong trade-offs for both clause space and variable space in resolution in Appendix E.

C.1 From Black Pebblings to Resolution Refutations

Given any black-only pebbling \mathcal{P} of a DAG G , we can mimic this pebbling in a resolution refutation of Peb_G by deriving that a literal v is true whenever the corresponding vertex in G is pebbled (this was perhaps first observed in [BSIW04]). This construction carries over also to substitution formulas $Peb_G[f_d]$ and we have the following theorem.

Theorem C.1. *Let f_d be a non-constant Boolean function of arity d and let G be a DAG with indegree at most ℓ and unique sink z . Then given any complete black pebbling \mathcal{P} of G , we can construct a resolution refutation $\pi : Peb_G[f_d] \vdash 0$ such that*

$$\begin{aligned} L(\pi) &\leq time(\mathcal{P}) \cdot \exp(O(d(\ell + 1))) \ , \\ W(\pi) &\leq d(\ell + 1) \ , \text{ and} \\ VarSp(\pi) &\leq space(\mathcal{P}) \cdot \exp(O(d(\ell + 1))) \ . \end{aligned}$$

Before presenting the proof, we note that in our applications we will have the function arity d and the DAG indegree ℓ fixed (we can for instance pick $d = \ell = 2$), which means that the bounds on length and space above turns into $L(\pi) \leq O(time(\mathcal{P}))$ and $VarSp(\pi) \leq O(space(\mathcal{P}))$. We also remark that for concrete functions f_d , such as for instance XOR over two variables, we can easily compute explicit upper bounds on the constants hidden in the asymptotic notation if we so wish, and that these constants are small.

Proof of Theorem C.1. The proof is by induction over the black pebbling \mathcal{P} . We maintain the invariant that if at time t we have black pebbles on the vertices in V , then π will contain exactly the clauses $\mathbb{C}_t = \{x[f_d] \mid x \in V\}$. Again, to simplify the notation in the proof we will use fractional time steps in π , making sure that it never takes more than $\exp(O(d(\ell + 1)))$ time steps to get from \mathbb{C}_{t-1} to \mathbb{C}_t .

Consider the pebbling move made in \mathcal{P} at time t :

1. If \mathcal{P} places a pebble on a source vertex s , we download the less than 2^d axioms in $s[f_d]$.
2. If \mathcal{P} places a pebble on a non-source vertex v with immediate predecessors u_1, \dots, u_ℓ , by induction we have $\{u_i[f_d] \mid i = 1, \dots, \ell\} \subseteq \mathbb{C}_{t-1}$. The argument in this case is very similar to the one in Section B.1.

First download the less than $2^{d(\ell'+1)}$ pebbling axioms in $(\bar{u}_1 \vee \dots \vee \bar{u}_{\ell'} \vee v)[f_d]$. Now

$$\{u_i[f_d] \mid i = 1, \dots, \ell'\} \cup \{(\bar{u}_1 \vee \dots \vee \bar{u}_{\ell'} \vee v)[f_d]\} \quad (20)$$

implies all clauses $D \in v[f_d]$. If we apply the restriction $\rho(-D)$ to the clause set (20) we can obtain a refutation in length and variable space at most $\exp(O(d(\ell+1)))$ (and trivially in width at most $d(\ell+1)$) by Observation A.1. Removing the restriction $\rho(-D)$, this refutation turns into a derivation of D . Doing this for all of the less than 2^d clauses $D \in v[f_d]$ completes the induction step.

3. If \mathcal{P} removes a pebble from any vertex v , we erase the less than 2^d clauses in $v[f_d]$ from memory.

At the end of the pebbling \mathcal{P} , we have $\mathbb{C}_\tau = \{z[f_d]\}$ for z the sink of G . We conclude the refutation by downloading all the sink axioms in $\bar{z}[f_d]$ and deriving the empty clause 0 in length $\exp(O(d))$, width d and variable space $\exp(O(d))$. \square

C.2 From Resolution Refutations to Black-White Pebblings

Let us now see how we can go in the other direction from resolution refutations to pebbling strategies.

Theorem C.2. *Let f be any non-authoritarian Boolean function and G be any DAG with unique sink and bounded indegree ℓ . Then from any resolution refutation $\pi : \text{Peb}_G[f] \vdash 0$ we can extract a black-white pebbling strategy \mathcal{P}_π for G such that $\text{time}(\mathcal{P}_\pi) \leq (\ell+1) \cdot L(\pi)$ and $\text{space}(\mathcal{P}_\pi) \leq Sp(\pi)$.*

Note, however, that Theorems C.1 and C.2 are not perfect converses, since we do not know any way of translating black-white pebbling strategies into resolution refutations that preserve the time and space properties (and our guess would be that there is no way of doing this in general for an arbitrary black-white pebbling, but this is an open problem).

The proof of Theorem C.2 is in three steps:

1. First, we convert $\pi : \text{Peb}_G[f] \vdash 0$ to a refutation π' of Peb_G such that $\text{SuppSize}(\pi') \leq Sp(\pi)$ and the number of axiom downloads in π' is upper-bounded by the number of axiom downloads in π . This is Theorem 3.2, which is the key technical contribution of this paper.
2. The refutation $\pi' : \text{Peb}_G \vdash 0$ can contain weakening moves however, and we do not want that, so we appeal to Proposition A.2 to get a refutation $\pi'' : \text{Peb}_G \vdash 0$ without any weakening steps. By Lemma A.6, without loss of generality we can assume that π'' is frugal (Definition A.5). This part of the proof just uses standard techniques, and the number of axiom downloads and the variable support size can only decrease when going from π' to π'' .
3. Finally, we show that π'' corresponds to a black-white pebbling strategy \mathcal{P} for G such that $\text{time}(\mathcal{P})$ is upper-bounded by the number of axiom downloads and $\text{space}(\mathcal{P})$ by the maximal number of variables occurring simultaneously in π'' . This final part relies heavily on the work [BS02] by the first author. Since we need a more detailed result than can be read off from that paper, however, we present the full construction below.

Putting together these three steps, Theorem C.2 clearly follows. What remains is thus to prove the following lemma.

Lemma C.3. *Let G be any DAG with unique sink and bounded indegree ℓ , and suppose that π is any resolution refutation of Peb_G without weakening that is also frugal. Then there is a black-white pebbling strategy \mathcal{P}_π for G such that $\text{space}(\mathcal{P}_\pi) \leq \text{SuppSize}(\pi)$ and $\text{time}(\mathcal{P}_\pi)$ is at most $(\ell+1)$ times the number of axiom downloads in π .*

Proof. Given a refutation $\pi = \{\mathbb{C}_0 = \emptyset, \mathbb{C}_1, \dots, \mathbb{C}_\tau = \{0\}\}$ of Peb_G , we translate every clause set \mathbb{C}_t into a black-white pebble configuration $\mathbb{P}_t = (B_t, W_t)$ using a slightly modified version of the ideas in [BS02], and then show that $\mathcal{P} = \{\mathbb{P}_0, \dots, \mathbb{P}_\tau\}$ is essentially a legal black-white pebbling of G as in the statement of the lemma. The translation will satisfy the invariant that $B_t \cup W_t = Vars(\mathbb{C}_t)$ which yields the upper bound on space in terms of variable support size. The first configuration $\mathbb{C}_0 = 0$ is thus translated into $\mathbb{P}_0 = (\emptyset, \emptyset)$.

Suppose inductively that (B_{t-1}, W_{t-1}) has been constructed from \mathbb{C}_{t-1} and consider all the variables $x \in Vars(\mathbb{C}_t)$ one by one. If $x \in Vars(\mathbb{C}_t) \cap B_{t-1}$, keep x in B_t . Otherwise, if $x \in Lit(\mathbb{C}_t)$ appears as a positive literal, add x to B_t . Otherwise, if $\bar{x} \in Lit(\mathbb{C}_t)$, add x to W_t . This is our translation of \mathbb{C}_t into black pebbles B_t and white pebbles W_t .

To see that this translation yields a legal pebbling, consider the derivation rule applied to get from \mathbb{C}_{t-1} to \mathbb{C}_t .

Axiom download Suppose that we download the pebbling or source axiom for a vertex v with immediate predecessors $u_1, \dots, u_{\ell'}$ (where we can have $\ell' = 0$). All predecessors u_i not having pebbles on them at time $t-1$ get white pebbles. Then v gets a black pebble, if it is not already pebbled. Note that this is a legal pebble placement since all immediate predecessors of v have pebbles at this point. We remark that to black-pebble v , we might have to remove a white pebble from v first, but since all immediate predecessors have pebbles on them this poses no problems. Also, downloading the sink axiom can at most place a white pebble on the sink z , which is in order. By the bound on the indegree, this step involves placing at most $\ell + 1$ pebbles.

Inference In this case $Vars(\mathbb{C}_{t-1}) = Vars(\mathbb{C}_t)$, so nothing happens.

Erasure Suppose that the clause erased in C . Just apply the translation function. Suppose that this results in a pebble on x disappearing. Then we have $x \in Vars(C)$ but $x \notin Vars(\mathbb{C}_t)$. Before being erased, C has been resolved with some other clause (recall that π is frugal). But as long as we did not resolve over the variable x , we will still have $x \in Vars(\mathbb{C}_t)$, and hence \mathbb{C} must have been resolved over x at some time $t' < t$. At this time x appeared both positively and negatively in $\mathbb{C}_{t'}$, and in view of how we defined the translation from clauses to pebbles, this means that the vertex x has contained a black pebble in the interval $[t', t-1]$. Thus the pebble disappearing at time t is black, and black pebbles can always be removed freely.

To conclude the proof, note that during the course of the refutation all axioms must have been downloaded at least once, since Peb_G is easily seen to be minimally unsatisfiable. In particular, this means that the sink z is black-pebbled at some time during the proof, and we can decide to keep the black pebble on z from that moment onwards. (This potentially adds one pebble extra to the pebbling space, but this is fine since the inequality in Theorem 3.2 is strict so there is margin for this.)

Since every time an axiom is downloaded it must also be erased at some later time, we get the time bound of $(\ell + 1)$ times the number of axiom downloads (and in fact it is easy to see that this bound can be improved by taking into account the inference steps, when nothing happens in the pebbling). The lemma follows. \square

As was discussed above, Lemma C.3 completes the proof of Theorem C.2.

C.3 Obtaining Resolution Trade-offs from Pebbling

Combining Theorems C.1 and C.2, we can now prove that if we can find DAGs G with appropriate pebbling trade-off properties, such DAGs immediately yield trade-off results in resolution. And as we will see in Appendix D, there are (explicitly constructible) such DAGs.

In order not to clutter the statement of the next theorem, we assume that the arity d of the Boolean function f and the indegree ℓ of the DAG are fixed, so that any dependence on d and ℓ can be hidden in the asymptotical notation. (This is not much of a restriction since we will have $d = \ell = 2$ in the applications that we care about.)

Theorem C.4. *Let d and ℓ be universal constants, and let f be some universally fixed non-authoritarian Boolean function of arity d . Suppose that G is a DAG with n vertices, unique sink z , and bounded indegree ℓ , and that $g, h : \mathbb{N}^+ \mapsto \mathbb{N}^+$ are functions satisfying the following properties:*

- *For every $s \geq \text{Peb}(G)$ there is a complete black pebbling \mathcal{P} of G with $\text{space}(\mathcal{P}) \leq s$ and $\text{time}(\mathcal{P}) \leq g(s)$.*
- *For every $s \geq \text{BW-Peb}(G)$ and every complete black-white pebbling \mathcal{P} of G with $\text{space}(\mathcal{P}) \leq s$ it holds that $\text{time}(\mathcal{P}) \geq h(s)$.*

Then the following holds for $\text{Peb}_G[f]$:

1. $\text{Peb}_G[f]$ is a k -CNF formula for some fixed $k = k(d, \ell, f)$ and has size $O(n)$.
2. $\text{Peb}_G[f]$ is refutable in length $L(\text{Peb}_G[f] \vdash 0) = O(n)$ and width $W(\text{Peb}_G[f] \vdash 0) = O(1)$ simultaneously, and is also refutable in variable space $\text{VarSp}(\text{Peb}_G[f] \vdash 0) = O(\text{Peb}(G))$.
3. For every $s \geq \text{Peb}(G)$ there is a resolution refutation $\pi_s : \text{Peb}_G[f] \vdash 0$ in length $L(\pi_s) = O(g(s))$ and variable space $\text{VarSp}(\pi_s) = O(s)$.
4. The clause space of any resolution refutation is lower-bounded by $\text{Sp}(\text{Peb}_G[f] \vdash 0) \geq \text{BW-Peb}(G)$, and for every $s \geq \text{BW-Peb}(G)$ and every refutation $\pi_s : \text{Peb}_G[f] \vdash 0$ in clause space $\text{Sp}(\pi_s) \leq s$, it holds that $L(\pi_s) = \Omega(h(s))$.

All hidden constants in the asymptotical notation depend only on d, ℓ , and f , and are independent of G .

Proof. Item 1 is an easy consequence of Definition 2.4. Items 2 and 3 both follow from Theorem C.1 (to get item 2, consider the trivial pebbling that black-pebbles all vertices of G in topological order). Finally, Theorem C.2 yields item 4. □

This theorem will be of particular interest when we can find graph families $\{G_n\}_{n=1}^\infty$ with $\text{Peb}(G_n) = \Theta(\text{BW-Peb}(G_n))$ having trade-off functions $g_n(s) = \Theta(h_n(s))$. For such families of DAGs, Theorem C.4 yields asymptotically tight trade-offs in resolution for both clause space and variable space with respect to length (since the upper bounds are in terms of variable space and the lower bounds in terms of clause space).

D Some Old and New Pebbling Results

Having come this far in the paper, we know that if we can find graphs with trade-off results for black-white pebbling and matching upper bounds for black pebbling, we can construct CNF formulas from these graphs with similar time-space trade-off properties in resolution. And indeed, as we show in this section, we can find graphs satisfying these properties (or in one case graphs that come sufficiently close for us to be able to get the desired result via some extra work).

First, we present some auxiliary definitions, notation and terminology in Section D.1. Then, in Section D.2, we prove a strong trade-off result for a very simple but surprisingly versatile family of graphs. Our results build on [CS80, CS82] and extend the results there from black-only to black-white pebbling. Finally, in Section D.3 we review a number of results from [LT82] that will also enable us to get strong trade-offs in resolution.

We remark that all the pebbling trade-off results presented in this section are for explicitly constructible graphs.

D.1 Pebbling Preliminaries

We will use the following notational conventions:

- n denotes the size (i.e., the number of vertices) of a DAG, or, in some cases where it is more convenient, the size to within a (small) constant factor.
- ℓ denotes the maximal indegree of a DAG.
- s denotes pebbling space (although s_1, s_2, \dots will sometimes denote source vertices of DAGs).
- $S(G)$ denotes the source vertices of G and $Z(G)$ denotes the sink vertices of G .

We say that the pebbling move *at time* σ is the move resulting in the pebble configuration \mathbb{P}_σ .

D.1.1 Technical Definitions and Some Observations

We need to generalize our definition of pebbling slightly to distinguish slightly different variants of pebbings and also to allow pebbings of graphs with more than one sink.

Definition D.1 (Conditional, persistent and visiting pebbings). Suppose that G is a DAG with sources S and sinks Z (one or many). Let the pebble game rules be as in Definition A.7, and define pebbling space in the same way.

We say that a pebbling $\mathcal{P} = \{\mathbb{P}_0, \dots, \mathbb{P}_\tau\}$ is *conditional* if $\mathbb{P}_0 \neq (\emptyset, \emptyset)$ and *unconditional* otherwise. Note that complete pebbings, or pebbling strategies, are always unconditional.

A complete black-white pebbling *visiting* Z is a pebbling $\mathcal{P} = \{\mathbb{P}_0, \dots, \mathbb{P}_\tau\}$ such that $\mathbb{P}_0 = \mathbb{P}_\tau = (\emptyset, \emptyset)$ and such that for every $z \in Z$, there exists a time $t_z \in [\tau]$ such that $z \in B_{t_z} \cup W_{t_z}$. The minimum space of such a visiting pebbling is denoted $BW\text{-Peb}^\emptyset(G)$, and for black pebbling we have the measure $\text{Peb}^\emptyset(G)$.

A *persistent* pebbling of G is a pebbling \mathcal{P} such that $\mathbb{P}_\tau = (Z, \emptyset)$. The minimum space of any complete persistent black-white or black-only pebbling of G is denoted $BW\text{-Peb}^z(G)$ and $\text{Peb}^z(G)$, respectively.

That is, a visiting pebbling touches all sinks but leaves the graph empty at time τ , whereas a persistent pebbling leaves black pebbles on all sinks at the end of the pebbling. If G is a DAG with m sinks, then it clearly holds that $BW\text{-Peb}^z(G) \leq BW\text{-Peb}^\emptyset(G) + m$ and $\text{Peb}^z(G) \leq \text{Peb}^\emptyset(G) + m$.

Intuitively, the pebbings that seem most natural and interesting are persistent pebbings of DAGs with a single sink. In our proofs, however, we will mostly be focusing on visiting pebbings. The reason that visiting pebbings will show up over and over again is that the graphs of interest will often be constructed in terms of smaller subgraph components with useful pebbling properties, and that for such subgraphs we have the following fact.

Observation D.2. *Suppose that G is a DAG and that \mathcal{P} is any complete pebbling of G . Let $U \subseteq V(G)$ be any subset of vertices of G and let $H = H(G, U)$ denote the induced subgraph with vertices $V(H) = U$ and edges $E(H) = \{(u, v) \in E(G) \mid u, v \in U\}$. Then the pebbling \mathcal{P} restricted to the vertices in U is a complete visiting pebbling strategy for H .*

Proof. It is easy to verify that if we only perform those pebbling moves in \mathcal{P} that pertain to vertices in U , then these moves constitute a legal pebbling on H in the sense of Definition A.7. Moreover, any complete pebbling of G must pebble all vertices in G , so \mathcal{P} restricted to U will pebble all vertices in H including the sinks of H . \square

To get trade-offs in resolution for minimally unsatisfiable k -CNF formulas, we want DAGs with unique sinks. Most pebbling results in Section D are more natural to state and prove for DAGs with multiple sinks, however, but this small technicality is easily taken care of. We do this next.




Figure 1: Schematic illustration of single-sink version \widehat{G} of graph G .

Definition D.3 (Single-sink version). Let G be a DAG with sinks $Z(G) = \{z_1, \dots, z_m\}$ for $m > 1$. The *single-sink version* \widehat{G} of G consists of all vertices and edges in G plus the extra vertices z_1^*, \dots, z_{m-1}^* and the edges $(z_1, z_1^*), (z_2, z_1^*), (z_1^*, z_2^*), (z_3, z_2^*), (z_2^*, z_3^*), (z_4, z_3^*),$ et cetera up to $(z_{m-2}^*, z_{m-1}^*), (z_m, z_{m-1}^*)$.

That is, \widehat{G} consists of G with a binary tree of minimal size added on top of the sinks. See Figure 1 for a small example. The following observation is immediate.

Observation D.4. Let G be a DAG with sinks $Z(G) = \{z_1, \dots, z_m\}$ for $m > 1$. Then for any flavour of pebbling (visiting or persistent) it holds that $BW\text{-Peb}(\widehat{G}) \leq BW\text{-Peb}(G) + 1$ and $\text{Peb}(\widehat{G}) \leq \text{Peb}(G) + 1$. Moreover, if there is a pebbling strategy \mathcal{P} (visiting or persistent) for G in space s that can pebble the sinks in arbitrary order, then there is a pebbling strategy $\widehat{\mathcal{P}}$ of the same type for \widehat{G} with $\text{time}(\widehat{\mathcal{P}}) \leq \text{time}(\mathcal{P}) + 2m$ and $\text{space}(\widehat{\mathcal{P}}) \leq \text{space}(\mathcal{P}) + 1$.

To simplify the proofs of our lower bounds, we want that the pebblings under consideration do not perform any obviously redundant moves. The following definition is a generalization of [GLT80] from black-only to black-white pebbling.

Definition D.5 (Frugal pebbling). Let \mathcal{P} be a complete pebbling of a DAG G . To every pebble placement on a vertex v at time σ we associate the *pebbling interval* $[\sigma, \tau)$, where τ is the first time after σ when the pebble is removed from v again (or $\tau = \infty$, say, if this never happens).

If a sink $z_i \in Z(G)$ is pebbled for the first time at time σ , then the pebble on z_i is *essential* during the pebbling interval $[\sigma, \tau)$. A pebble on a non-sink vertex v is essential during $[\sigma, \tau)$ if either an essential black pebble is placed on an immediate successor of v during (σ, τ) or an essential white pebble is removed from an immediate successor of v during (σ, τ) . Any other pebble placements on any vertices are non-essential.

The pebbling strategy \mathcal{P} is *frugal* if all pebbles in \mathcal{P} are essential at all times.

Without loss of generality, we can assume that all pebblings we deal with are frugal.

Lemma D.6. For any complete pebbling \mathcal{P} (black or black-white, visiting or persistent) there is a frugal pebbling \mathcal{P}' of the same type such that $\text{time}(\mathcal{P}') \leq \text{time}(\mathcal{P})$ and $\text{space}(\mathcal{P}') \leq \text{space}(\mathcal{P})$.

Proof. Just delete any non-essential pebbles from \mathcal{P} and verify that what remains is a legal pebbling. \square

One minor technical snag is that we will need to assume not only that complete pebblings are frugal, but that this also holds for *conditional pebblings* (Definition D.1). This is no real problem, however, since we can always assume that the conditional pebblings we are dealing with are subpebblings of some larger, unconditional pebbling. More formally, we can define a conditional pebbling to be frugal if it satisfies the condition in Definition D.5 that any pebble placed on a non-sink vertex v stays until either a black pebble is placed on an immediate successor of v or a white pebble is removed from an immediate successor of v .

D.1.2 Some Upper and Lower Bounds

If we do not care about space, the easiest way to pebble a DAG is to place black pebbles on the vertices in topological order (and then remove all pebbles from non-sink vertices). Since we will have reason to use this pebbling strategy on occasion in what follows, we give it a name for reference.

Observation D.7 (Trivial pebbling). *Any DAG G can be completely, persistently black-pebbled in space at most $|V(G)|$ and time at most $2 \cdot |V(G)|$ simultaneously.*

Another easy upper bound on the black pebbling price can be given in terms of the fan-in and depth of the DAG.

Definition D.8 (Depth). The *depth* of a DAG G is the length of a longest path from a source to a sink in G .

Observation D.9. *Any DAG G with maximal indegree ℓ and depth d has a black pebbling strategy in space at most $d\ell + 1$.*

Proof. By induction over the depth. The base case is immediate. For a graph of depth $d + 1$, pebble the sinks one by one. For each sink we can pebble its immediate predecessors with $d\ell + 1$ pebbles each by induction. Placing black pebbles on the immediate predecessors one by one and leaving them there, we never use more than $(d\ell + 1) + (\ell - 1)$ pebbles simultaneously. Finally, keeping the at most ℓ pebbles on the predecessors, pebble the sink. \square

A simple but important lemma that lies at the heart of most black-white pebbling lower bounds follows next.

Lemma D.10 ([GT78]). *Suppose that $Q : u \rightsquigarrow v$ is a path in a DAG G and that $\mathcal{P} = \{\mathbb{P}_\sigma, \mathbb{P}_{\sigma+1}, \dots, \mathbb{P}_\tau\}$ is a pebbling such that the whole path Q is completely free of pebbles at times σ and τ but the endpoint v is pebbled at some point in the time interval (σ, τ) . Then the starting point u is pebbled during (σ, τ) as well.*

Proof. By induction over the length of the path Q . \square

A common graph in many pebbling constructions is the *pyramid*.

Definition D.11 (Pyramid graph). The *pyramid graph* Π_h of height $h \geq 1$ is a layered DAG with $h + 1$ levels, where there is one vertex on the highest level (the sink z), two vertices on the next level et cetera down to $h + 1$ vertices at the lowest level 0. The i th vertex at level L has incoming edges from the i th and $(i + 1)$ st vertices at level $L - 1$.

See Figure 2 for an example pyramid. The pebbling price of pyramids is well understood.

Theorem D.12. *The black pebbling price of a pyramid of height h is $\text{Peb}(\Pi_h) = h + 2$ and there is a linear-time pebbling achieving this bound. The black-white pebbling price is $\text{BW-Peb}^0(\Pi_h) = h/2 + O(1)$, and for even height there is the exact bound $\text{BW-Peb}^0(\Pi_{2h}) = h + 2$.*




Figure 2: Pyramid Π_6 of height 6.

Proof sketch. The lower bound for black pebbling is from Cook [Coo74], and it is easy to construct a linear-time pebbling matching this bound by pebbling the pyramid bottom-up, layer by layer.

The black-white pebbling strategy for pyramids in space $h/2 + O(1)$ can be obtained from the strategy for trees in Lengauer and Tarjan [LT80], and Klawe [Kla85] showed that $h/2 + O(1)$ is also a lower bound. The exact bound for pyramids of even height can be read off the exposition of Klawe’s proof in [NH08a]. \square

Another important building block in many pebbling results are so-called superconcentrators.

Definition D.13 (Superconcentrator). A directed acyclic graph G is an N -superconcentrator if it has N sources $S = \{s_1, \dots, s_N\}$, N sinks $Z = \{z_1, \dots, z_N\}$, and for any subsets S' and Z' of sources and sinks of size $|S'| = |Z'| = k$ it holds that there are k vertex-disjoint paths between S' and Z' in G .

Note that we do not assume that we can specify which source in S' should be connected to which sink in Z' .

For our pebbling purposes, we will be interested in superconcentrators with number of vertices and edges linear in N (in addition, we will want them to have bounded indegree, but this extra requirement is easy to take care of). Valiant [Val76] proved the existence of such graphs. Gabber and Galil [GG81] provided the first explicit construction of linear-size superconcentrators, based on an earlier non-explicit construction by Pippenger [Pip77]. We remark that the superconcentrators in [GG81] have logarithmic depth. The currently best known construction (i.e., with lowest edges-to-vertices ratio) that we are aware of is due to Alon and Capalbo [AC03].

Here is an important lemma that explains why superconcentrators are good building blocks if we want to construct graphs that are hard to pebble.

Lemma D.14 (Basic Lower Bound Argument for superconcentrators ([LT82])). *Suppose that the DAG G is an N -superconcentrator and that $\mathcal{P} = \{\mathbb{P}_\sigma, \mathbb{P}_{\sigma+1}, \dots, \mathbb{P}_\tau\}$ is a conditional black-white pebbling such that $\text{space}(\mathbb{P}_\sigma) \leq s_\sigma$, $\text{space}(\mathbb{P}_\tau) \leq s_\tau$, and \mathcal{P} pebbles at least $s_\sigma + s_\tau + 1$ sinks during the closed time interval $[\sigma, \tau]$. Then \mathcal{P} pebbles and unpebbles at least $N - s_\sigma - s_\tau$ different sources during the open time interval (σ, τ) .*

Proof. Suppose not. Then \mathcal{P} pebbles some set of $s_\sigma + s_\tau + 1$ sinks without pebbling some set of $s_\sigma + s_\tau + 1$ sources. Fix such sets of sources and sink vertices and consider the vertex-disjoint paths from sources to sinks. Then at least one path is empty both at time σ and at time τ and the end point of the path is pebbled during the interval (σ, τ) but not the starting point. This contradicts Lemma D.10. \square

Following Lengauer and Tarjan, we will refer to this as the *Basic Lower Bound Argument* lemma, or just BLBA-lemma, for superconcentrators. We immediately get the following corollary.




Figure 3: Base case Γ_1^3 for Carlson-Savage graph with 3 spines and sinks.

Corollary D.15 ([LT82]). *Any complete black-white pebbling of an N -superconcentrator in space at most s has to pebble at least $\Omega(N^2/s)$ sources (so, in particular, this is a lower bound on the pebbling time).*

D.2 A New Pebbling Trade-off Result

In this section we present the third main contribution of this paper, which is a graph family that provides us with a number of interesting time-space trade-offs for different parameter settings. These trade-offs have the property that the lower bounds are in terms of black-white pebbling while the upper bounds are in terms of black-only pebbling, and thanks to this we can apply the machinery of Theorem 3.2 on page 7 and Theorem C.4 on page 27 on these graphs to derive corresponding trade-offs in proof complexity for resolution.

D.2.1 Definition of Graph Family and Statement of Trade-off

Our graph family is built on a construction by Carlson and Savage [CS80, CS82]. Carlson and Savage only prove their trade-off for black pebbling, however, and in order to get results for black-white pebbling we have to modify the construction somewhat and also apply some new ideas in the proofs.

The next definition will hopefully be easier to parse if the reader first studies the illustrations in Figures 3 and 4.

Definition D.16 (Carlson-Savage graph). We define a two-parameter graph family Γ_r^c , for $c, r \in \mathbb{N}^+$, by induction over r . The base case Γ_1^c is a DAG consisting of two sources s_1, s_2 and c sinks $\gamma_1, \dots, \gamma_c$ with directed edges (s_i, γ_j) , for $i = 1, 2$ and $j = 1, \dots, c$, i.e., edges from both sources to all sinks. The graph Γ_{r+1}^c is a DAG with c sinks which is built from the following components:

- c disjoint copies $\Pi_{2r}^1, \dots, \Pi_{2r}^c$ of a pyramid (Definition D.11) of height $2r$, where we let z_1, \dots, z_c denote the pyramid sinks.
- one copy of Γ_r^c , for which we denote the sinks by $\gamma_1, \dots, \gamma_c$.
- c disjoint and identical *spines*, where each spine is composed of cr *sections*, and every section contains $2c$ vertices. We let the vertices in the i th section of a spine be denoted $v[i]_1, \dots, v[i]_{2c}$.

The edges in Γ_{r+1}^c are as follows:

- All “internal edges” in $\Pi_{2r}^1, \dots, \Pi_{2r}^c$ and Γ_r^c are present also in Γ_{r+1}^c .
- For each spine, there are edges $(v[i]_j, v[i]_{j+1})$ for all $j = 1, \dots, 2c - 1$ within each section i and edges $(v[i]_{2c}, v[i+1]_1)$ from the end of a section to the beginning of next for $i = 1, \dots, cr - 1$, i.e., for all sections but the final one, where $v[cr]_{2c}$ is a sink.
- For each section i in each spine, there are edges $(z_j, v[i]_j)$ from the j th pyramid sink to the j th vertex in the section for $j = 1, \dots, c$, as well as edges $(\gamma_j, v[i]_{c+j})$ from the j th sink in Γ_r^c to the $(c + j)$ th vertex in the section for $j = 1, \dots, c$.

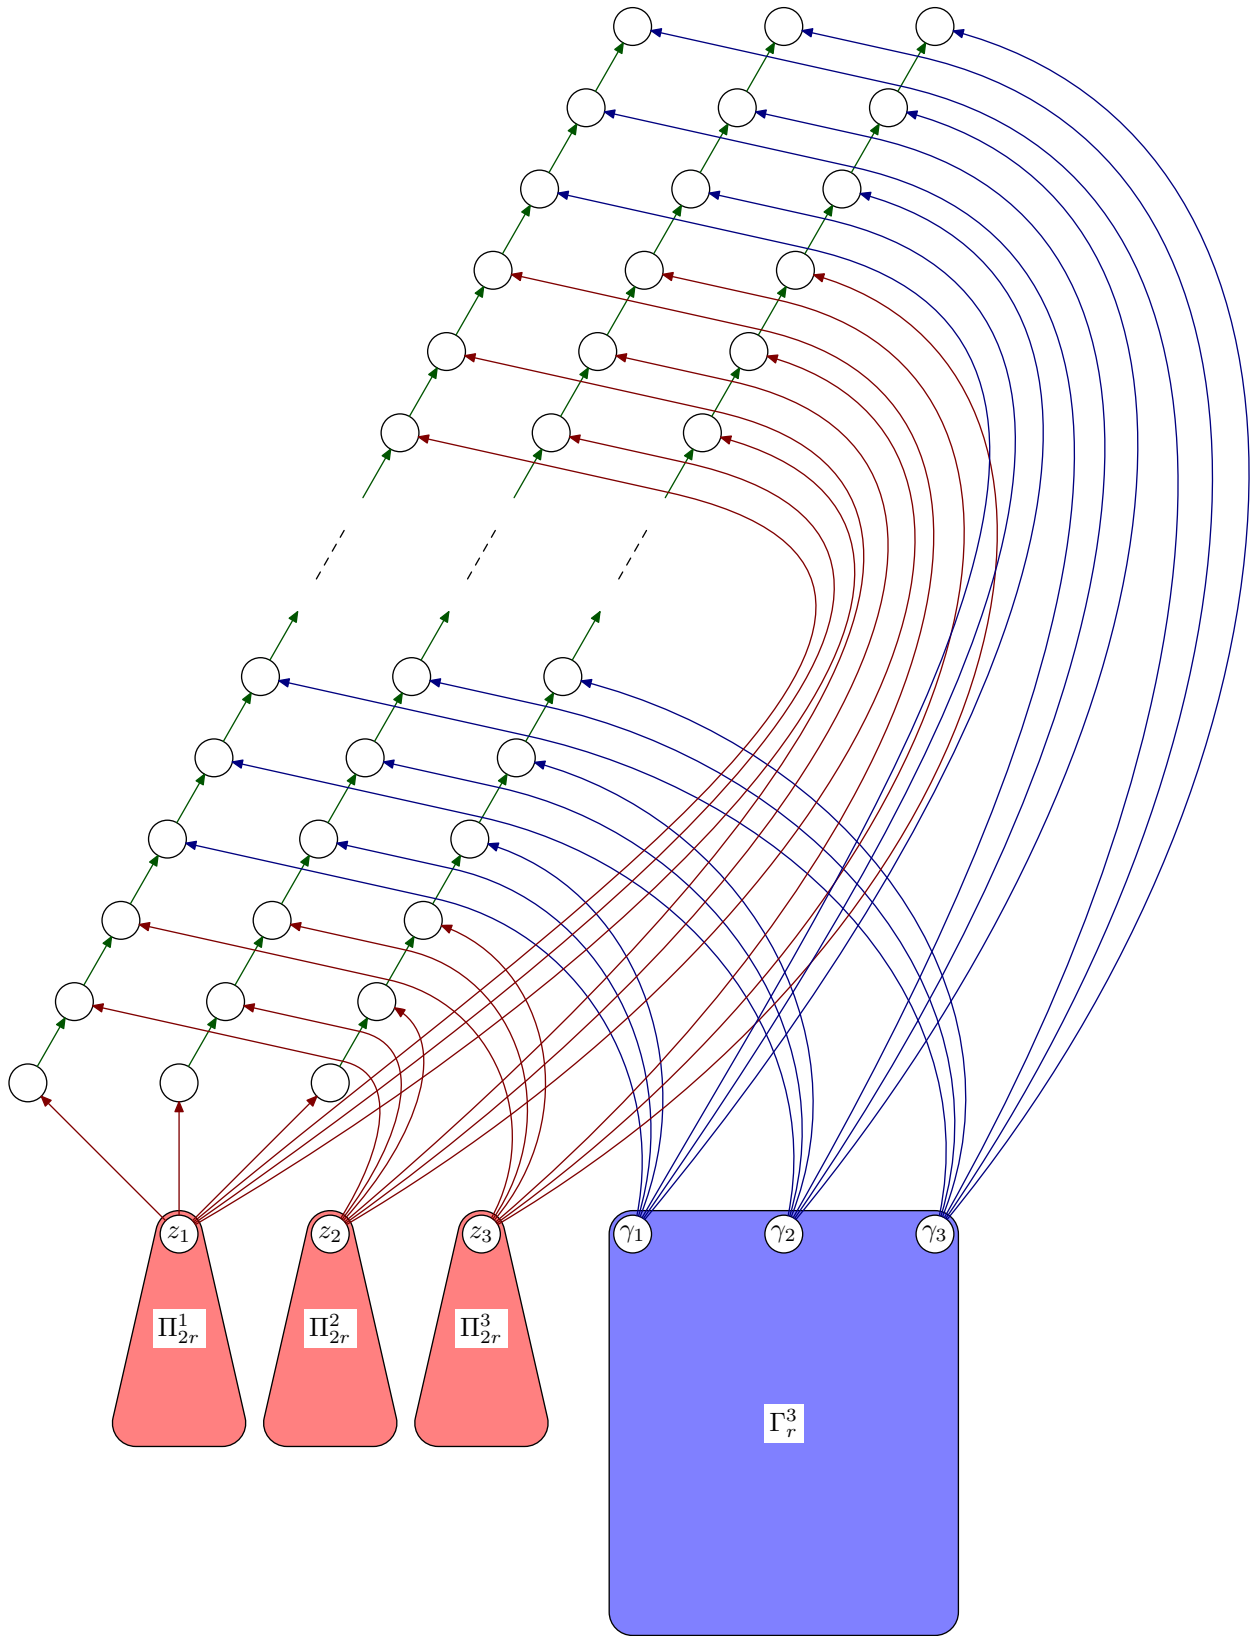


Figure 4: Inductive definition of Carlson-Savage graph Γ_{r+1}^3 with 3 spines and sinks.

We now make the formal statements of the trade-off properties that these DAGs possess. The proofs of all the statements are postponed to Section D.2.2. First, we collect some basic properties.

Lemma D.17. *The graphs Γ_r^c are of size $|V(\Gamma_r^c)| = \Theta(cr^3 + c^3r^2)$, and have black-white pebbling price $BW\text{-Peb}^\emptyset(\Gamma_r^c) = r + 2$ and black pebbling price $\text{Peb}^\emptyset(\Gamma_r^c) = 2r + 1$.*

This tells us that the minimum pebbling space required grows linearly with the recursion depth r but is independent of the number of spines c of the DAG.

Next, we show that there is a linear-time completely black pebbling of Γ_r^c in space linear in the sum of the parameters. This is in fact a strict improvement (though easily obtained) of the corresponding result in [CS80, CS82].

Lemma D.18. *The graph Γ_r^c has a persistent black pebbling strategy \mathcal{P} in time linear in the size of the DAG and with space $O(c + r)$.*

The proof is by induction, and the idea in the induction step for Γ_{r+1}^c is to make a persistent pebbling of Γ_r^c in space $O(c + r)$, then pebble the pyramids $\Pi_{2r}^1, \dots, \Pi_{2r}^c$ one by one in linear time and space $O(r)$, and finally, using the $2c$ black pebbles on $z_1, \dots, z_c, \gamma_1, \dots, \gamma_c$ that we have left in place, to pebble all c spines in parallel with $O(c)$ extra pebbles.

The main result of this section is the following theorem, which allows us to get a variety of pebbling trade-off results if we choose the parameters c and r appropriately.

Theorem D.19. *Suppose that \mathcal{P} is a complete visiting black-white pebbling of Γ_r^c with*

$$\text{space}(\mathcal{P}) < BW\text{-Peb}^\emptyset(\Gamma_r^c) + s = (r + 2) + s$$

for $0 < s \leq c/8 - 1$. Then the time required to perform \mathcal{P} is lower-bounded by

$$\text{time}(\mathcal{P}) \geq \left(\frac{c - 2s}{4s + 4} \right)^r \cdot r! .$$

As has already been noted, we defer the proof of Theorem D.19 to Section D.2.2, but let us nevertheless try to provide some intuition as to why the theorem should be true.

For simplicity, let us focus on black-only pebbling strategies. Inductively, suppose that the trade-off in Theorem D.19 has been proven for Γ_r^c and consider Γ_{r+1}^c . Any pebbling strategy for this DAG will have to pebble through all sections in all spines. Consider the first section anywhere, let us say on spine j , that has been completely pebbled, i.e., there have been pebbles placed on and removed from all vertices in the section. Let us say that this happens at time τ_1 . But this means that Γ_r^c and all pyramids $\Pi_{2r}^1, \dots, \Pi_{2r}^c$ must have been completely pebbled during this part of the pebbling as well. Fix any pyramid and consider some point in time $\sigma_1 < \tau_1$ when the number of pebbles in this pyramid reaches the space $r + O(1)$ required by the known lower bound on pyramid pebbling price. At this point, the rest of the graph must contain very few pebbles. In particular, there are very few pebbles on the subgraph Γ_r^c at time σ_1 , so we can think of Γ_r^c as being completely empty of pebbles for all practical purposes.

Let us now shift the focus to the next section in the spine j that is completed, say, at time $\tau_2 > \tau_1$. Again, we can argue that some pyramid is completely pebbled in the time interval $[\tau_1, \tau_2]$, and thus has $r + O(1)$ pebbles on it at some time $\sigma_2 > \tau_1 > \sigma_1$. This means that we can think of Γ_r^c as being completely empty at time σ_2 as well.

But note that all sinks in the subgraph Γ_r^c must have been pebbled in the time interval $[\sigma_1, \sigma_2]$, and since we know that Γ_r^c is (almost) empty at times σ_1 and σ_2 , this allows us to apply the induction hypothesis. Since \mathcal{P} has to pebble through a lot of sections in different spines, we will be able to repeat the above

argument many times and apply the induction hypothesis on Γ_r^c in each round. Adding up all the lower bounds obtained in this way, we see that the induction step goes through.

This is essentially the proof in [CS80, CS82] for black pebbling, modulo a number of technical details that we glossed over. For black-white pebbling, these technical complications grow more serious. The main problem is that in contrast to a black pebbling, that has to proceed through the DAG in some kind of bottom-up fashion, a black-white pebbling can place and remove pebbles anywhere in the DAG at any time. Therefore, it is more difficult to control the progress of a black-white pebbling, and we have to work harder in the proof of our theorem.

Also, it should be noted that the added complications when going from black to black-white pebbling result in our bounds for black-white pebbling being slightly worse than the ones in [CS80, CS82] for black pebbling only. More specifically, Carlson and Savage are able to prove their results for DAGs having only $\Theta(r)$ sections per spine, whereas we need $\Theta(cr)$ sections. This blows up the number of vertices, which in turn weakens the trade-offs measured in terms of graph size.

It would be interesting to find out whether our proof, presented below, could in fact be made to work for graphs with only $O(r)$ sections per spine. If so, this would immediately improve all the trade-off results for resolution in Appendix E that we obtain based on the graphs in Definition D.16.

D.2.2 Proofs of Lemma D.17, Lemma D.18, and Theorem D.19

Before proving the results claimed in Section D.2.1, we establish a couple of useful auxiliary lemmas. The first lemma below gives us information about how the spines in Γ_r^c are pebbled. We will use this information repeatedly in what follows.

Lemma D.20. *Suppose that G is a DAG and that v is a vertex in G with a path Q to some sink $z_i \in Z(G)$ such that all vertices in $Q \setminus \{z_i\}$ have outdegree 1. Then any frugal black-white pebbling strategy pebbles v exactly once, and the path Q contains pebbles during one contiguous time interval.*

Proof. By induction from the sink backwards. The induction base is immediate. For the inductive step, suppose v has immediate successor w and that w is pebbled exactly once.

If w is black-pebbled at time σ , then v has been pebbled before this and the first pebble placed on v stays until time σ . No second placement of a pebble on v after time σ can be essential since v has no other immediate successor than w . If w is white-pebbled and the pebble is removed at time σ , then the first pebble placed on v stays until time σ and no second placement of a pebble on v after time σ can be essential.

Thus each vertex on the path is pebbled exactly once, and the time intervals when a vertex v and its successor w have pebbles on them overlap. The lemma follows. \square

The second lemma speaks about subgraphs H of a DAG G whose only connection to the rest of the DAG $G \setminus H$ are via the sink of H . Note that the pyramids in Γ_r^c satisfy this condition.

Lemma D.21. *Let G be a DAG and H a subgraph in G such that H has a unique sink z_h and the only edges between $V(H)$ and $V(G) \setminus V(H)$ emanate from z_h . Suppose that \mathcal{P} is any frugal complete pebbling of G having the property that H is completely empty of pebbles at some given time τ' but at least one vertex of H has been pebbled during the time interval $[0, \tau']$. Then \mathcal{P} pebbles H completely during the interval $[0, \tau']$.*

Proof. Suppose that $v \in V(H)$ is pebbled at time $\sigma' < \tau'$. As in the proof of Lemma D.10, we can argue by induction over the length of the longest path from v to the sink z_h of H that z_h must also be pebbled before time τ' . Note that such a path exists since the sink z_h is unique, and that any path starting in v must hit z_h sooner or later, since this vertex is the only way out of H into the rest of G . Since H is empty at times 0 and τ' , we conclude that \mathcal{P} makes a complete visiting pebbling of H during $[0, \tau']$. \square

Let us now establish that the size and pebbling price of Γ_r^c are as claimed.

Proof of Lemma D.17. The base case Γ_1^c for the Carlson-Savage graph in Definition D.16 has size $c + 2$. A pyramid of height h has $(h + 1)(h + 2)/2$ vertices, so the c pyramids of height $2(r - 1)$ in Γ_r^c contribute $cr(2r - 1)$ vertices. The c spines with cr sections of $2c$ vertices each contribute a total of $2c^3r$ vertices. And then there are the vertices in Γ_{r-1}^c . Summing up, the total number of vertices in Γ_r^c is

$$(c + 2) + \sum_{i=2}^r (ci(2i - 1) + 2c^3i) = \Theta(cr^3 + c^3r^2) \quad (21)$$

as is stated in the lemma.

Clearly, $BW-Peb^\emptyset(\Gamma_1^c) = Peb^\emptyset(\Gamma_1^c) = 3$, since pebbling a vertex with fan-in 2 requires 3 pebbles and Γ_1^c can be completely pebbled in this way by placing pebbles on the two sources and then pebble and unpebble the sinks one by one.

Suppose inductively that $BW-Peb^\emptyset(\Gamma_r^c) = r + 2$ and consider Γ_{r+1}^c . It is straightforward to see that $BW-Peb^\emptyset(\Gamma_{r+1}^c) \leq r + 3$. Every pyramid Π_{2r}^j can be completely pebbled with $r + 2$ pebbles (Theorem D.12). We can pebble each spine bottom-up in the following way, section by section. Suppose by induction that we have a black pebble on the last vertex $v[i - 1]_{2c}$ in the $(i - 1)$ st section. Keeping the pebble on $v[i - 1]_{2c}$, perform a complete visiting pebbling of Π_{2r}^1 . At some point during this pebbling we must have a pebble on the pyramid sink z_1 and at most r other pebbles on the pyramid (simply because without loss of generality some pebble placement on z_1 must be followed by a removal or placement of a pebble on some other vertex). At this time, place a black pebble on $v[i]_1$ and remove the pebble from $v[i - 1]_{2c}$. Complete the pebbling of Π_{2r}^1 , leaving the pyramid empty. Performing complete visiting pebbings of $\Pi_{2r}^2, \dots, \Pi_{2r}^c$ in the same way allows us to move the black pebble along $v[i]_2, \dots, v[i]_c$, never exceeding total pebbling space $r + 3$. It is easy to see that in the same way, for every visiting pebbling \mathcal{P} of Γ_r^c there must exist times σ_i for all $i = 1, \dots, c$, when $space(\mathbb{P}_{\sigma_i}) < space(\mathcal{P})$ and the sink γ_i contains a pebble. Performing a minimum-space pebbling of Γ_r^c , possibly c times if necessary, this allows us to advance the black pebble along $v[i]_{c+1}, \dots, v[i]_{2c}$, never exceeding total pebbling space $r + 3$. This shows that Γ_{r+1}^c can be completely pebbled with $r + 3$ pebbles. A simple pattern-matching of this argument for black pebbling (appealing to Theorem D.12 for the black pebbling price of pyramids) also yields $Peb^\emptyset(\Gamma_r^c) \leq 2r + 3$.

To prove that there are matching lower bounds for the pebbling constructed above, it is sufficient to show that some pyramid Π_{2r}^j must be completely pebbled while there is at least one pebble on Γ_{r+1}^c outside of Π_{2r}^j . To see why, note that if we can prove this, then simply by using the fact that $BW-Peb^\emptyset(\Pi_{2r}) = r + 2$ and $BW-Peb^\emptyset(\Pi_{2r}^j) = 2r + 2$ and adding an additive constant 1 for the pebble outside of Π_{2r}^j we have the matching lower bounds that we need. We present the argument for black-white pebbling, which is the harder case. The black-only pebbling case is handled completely analogously.

Suppose in order to get a contradiction that there is a complete visiting pebbling strategy \mathcal{P} for Γ_{r+1}^c in space $r + 2$. By Observation D.2, \mathcal{P} performs a complete visiting pebbling of every pyramid Π_{2r}^j . Consider the first time τ_1 when some pyramid has been completely pebbled and let this pyramid be $\Pi_{2r}^{j_1}$. Then at some time $\sigma_1 < \tau_1$ there are $r + 2$ pebbles on $\Pi_{2r}^{j_1}$ and the rest of the graph Γ_{r+1}^c must be empty of pebbles at this point.

We claim that this implies that no vertex in Γ_{r+1}^c outside of the pyramid $\Pi_{2r}^{j_1}$ has been pebbled before time σ_1 . Let us prove this crucial fact by a case analysis.

1. No vertex v in any other pyramid $\Pi_{2r}^{j'}$ can have been pebbled before time σ_1 . For if so, Lemma D.21 says that $\Pi_{2r}^{j'}$ has been completely pebbled before time σ_1 , contradicting that $\Pi_{2r}^{j_1}$ is the first pyramid completely pebbled by \mathcal{P} .

2. No vertex on a spine has been pebbled before time σ_1 . This is so since Lemma D.20 tells us that if some vertex on a spine has been pebbled, then the whole spine must have been pebbled in view of the fact that it is empty at time σ_1 . But then Lemma D.10 implies that all pyramid sinks must have been pebbled. This case has already been excluded.
3. Finally, no vertex v in Γ_r^c can have been pebbled before time σ_1 . Otherwise the frugality of \mathcal{P} implies (by pattern matching on the arguments in the proofs of Lemmas D.10 and D.20) that some successor of v must have been pebbled as well, and some successor of that successor et cetera, all the way up to where Γ_r^c connects with the spines. But we have ruled out the possibility that a spine vertex has been pebbled.

This establishes the claim, and we are now almost done. Before clinching the argument, we need to make a couple of observations. Note first that by frugality, we can conclude that at some time in the interval (σ_1, τ_1) some vertex in some spine must be pebbled. This is so since the pyramid sink z_{j_1} has been pebbled before time τ_1 and all of $\Pi_{2r}^{j_1}$ is empty at time τ_1 but all spines are empty at time $\sigma_1 < \tau_1$. But then Lemma D.20 tells us that there will remain a pebble on this spine until all of the spine has been completely pebbled.

Consider now the second pyramid $\Pi_{2r}^{j_2}$ completely pebbled by \mathcal{P} , say, at time τ_2 . At some point in time $\sigma_2 < \tau_2$ we have $r + 2$ pebbles on $\Pi_{2r}^{j_2}$, and moreover $\sigma_2 > \tau_1$ since $\Pi_{2r}^{j_2}$ is empty at time τ_1 . But now it must hold that either there is a pebble on a spine at this point, or, if all spines are completely empty, that some spine has been completely pebbled. If some spine has been completely pebbled, however, this in turn implies (appealing to Lemma D.10 again) that there must be some pebble somewhere in some other pyramid $\Pi_{2r}^{j'}$ at time σ_2 . Thus the pebbling space exceeds $r + 2$ and we have obtained our contradiction. The lemma follows. \square

Studying the pebbling strategies in the proof of Lemma D.17, it is not hard to see that they are terribly inefficient. The subgraphs in Γ_r^c will be pebbled over and over again, and for every step in the recursion the time required multiples. We next show that if we are just a bit more generous with the pebbling space, then we can get down to linear time.

Proof of Lemma D.18. We want to prove that Γ_r^c has a persistent black pebbling strategy \mathcal{P} in linear time and in space $O(c + r)$. Suppose that there is such a pebbling strategy \mathcal{P}_r for Γ_r^c . We show how to construct a pebbling \mathcal{P}_{r+1} for Γ_{r+1}^c inductively. Note that the base case for Γ_1^c is trivial.

The construction of \mathcal{P}_{r+1} is very straightforward. First use \mathcal{P}_r to make a persistent pebbling of Γ_r^c in space $O(c + r)$. At the end of \mathcal{P}_r , we have c pebbles on the sinks $\gamma_1, \dots, \gamma_c$. Keeping these pebbles in place, pebble the pyramids $\Pi_{2r}^1, \dots, \Pi_{2r}^c$ persistently one by one in linear time and space $O(r)$. We leave pebbles on all pyramid sinks z_1, \dots, z_c . This stage of the pebbling only requires space $O(c + r)$ and at the end we have $2c$ black pebbles on all pyramid sinks z_1, \dots, z_c and all sinks $\gamma_1, \dots, \gamma_c$ of Γ_r^c . Keeping all these pebbles in place, we can pebble all c spines in parallel in linear time with $c + 1$ extra pebbles. \square

It remains to prove the trade-off result in Theorem D.19. It is clear that this theorem follows if we can prove the next, slightly stronger, statement.

Lemma D.22. *Suppose that $\mathcal{P} = \{\mathbb{P}_\sigma, \dots, \mathbb{P}_\tau\}$ is a conditional black-white pebbling on Γ_r^c and that s is a constant satisfying the following properties:*

1. $0 < s \leq c/8 - 1$.
2. \mathcal{P} pebbles all sinks in Γ_r^c during the time interval $[\sigma, \tau]$.
3. $\max\{\text{space}(\mathbb{P}_\sigma), \text{space}(\mathbb{P}_\tau)\} < s$.

$$4. \text{space}(\mathcal{P}) < \text{BW-Peb}^\emptyset(\Gamma_r^c) + s = (r + 2) + s.$$

Then it holds that $\text{time}(\mathcal{P}) = \tau - \sigma \geq \left(\frac{c-2s}{4s+4}\right)^r \cdot r!$.

We will have to spend some time working on this lemma before the proof is complete, though. We start by establishing two additional auxiliary lemmas that upper-bound how many pyramids and spine sections can contain pebbles simultaneously at any one given time in a pebbling subjected to space constraints as in Lemma D.22. The claims in the two lemmas are very similar in spirit, as are the proofs, so we state the lemmas together and then present the proofs together.

Lemma D.23. *Suppose that $\mathcal{P} = \{\mathbb{P}_\sigma, \dots, \mathbb{P}_\tau\}$ is a conditional black-white pebbling on Γ_r^c and that s is a constant satisfying the conditions in Lemma D.22. Then at all times during the pebbling \mathcal{P} strictly less than $4(s + 1)$ pyramids Π_{2r}^j contain pebbles simultaneously.*

Lemma D.24. *Suppose that $\mathcal{P} = \{\mathbb{P}_\sigma, \dots, \mathbb{P}_\tau\}$ is a conditional black-white pebbling on Γ_r^c and that s is a constant satisfying the conditions in Lemma D.22. Then at all times during the pebbling \mathcal{P} strictly less than $4(s + 1)$ spine sections contain pebbles simultaneously.*

Note that Lemma D.24 provides a total bound on the number of pebbled sections in all c spines. There might be spines containing several sections being pebbled simultaneously (in fact, this is exactly what one would expect a black-white pebbling to do to optimize the time given the space constraints), but what Lemma D.24 says that if we fix an arbitrary time $t \in [\sigma, \tau]$, add up the number of sections containing pebbles at time t in each spine, and sum over all spines, we never exceed $4(s + 1)$ sections in total at any point in time $t \in [\sigma, \tau]$.

Proof of Lemma D.23. Suppose that on the contrary, there is some time $t^* \in (\sigma, \tau)$ when at least $4s + 4$ pyramids Π^j in Γ_r^c contain pebbles. Of these pyramids, at least $2s + 4$ are empty both at time σ and at time τ since $\text{space}(\mathbb{P}_\sigma) < s$ and $\text{space}(\mathbb{P}_\tau) < s$. By Lemma D.21, these pyramids, which we denote Π^1, \dots, Π^{2s+4} , are completely pebbled. We conclude that for every Π^j , $j = 1, \dots, 2s + 4$, there is an interval $[\sigma_j, \tau_j]$ such that $t^* \in (\sigma_j, \tau_j)$ and Π^j is empty at times σ_j and τ_j but contains pebbles throughout the interval (σ_j, τ_j) during which it is completely pebbled.

For each Π^j there must exist some time $t_j^* \in (\sigma_j, \tau_j)$ when there are at least $r + 1 = \text{BW-Peb}^\emptyset(\Pi^j)$ pebbles. Fix such a time t_j^* for every pyramid Π^j and assume that all t_j^* , $j = 1, \dots, 2s + 4$, are sorted in increasing order. We have two possible cases:

1. At least half of all t_j^* occur before (or at) time t^* , i.e., they satisfy $t_j^* \leq t^*$. If so, look at the largest $t_j^* \leq t^*$. At this time there are at least $r + 1$ pebbles on Π^j and at least $\frac{2s+4}{2} - 1 = s + 1$ pebbles on other pyramids, which means that $\text{space}(\mathbb{P}_{t_j^*}) \geq (r + 2) + s$. In other words, \mathcal{P} exceeds the space restrictions contradicting our assumptions.
2. At least half of all t_j^* occur after time t^* , i.e., they satisfy $t_j^* > t^*$. If we consider the smallest t_j^* larger than t^* we can again conclude that $\text{space}(\mathbb{P}_{t_j^*}) \geq (r + 1) + (s + 1)$, which is a contradiction.

Hence, if \mathcal{P} is a pebbling that complies with the restrictions in Lemma D.22, it can never be the case that $4s + 4$ pyramids Π^j in Γ_r^c contain pebbles simultaneously during \mathcal{P} . \square

Proof of Lemma D.24. Suppose in order to get a contradiction that at some time $t^* \in (\sigma, \tau)$ at least $4s + 4$ sections contain pebbles. At least $2s + 4$ of these sections are empty at times σ and τ . Let us denote these sections R^1, \dots, R^{2s+4} . Appealing to Lemma D.20, we conclude that there are intervals $[\sigma_j, \tau_j]$ for $j = 1, \dots, 2s + 4$, such that $t^* \in (\sigma_j, \tau_j)$ and R^j is empty at times σ_j and τ_j but contains pebbles throughout the interval (σ_j, τ_j) during which it is completely pebbled.

By Lemma D.23 we know that less than $4s + 4$ pyramids contain pebbles at time σ_j and similarly at time τ_j . Since all c pyramids in Γ_r^c must have their sinks pebbled during (σ_j, τ_j) by y but we have $2 \cdot (4s + 4) < c$ by the assumptions in Lemma D.22, we conclude from Lemma D.21 that for every interval (σ_j, τ_j) we can find some pyramid Π^j that is completely pebbled during this interval. This, in turn, implies that there is some time $t_j^* \in (\sigma_j, \tau_j)$ when the pyramid Π^j contains at least $BW\text{-Peb}^\emptyset(\Pi^j) = r + 1$ pebbles. (We note that many t_j^* can be equal and even refer to the same pyramid which has just happened to receive a lot of different labels, but this is not a problem as we shall see.)

As in the proof of Lemma D.23, we now sort the $t_j^*, j = 1, \dots, 2s + 4$, in increasing order and consider the two possible cases. If at least half of all t_j^* satisfy $t_j^* \leq t^*$, we look at the largest $t_j^* \leq t^*$. At this time there are at least $r + 1$ pebbles on Π^j and at least $\frac{2s+4}{2} = s + 2$ pebbles on different sections, which means that $\text{space}(\mathbb{P}_{t_j^*}) \geq rs + 3$ exceeds the space restrictions. If, on the other hand, at least half of all t_j^* satisfy $t_j^* > t^*$, then for the smallest t_j^* larger than t^* we can again conclude that $\text{space}(\mathbb{P}_{t_j^*}) \geq r + s + 3$, which is a contradiction. The lemma follows. \square

Putting together everything that has been proven so far in this section, we are able to establish the pebbling trade-off result.

Proof of Lemma D.22. Suppose that $\mathcal{P} = \{\mathbb{P}_\sigma, \dots, \mathbb{P}_\tau\}$ is a conditional black-white pebbling on Γ_r^c pebbling all sinks and that $\max\{\text{space}(\mathbb{P}_\sigma), \text{space}(\mathbb{P}_\tau)\} < s$ and $\text{space}(\mathcal{P}) < (r+2)+s$ for $0 < s \leq c/8 - 1$. Let us define

$$T(c, r, s) = \left(\frac{c - 2s}{4s + 4}\right)^r \cdot r! \quad . \quad (22)$$

We show that $\text{time}(\mathcal{P}) \geq T(c, r, s)$ by induction over r .

For $r = 1$, the assumptions in the lemma imply that more than $c - 2s$ sinks are empty at times σ and τ . These sinks must be pebbled, which trivially requires strictly more than $c - 2s > \left(\frac{c-2s}{4s+4}\right) = T(c, 1, s)$ time steps.

Assume that the lemma holds for Γ_{r-1}^c and consider any pebbling of Γ_r^c . Less than $2s$ spines contain pebbles at time σ or time τ . All the other strictly more than $c - 2s$ spines are empty at times σ and τ but must be completely pebbled during $[\sigma, \tau]$ by Lemma D.10.

Consider the first time σ' when any spine gets a pebble for the first time. Let us denote this spine by Q' . By Lemma D.20 we know that Q' contains pebbles during a contiguous time interval until it is completely pebbled and emptied at, say, time τ' . During this whole interval $[\sigma', \tau']$ less than $4s + 4$ sections contain pebbles at any one given time, so in particular less than $4s + 4$ spines contain pebbles. Moreover, Lemma D.20 says that every spine containing pebbles will remain pebbled until completed. What this means is that if we order the spines with respect to the time when they first receive a pebble in groups of size $4s + 4$, no spine in the second group can be pebbled until the at least one spine in the first group has been completed.

We remark that this divides the spines that are empty at the beginning and end of \mathcal{P} into strictly more than

$$\frac{c - 2s}{4s + 4} \quad (23)$$

groups. Furthermore, we claim that completely pebbling just one empty spine requires at least

$$r \cdot T(c, r - 1, s) \quad (24)$$

time steps. Given this claim we are done, since combining (23) and (24) we can deduce that the total pebbling time is lower-bounded by

$$\frac{c - 2s}{4s + 4} r \cdot T(c, r - 1, s) = T(c, r, s) \quad (25)$$

since at least one spine from each group is pebbled in a time interval totally disjoint from the time intervals for all spines in the next group.

It remains to establish the claim. To this end, fix any spine Q^* empty at times σ^* and τ^* but completely pebbled in $[\sigma^*, \tau^*]$. Consider the first time $\tau_1 \in [\sigma^*, \tau^*]$ when any section in Q^* , let us denote it by R_1 , has been completely pebbled (i.e., all vertices has been touched by pebbles but are now empty again). During $[\sigma^*, \tau_1]$ all pyramid sinks z_1, \dots, z_c are pebbled (Lemma D.10), and since less than $2 \cdot (4s+4) < c$ pyramids contain pebbles at times σ^* or τ_1 (Lemma D.23), at least one pyramid is pebbled completely (Lemma D.21), which requires $r+1$ pebbles. Moreover, there is at least one pebble on R_1 during this whole interval. Hence, there is a time $\sigma_1 \in [\sigma^*, \tau_1]$ when there are strictly less than $(r+2) + s - (r+1) - 1 = s$ pebbles on Γ_{r-1}^c . Also, at this time σ_1 less than $4s+4$ sections contain pebbles (Lemma D.24), and in particular this means that there are pebbles on less than $4s+3$ other section of our spine Q^* . This puts an upper bound on the number of sections of Q^* pebbled this far, since every section is completely pebbled during a contiguous time interval before being emptied again, and we chose to focus on the first section R_1 in Q^* that was finished.

Look now at the first section R_2 in Q^* other than the less than $4s+4$ sections containing pebbles at time σ_1 that is completely pebbled, and let the time when R_2 is finished be denoted τ_2 (clearly, $\tau_2 > \tau_1$). During $[\sigma_1, \tau_2]$ all sinks of Γ_{r-1}^c must have been pebbled, and at time $\tau_2 - 1$ less than $4s+3$ other section in Q^* contain pebbles.

Wrapping up, consider the first new section R_3 in our spine Q^* to be completely pebbled among those that has not yet been touched at time $\tau_2 - 1$. Suppose that R_3 is finished at time τ_3 . Then during $[\tau_2, \tau_3]$ some pyramid is completely pebbled, and thus there must exist a time $\sigma_3 \in (\tau_2, \tau_3)$ when there are at least $r+1$ pebbles on this pyramid and at least one pebble on the spine Q^* , leaving less than s pebbles for Γ_{r-1}^c . But this means that we can apply the induction hypothesis on the interval $[\sigma_1, \sigma_3]$ and deduce that $\sigma_3 - \sigma_1 \geq T(c, r-1, s)$. Note also that at time σ_3 less than $8s+8 < c$ sections in Q^* have been finished.

Continuing in this way, for every group of $8s+8 < c$ finished sections in Q^* we get one pebbling of Γ_{r-1}^c in space less than $BW\text{-Peb}^0(\Gamma_{r-1}^c) + s$ and with less than s pebbles in the start and end configurations, which allows us to apply the induction hypothesis a total number of at least $\frac{cr}{8s+8} > r$ times. (Just to argue that we get the constants right, note that $8s+8 < c$ implies that after the final pebbling of the sinks of Γ_{r-1}^c has been done, there is still some empty section left in Q^* . When this final section is taken care of, we will again get at least $r+1$ pebbles on some pyramid while at least one pebble resides on Q^* , so we get the space on Γ_{r-1}^c down below s as is needed for the induction hypothesis.)

This proves our claim that pebbling one spine takes time at least $r \cdot T(c, r-1, s)$. The lemma follows. \square

As we already noted, this completes the proof of Theorem D.19, since this theorem follows immediately from Lemma D.22 for the special case when $\mathbb{P}_\sigma = \mathbb{P}_\tau = (\emptyset, \emptyset)$.

D.3 Recapitulation of Some Known Pebbling Trade-off Results

All the material in Section D.3 is from [LT82]. The statements of the results below differ slightly in the constants in that paper, though, since there are some (minor) technical differences in the definitions as compared to the present paper.

D.3.1 Pebbling Trade-offs for Constant Space

This material is from [LT82, Section 2].

Definition D.25 (Permutation graph ([LT82])). Let π denote some permutation of $\{0, 1, \dots, n-1\}$. Then the *permutation graph* Δ_π^n over n elements with respect to the permutation π is defined as follows. Δ_π^n has $2n$ vertices divided into a *lower row* u_0, u_1, \dots, u_{n-1} and an *upper row* w_0, w_1, \dots, w_{n-1} . For all




Figure 5: Bit reversal graph Δ_{rev}^8 on 8 elements.




Figure 6: Bit reversal graph Δ_{rev}^{16} on 16 elements.

$i = 0, 1, \dots, n-2$, there are directed edges (u_i, u_{i+1}) and (w_i, w_{i+1}) , and for all $i = 0, 1, \dots, n-1$, there are edges $(u_i, w_{\pi(i)})$ from the lower row to the upper row.

The only source vertex in Δ_{π}^n is u_0 and w_{n-1} is the unique sink. The maximal indegree is 2.

Any DAG of fan-in 2 must have pebbling price at least 3. It is not too hard to see that permutation graphs Δ_{π}^n have pebbling strategies in this minimal space: keeping one pebble on vertex w_i of the upper row, move two pebbles consecutively on the lower row until $u_{\pi^{-1}(i+1)}$ is reached, and then pebble w_{i+1} . This pebbling requires quadratic time, however. The pebbling sketched above can be generalized to yield the following result.

Lemma D.26 ([LT82]). *Let Δ_{π}^n be the permutation graph over n elements for any permutation π . Then the black pebbling price of π is $\text{Peb}(\Delta_{\pi}^n) = 3$, and for any space $s \geq 3$ there is a black pebbling strategy \mathcal{P} for Δ_{π}^n with $\text{space}(\mathcal{P}) \leq s$ and $\text{time}(\mathcal{P}) \leq \frac{2n^2}{s-2} + 2n$.*

We will be particularly interested in permutations defined in terms of reversing the binary representation of the integers $\{0, 1, \dots, n-1\}$.

Definition D.27 (Bit reversal graph ([LT82])). The m -bit reversal of the non-negative integer i , $i < 2^m - 1$, is the non-negative integer $\text{rev}_m(i)$ obtained by writing the m -bit binary representation of i in reverse order. The bit reversal graph $\Delta_{\text{rev}_m}^{2^m}$ is the permutation graph over $n = 2^m$ with respect to the permutation rev_m .

For instance, we have $\text{rev}_3(1) = 4$, $\text{rev}_3(2) = 2$, and $\text{rev}_3(3) = 6$. We will denote the bit reversal graph by Δ_{rev}^n for simplicity, implicitly assuming that $n = 2^m$. Two examples of bit reversal graphs can be found in Figures 5 and 6.

For bit reversal graphs, the trade-off in Lemma D.26 for black pebbling is asymptotically tight.

Theorem D.28 ([LT82]). *Suppose that \mathcal{P} is any complete black pebbling of Δ_{rev}^n with $\text{space}(\mathcal{P}) = s$ for $s \geq 3$. Then $\text{time}(\mathcal{P}) \geq \frac{n^2}{8s}$.*

Note, in particular, that if we want to black-pebble Δ_{rev}^n in linear time, then linear space is needed. If we are also allowed to use white pebbles, however, the argument in the proof of Theorem D.28 breaks down, and the best lower bound we can get is as follows.

Theorem D.29 ([LT82]). *Suppose that \mathcal{P} is any complete black-white pebbling of Δ_{rev}^n with $\text{space}(\mathcal{P}) = s$ for $s \geq 3$. Then $\text{time}(\mathcal{P}) \geq \frac{n^2}{18s^2} + 2n$.*

The reason for the discrepancy between Theorem D.28 and Theorem D.29 turns out to be that in fact, it is possible to do better using white pebbles in addition to the black ones (essentially mimicking the proof of the lower bound in Theorem D.29). In particular, there is a linear-time black-white pebbling strategy for Δ_{rev}^n using only order of \sqrt{n} pebbles.

Theorem D.30 ([LT82]). *Let Δ_{rev}^n be the bit reversal graph over $n = 2^m$ elements. Then for all $s \geq 3$, there is a complete black-white pebbling \mathcal{P} of Δ_{rev}^n with $\text{space}(\mathcal{P}) \leq s$ and $\text{time}(\mathcal{P}) \leq 144 \frac{n^2}{s^2} + 18n$.*

On a high level, the reason that black-white pebblings can do much better than black-only pebblings on bit reversal graphs is that these graphs have such a regular structure. Lengauer and Tarjan raise the question whether there are other permutations for which the lower bound in Theorem D.28 holds also for black-white pebbling, and conjecture that the answer is yes. To the best of our knowledge, this problem is still open. One obvious question is whether anything interesting can be said about what holds for a random permutation in this respect. If the conjecture turns out to be true for a random permutation (with high probability, say), then such a result, although non-constructive, would be interesting.

D.3.2 DAGs Yielding Robust Pebbling Trade-offs

To get robust pebbling trade-offs, i.e., trade-offs that hold over a large space interval, we use a DAG family studied in [LT82, Section 4].

Definition D.31 (Stack of superconcentrators ([LT82])). Let SC_m denote any (explicitly constructible) linear-size m -superconcentrator with bounded indegree and depth $\log m$. We let Φ_r^m denote the graph constructed by stacking r copies SC_m^1, \dots, SC_m^r of SC_m on top of one another, with the sinks $z_1^j, z_2^j, \dots, z_m^j$ of SC_m^j connected to the sources $s_1^{j+1}, s_2^{j+1}, \dots, s_m^{j+1}$ of SC_m^{j+1} by edges (z_i^j, s_i^{j+1}) for all $i = 1, \dots, m$ and all $j = 1, \dots, r - 1$.

Clearly, Φ_r^m has size $\Theta(rm)$. Figure 7 gives a schematic illustration of the construction.

Lengauer and Tarjan establish fairly detailed trade-off results for stacks of superconcentrators using different explicit and non-explicit constructions for the superconcentrator building blocks. All of these results can be translated into corresponding trade-off results in resolution. For simplicity and conciseness, however, we only state a special case of their results in this conference version, deferring a more detailed treatment to the coming full-length version of the paper.

Theorem D.32 ([LT82]). *Let Φ_r^m denote a stack of (explicitly constructible) linear-size m -superconcentrator with bounded indegree and depth $\log m$. Then the following holds:*

1. $\text{Peb}(\Phi_r^m) = O(r \log m)$.
2. *There is a linear-time black pebbling strategy \mathcal{P} for Φ_r^m with $\text{space}(\mathcal{P}) = O(m)$.*
3. *If \mathcal{P} is a black-white pebbling strategy for Φ_r^m in space $s \leq m/20$, then $\text{time}(\mathcal{P}) \geq m \cdot \left(\frac{rm}{64s}\right)^r$.*

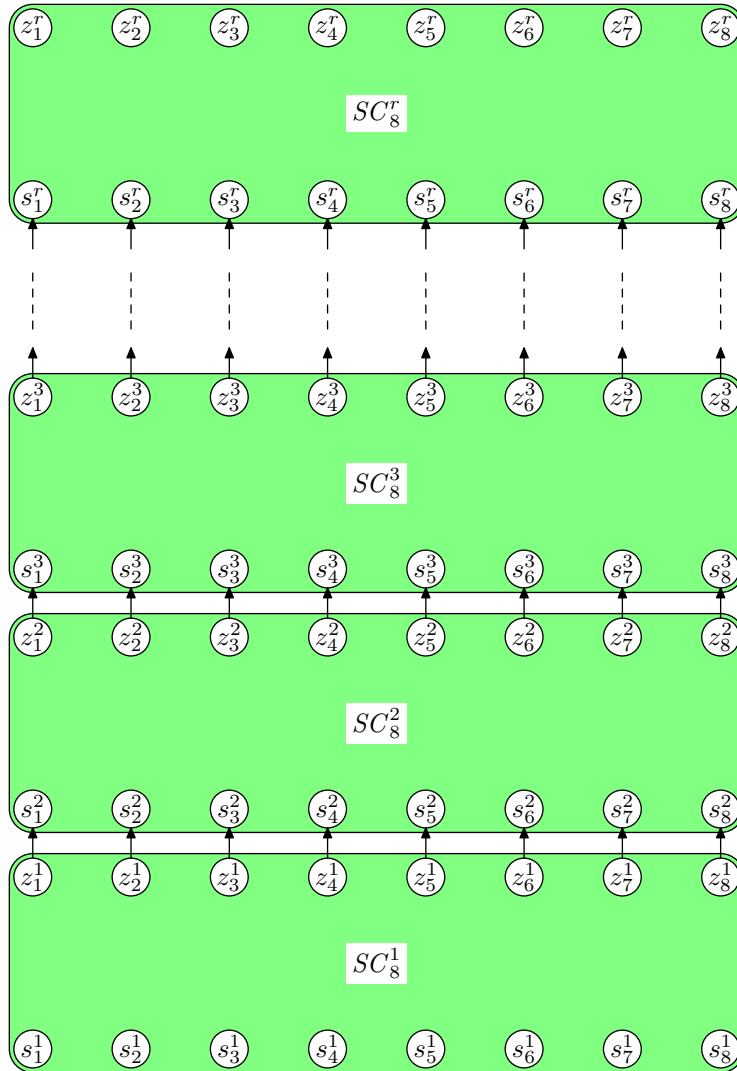


Figure 7: Schematic illustration of stack of superconcentrators Φ_r^8 .

Proof sketch. The upper bound on black pebbling price follows from Observation D.9, since the depth of Φ_r^m is $O(r \log m)$.

The linear-time black pebbling strategy is obtained by applying the trivial pebbling strategy in Observation D.7 consecutively to each superconcentrator, keeping pebbles on the sinks of SC_m^j while pebbling SC_m^{j+1} .

The reason that the final trade-off result holds is, very loosely put, that the lower bounds in Lemma D.14 and Corollary D.15 propagate through the stack of superconcentrators and get multiplied at each level. If the pebbling strategy is restricted to keeping s/r pebbles on each copy SC_m^j of the superconcentrator, this is not hard to prove directly from Lemma D.14. Establishing that this intuition holds also in the general case, when pebbles may be unevenly distributed over the superconcentrator copies, is much more technically challenging, however. \square

D.3.3 Exponential Pebbling Trade-offs

To get exponential trade-offs, i.e., trade-offs with lower bounds on the length on the form 2^{n^ϵ} for some constant $\epsilon > 0$, the graphs in Section D.3.2 are not sufficient. Instead, we need to appeal to stronger results from [LT82, Section 5].

Theorem D.33 ([LT82]). *For every $\ell \in \mathbb{N}^+$ there exist constants $c, c' > 1$ such that the following holds for all sufficiently large n . Let G be a DAG with n vertices and maximal indegree ℓ . Then for any space constraint s satisfying $cn/\log n \leq s \leq n$, there is a black pebbling strategy \mathcal{P} for G with $\text{space}(\mathcal{P}) \leq s$ and $\text{time}(\mathcal{P}) \leq s \cdot 2^{c'n/s}$.*

By stacking superconcentrators of *defferent* sizes on top of one another, Lengauer and Tarjan are able to prove a matching lower bound. We refer to [LT82, Section 5] for the details of the construction.

Theorem D.34 ([LT82]). *There exists a constant $\epsilon > 0$ such that the following holds for all sufficiently large integers n, s satisfying $cn/\log n \leq s \leq n$: There exists a DAG G with maximal indegree 2 and number of vertices at most n such that any black-white pebbling strategy \mathcal{P} for G with $\text{space}(\mathcal{P}) \leq s$ must have $\text{time}(\mathcal{P}) \geq s \cdot 2^{\epsilon n/s}$.*

Note that the graph G in Theorem D.34 depends on the pebbling space parameter s . Lengauer and Tarjan conjecture that no single graph gives an exponential time-space tradeoff for the whole range of $s \in [n/\log n, n]$, but to the best of our knowledge this problem is still open.

E Time-Space Trade-offs for Resolution

We have finally reached the point where we can state and prove our results for resolution. Given all the work done so far, the proofs are mostly simple variations of the following pattern: pick some graph family in Appendix D, make the appropriate choices of parameters, consider the corresponding pebbling contradiction CNF formulas, do f -substitution for some non-authoritarian function f , and apply Theorem C.4 (which we obtained with the help of Theorem 3.2).

Note that all the pebbling trade-off results are for explicit formulas (since they are pebbling formulas over explicitly constructible graphs). Note also that all trade-offs hold for variable space and clause space simultaneously, since the upper bounds are for variable space and the lower bounds for clause space.

E.1 Trade-offs for Constant Space

Time-space trade-offs in resolution occur even for formulas refuted in (very small) constant space. What is more, here we can specify the whole trade-off curve. (The rest of the results are more of a threshold type.)

Theorem E.1. *There are explicitly constructible families of minimally unsatisfiable k -CNF formulas $\{F_n\}_{n=1}^\infty$ of size $\Theta(n)$ such that:*

1. *Every formula F_n is refutable in resolution in length $L(F_n \vdash 0) = O(n)$ and also in variable space $\text{VarSp}(F_n \vdash 0) = O(1)$ (but not simultaneously).*
2. *For any $s > 0$ there is a resolution refutation $\pi_n : F_n \vdash 0$ in simultaneous variable space $\text{VarSp}(\pi_n) = O(s)$ and length $L(\pi) = O((n/s)^2 + n)$.*
3. *Any resolution refutation $\pi_n : F_n \vdash 0$ in clause space $\text{Sp}(\pi_n) = s$ for $s \geq \text{Sp}(F_n \vdash 0)$ must have length $L(\pi) \Omega((n/s)^2 + n)$.*

The constants hidden in the asymptotic notation are independent of n and s .

Proof. Fix any non-authoritarian function f and consider the pebbling formulas $\text{Peb}_{\Delta_{\text{rev}}^m}[f]$ defined over bit reversal DAGs (Definition D.27) for $m = \log n$.

Appealing to Theorem C.4 will get us a long way but not quite to our final destination. More precisely, the upper bounds on length and space follow from Lemma D.26 in this way, and the lower bound in the trade-off follows from Theorem D.29.

However, we cannot get the upper bound in the same manner, since Theorem D.28 tells us that there *cannot* exist black pebblings with parameters matching the lower bounds for black-white pebblings. However, in this particular case it turns out that we can mimic the asymptotically optimal black-white pebbling in Theorem D.30 in resolution in a space-preserving way. The details are not hard and will appear in the full-length version of this paper. \square

E.2 Superpolynomial Trade-offs for any Non-constant Space

It is clear that we can never get superpolynomial trade-offs from DAGs pebbleable in constant space, since such graphs must have constant-space pebbling strategies in polynomial time by a simple counting argument. However, perhaps somewhat surprisingly, as soon as we study *arbitrarily slowly* growing space, we can obtain superpolynomial trade-offs for formulas whose refutation space grows this slowly. This is a consequence of our new pebbling trade-off result in Section D.2.

Theorem E.2. *Let $g(n)$ be any arbitrarily slowly growing monotone function $\omega(1) = g(n) = O(n^{1/7})$, and let $\epsilon > 0$ be an arbitrarily small positive constant. Then there are explicitly constructible families of minimally unsatisfiable k -CNF formulas $\{F_n\}_{n=1}^\infty$ of size $\Theta(n)$ such that:*

1. *Every formula F_n is refutable in resolution in length $L(F_n \vdash 0) = O(n)$ and also in variable space $\text{VarSp}(F_n \vdash 0) = O(g(n))$ (but not simultaneously).*
2. *There are refutations $\pi_n : F_n \vdash 0$ in simultaneous variable space $\text{VarSp}(\pi_n) = O\left(\sqrt[3]{n/g^2(n)}\right)$ and length $L(\pi_n) = O(n)$.*
3. *There is a constant $K > 0$ such that any resolution refutation $\pi_n : F_n \vdash 0$ in clause space $\text{Sp}(\pi_n) \leq K(n/g^2(n))^{1/3-\epsilon}$ must have length $L(\pi_n)$ superpolynomial in n .*

The constant K as well as the constants hidden in the asymptotic notation are independent of n (but depend on g and ϵ).

We remark that the upper-bound condition $g(n) = O(n^{1/7})$ is very mild and is there only for technical reasons in this theorem. If we allow the minimal space to grow as fast as n^ϵ for some $\epsilon > 0$, then there are other pebbling trade-off results that can give even stronger results for resolution than the one stated above (see, for instance, Section E.4). Thus the interesting part is that $g(n)$ is allowed to grow arbitrarily slowly.

Proof of Theorem E.2. Consider the graphs Γ_r^c in Definition D.16. We want to choose the parameters c and r in a suitable way so that get a family of graphs in size $n = \Theta(cr^3 + c^3r^2)$ (using the bound on the size of Γ_r^c from Lemma D.17). If we choose

$$r = r(n) = g(n) \quad (26)$$

for $g(n) = O(n^{1/7})$, this forces

$$c = c(n) = \Theta\left(\sqrt[3]{n/g^2(n)}\right). \quad (27)$$

Consider the graph family $\{G_n\}_{n=1}^\infty$ defined by $G_n = \Gamma_{r(n)}^{c(n)}$ as in (26) and (27), which is a family of size $\Theta(n)$. Construct the single-sink version \widehat{G}_n of G_n , fix any non-authoritarian function f , consider the pebbling formulas $F_n = \text{Peb}_{\widehat{G}_n}[f]$, and appeal to the translation between pebbling and resolution in Theorem C.4.

Lemma D.17 yields that $\text{VarSp}(F_n \vdash 0) = O(g(n))$. Also, the persistent black pebbling of G_n in Lemma D.18 yields a linear-time refutation $\pi_n : F_n \vdash 0$ with $\text{VarSp}(\pi_n) = O(\sqrt[3]{n/g^2(n)})$.

Now set the parameter s in Theorem D.19 to $s = c^{1-\epsilon'}$ for $\epsilon' = 3\epsilon$. Then for large enough n we have $s \leq c/8 - 1$ and Theorem D.19 can be applied. Combining the pebbling trade-off there with Theorem C.4, we get that if the clause space is less than $(n/g^2(n))^{1/3-\epsilon}$, then the required length of the refutation grows as $(\Omega(cc'))^r = (\Omega(n/g^2(n)))^{\epsilon g(n)}$ which is superpolynomial in n for any $g(n) = \omega(1)$. The theorem follows. \square

E.3 Robust Superpolynomial Trade-offs

We now know that there are polynomial trade-offs in resolution for constant space, and that going ever so slightly above constant space we can get superpolynomial trade-offs. The next question we want to focus on is how robust trade-offs we can get. That is, over how large a range of space does the trade-off hold? Given minimal refutation space s , how much larger space is needed in order to obtain the linear length refutation that we know exists for any pebbling contradiction?

The answer is that we can get superpolynomial trade-offs that span almost the whole range between constant and linear space. We present to different results

Theorem E.3. *There are explicitly constructible families of minimally unsatisfiable k -CNF formulas $\{F_n\}_{n=1}^\infty$ of size $\Theta(n)$ such that:*

1. *Every formula F_n is refutable in length $L(F_n \vdash 0) = O(n)$ and also in variable space $\text{VarSp}(F_n \vdash 0) = O(\log n)$.*
2. *There is a resolution refutation $\pi_n : F_n \vdash 0$ in variable space $\text{VarSp}(\pi_n) = O\left(\sqrt[3]{n/\log^2 n}\right)$ and length $L(\pi_n) = O(n)$.*
3. *There is a constant $K > 0$ such that any resolution refutation $\pi_n : F_n \vdash 0$ in clause space $\text{Sp}(\pi_n) \leq K \sqrt[3]{n/\log^2 n}$ must have length $L(\pi_n) = n^{\Omega(\log \log n)}$.*

The constant K as well as the constants hidden in the asymptotic notation are independent of n .

Proof. Consider the graphs Γ_r^c in Definition D.16 with parameters chosen so that $c = 2^r$. Then the size of Γ_r^c is $\Theta(r^2 2^{3r})$ by Lemma D.17. Let $r(n) = \max\{r : r^2 2^{3r} \leq n\}$ and define the graph family $\{G_n\}_{n=1}^\infty$ by $G_n = \Gamma_{r(n)}^{2^r}$ for $r = r(n)$. Finally, construct the single-sink version \widehat{G}_n of G_n , fix any non-authoritarian function f and consider the pebbling formulas $F_n = \text{Peb}_{\widehat{G}_n}[f]$ with the help of Theorem C.4.

Translating from G_n back to Γ_r^c we have parameters $r = \Theta(\log n)$ and $c = \Theta((n/\log^2 n)^{1/3})$, so Lemma D.17 yields that $\text{VarSp}(F_n \vdash 0) = O(\log n)$. Also, the persistent black pebbling of G_n in Lemma D.18 yields a linear-time refutation $\pi_n : F_n \vdash 0$ with $\text{VarSp}(\pi_n) = O((n/\log^2 n)^{1/3})$.

Setting $s = c/8 - 1$ in Theorem D.19 shows that there is a constant K such that if the clause space of a refutation $\pi_n : F_n \vdash 0$ drops below $K \cdot (n/\log^2 n)^{1/3} \leq (r + 2) + s$, then we must have

$$L(\pi_n) \geq O(1)^r \cdot r! = n^{\Omega(\log \log n)} \quad (28)$$

(where we used that $r = \Theta(\log n)$ for the final equality). The theorem follows. \square

Sacrificing a square at the lower end of the interval, we can improve the upper end to $n/\log n$.

Theorem E.4. *There are explicitly constructible families of minimally unsatisfiable k -CNF formulas $\{F_n\}_{n=1}^\infty$ of size $\Theta(n)$ such that:*

1. *Every formula F_n is refutable in resolution in length $L(F_n \vdash 0) = O(n)$ and also in variable space $\text{VarSp}(F_n \vdash 0) = O(\log^2 n)$.*
2. *There is a resolution refutation $\pi_n : F_n \vdash 0$ in variable space $\text{VarSp}(\pi_n) = O(n/\log n)$ and length $L(\pi_n) = O(n)$.*
3. *There is a constant $K > 0$ such that any resolution refutation $\pi_n : F_n \vdash 0$ in clause space $\text{Sp}(\pi_n) \leq Kn/\log n$ must have length $L(\pi_n) = n^{\Omega(\log \log n)}$.*

The constant K and the constants hidden in the asymptotic notation are independent of n .

Proof. Pick any non-authoritarian function f and consider the pebbling formulas $\text{Peb}_{\widehat{\Phi_r^m}}[f]$ defined over single-sink versions of stacks of superconcentrators Φ_r^m as in Definition D.31 with $m = 20T$ and $r = \lfloor n/T \rfloor$ for $T = \Theta(n/\log n)$. The theorem now follows by combining Theorem D.32 with Theorem C.4. \square

We remark that the results in Theorem E.4 can perhaps be considered to be slightly stronger than those in Theorem E.3, but they require a very much more involved graph construction with worse hidden constants than the very simple and clean construction underlying Theorem E.3.

E.4 Exponential Trade-offs

Superpolynomial trade-offs are all fine and well, but can we get *exponential* trade-offs? In this final subsection we answer this question in the affirmative.

The same counting argument that was mentioned in the beginning of Section E.2 tells us that we can never expect to get exponential trade-offs from DAGs with polylogarithmic pebbling price. However, if we move to graphs with pebbling price $\Omega(n^\epsilon)$ for some constant $\epsilon > 0$, pebbling formulas over such graphs can exhibit exponential trade-offs.

We obtain our first such exponential trade-off, which also exhibits a certain robustness, by again studying the DAGs in Definition D.16.

Theorem E.5. *There are explicitly constructible families of minimally unsatisfiable k -CNF formulas $\{F_n\}_{n=1}^\infty$ of size $\Theta(n)$ such that:*

1. *Every formula F_n is refutable in resolution in length $L(F_n \vdash 0) = O(n)$ and also in variable space $\text{VarSp}(F_n \vdash 0) = O(\sqrt[8]{n})$.*
2. *There is a resolution refutation $\pi_n : F_n \vdash 0$ in variable space $\text{VarSp}(\pi_n) = O(\sqrt[4]{n})$ and length $L(\pi_n) = O(n)$.*

3. There is a constant $K > 0$ such that any resolution refutation $\pi_n : F_n \vdash 0$ in clause space $Sp(\pi_n) \leq K \sqrt[4]{n}$ must have length $L(\pi_n) = (\sqrt[8]{n})!$.

The constant K as well as the constants hidden in the asymptotic notation are independent of n .

Proof. Combine Theorem C.4 and Theorem D.19 in the same way as in the other proofs above for Γ_r^c with $c = \sqrt[4]{n}$ and $r = \sqrt[8]{n}$. \square

We remark that there is nothing magic in our particular choice of parameters c and r in Theorem E.5. Other parameters could be plugged in instead and yield slightly different results.

Now that we know that there are robust exponential trade-offs for resolution, we want to obtain exponential trade-offs for formulas with their minimal refutation space being as large as possible.

The higher the lower bound on space is, the more interesting the trade-off gets. Remember that a SAT solver refuting the formula will very likely use at least linear space to do so. It is unclear why the SAT solver would work hard on optimizing lower order terms in the memory consumption and thus get stuck in a trade-off for relatively small space. Ideally, therefore, we would like to obtain trade-offs for superlinear space (if there are such trade-offs, that is). For such formulas, we would be more confident that the trade-off phenomena should also show up in practise (although the case can certainly be made that even sublinear space trade-offs could possibly be relevant for real life applications).

It is clear that pebbling contradictions can never yield any trade-off results in the superlinear regime, since they are always refutable in linear length and linear space simultaneously. Also, all trade-offs obtainable from the graphs in Definition D.16 will be for space far below linear. However, using results from Section D.3.3 we can get exponential trade-offs for space almost linear, or more precisely for space as large as $\Theta(n/\log n)$.

Theorem E.6. *There are explicitly constructible families of minimally unsatisfiable k -CNF formulas $\{F_n\}_{n=1}^\infty$ of size $\Theta(n)$ such that:*

1. Every formula F_n is refutable in length $L(F_n \vdash 0) = O(n)$ and variable space $VarSp(F_n \vdash 0) = O(n/\log n)$.
2. There is a resolution refutation $\pi_n : F_n \vdash 0$ in variable space $VarSp(\pi_n) = O(n)$ and length $L(\pi) = O(n)$.
3. There is a constant $K > 0$ such that any resolution refutation $\pi_n : F_n \vdash 0$ in clause space $Sp(\pi_n) \leq Kn/\log n$, where $Kn/\log n \geq Sp(F_n \vdash 0)$, must have length $L(\pi) = \exp(n^\epsilon)$.

All constants, including those hidden in the asymptotic notation, are independent of n .

Proof. Use Theorem D.34 and Theorem C.4. \square

We remark that again, Theorem D.34 in combination with Theorem D.33 can be used to obtain DAGs (and thus CNF formulas) with other trade-offs as well for different space parameters in the range between $n/\log n$ and n . For simplicity and conciseness, however, we only state the special case above in this conference version, deferring a more detailed treatment to the coming full-length version of the paper.