



Linear and Negative Resolution are weaker than Resolution

Josh Buresh-Oppenheim

Computer Science Department

University of Toronto

Toronto, ON

bureshop@cs.toronto.edu

David Mitchell

School of Computing Science

Simon Fraser University

Burnaby, BC

mitchell@cs.sfu.ca

Toniann Pitassi*

Computer Science Department

University of Toronto

Toronto, ON

toni@cs.toronto.edu

October 10, 2001

Abstract

We prove exponential separations between the sizes of particular refutations in negative, respectively linear, resolution and general resolution. Only a superpolynomial separation between negative and general resolution was previously known. Our examples show that there is no strong relationship between the size and width of refutations in negative and linear resolution.

* Research partially supported by NSF grant CCR-9457782, US-Israel BSF Grant 95-00238, and an NSERC University Faculty Award.

1 Introduction

Lower bounds for resolution have been greatly simplified in the past few years, due to a fundamental relationship between proof size and proof width. That is, it has been shown recently [BSW99] that for any 3CNF formula F in n underlying variables, F has a resolution proof of size S if and only if there is a resolution proof of F with maximal clause size $\sqrt{n \log S}$.

This relationship has greatly simplified most lower bounds for resolution, reducing the problem to showing a wide-clause lemma.

In this paper, we prove exponential separations for both linear and negative resolution. That is, we give examples of formulas that have linear size resolution proofs, but requiring exponential-size proofs in both linear and negative resolution. Previously, Goerdt [Goe92] has obtained a quasipolynomial separation for negative resolution. It appears that no separations were previously known between linear and unrestricted resolution.

We also prove that there is no analogous size-width tradeoff for either linear or negative resolution.

2 Definitions

2.1 Resolution

The *resolution principle* says that if C and D are clauses and x is a variable, then any assignment that satisfies both of the clauses $C \vee x$ and $D \vee \neg x$ also satisfies $C \vee D$. The clause $C \vee D$ is said to be a *resolvent* of the clauses $C \vee x$ and $D \vee \neg x$ derived by *resolving on* the variable x . A *resolution derivation* of a clause C from a CNF formula F consists of a sequence of clauses in which each clause is either a clause of F , or is a resolvent of two previous clauses, and C is the last clause in the sequence; it is a *refutation* of F if C is the empty clause Λ .

The *size* of a refutation is the number of resolvents in it. The *width* of a clause is the number of literals occurring in the clause. The width of a refutation is the maximum width of all clauses occurring in the refutation.

We can represent a resolution refutation as a directed acyclic graph (dag) where the nodes are the clauses in the refutation, each clause of F has out-degree 0, and any other clause has two arcs pointing to the two clauses that produced it. The arcs pointing to $C \vee x$ and $D \vee \neg x$ are labeled with the literals x and $\neg x$ respectively. It is well known that resolution is a *sound* and *complete* propositional proof system, i.e., a formula F is unsatisfiable if and only if there is a resolution refutation for F .

A *negative resolution* refutation of F is a resolution refutation with the additional restriction that all resolutions must be negative. A resolution step $C \vee x$ and $D \vee \neg x$ implies $C \vee D$ is negative whenever D contains only negative literals.

A *linear resolution* refutation of F is a resolution refutation with the additional restriction that the underlying dag must be linear. That is, the proof consists of a sequence of clauses C_1, C_2, \dots, C_m

such that C_m is the empty clause, for every $1 \leq i \leq m$, either C_i is an initial clause, or C_i is derived from C_{i-1} and an initial clause, or C_i is derived from C_{i-1} and C_j , for some $j < i - 1$.

Both negative and linear resolution are sound and complete.

An *assignment* for a formula F (sometimes called a *restriction*) is a Boolean assignment to some of the variables in the formula; the assignment is *total* if all the variables in the formula are assigned values. If C is a clause, and σ an assignment, then we write $C \upharpoonright \sigma$ for the result of applying the assignment to C , that is, $C \upharpoonright \sigma = 1$ if $\sigma(l) = 1$ for some literal l in C , otherwise, $C \upharpoonright \sigma$ is the result of removing all literals set to 0 by σ from C (with the convention that the empty clause is identified with the Boolean value 0). If F is a CNF formula, then $F \upharpoonright \sigma$ is the conjunction of all the clauses $C \upharpoonright \sigma$, C a clause in F .

If $R = C_1, \dots, C_k$ is a resolution derivation from a formula F , and σ an assignment to the variables in F , then we write $R \upharpoonright \sigma$ for the sequence $C_1 \upharpoonright \sigma, \dots, C_k \upharpoonright \sigma$.

Lemma 1: If R is a linear (negative) resolution derivation of C from a formula F , and σ an assignment, then there is a subsequence of $R \upharpoonright \sigma$ that is a linear (regular) resolution derivation of $C \upharpoonright \sigma$ from $F \upharpoonright \sigma$.

Proof: This is a straightforward induction on the length of the derivation from F . \square

2.2 Tautologies on graphs

Our hard formulas are from [BSW99]. They are a generalization of the implication graph formulas, first introduced by Raz and McKenzie [RM97], and also used in subsequent papers [BEGJ98, BOCIP00]. Let G be a directed, acyclic graph, with fan-in 2, n vertices, and a single sink vertex.

The implication graph formulas encode the following contradictory statement: “All of the source vertices are colored red, the sink is colored blue, and if both the predecessors of a vertex are red, so is the vertex itself.”

The formula associated with G , $Imp(G)$ has one variable, x_i , for every node i in G , and the following clauses: (1) for each source node i in G , (x_i) ; (2) for the sink node s in G , $(\neg x_s)$; (3) for every triple of nodes i, j , and k such that the edges (i, k) and (j, k) are present in G , we have the clause $(\neg x_i \vee \neg x_j \vee x_k)$.

The natural way to refute the above formula/clauses is to begin at the source vertices, and derive successively that each layer of vertices must be true, until finally we can conclude that each sink vertex must be true. This gives us the desired contradiction since the sink vertex is false. For any graph D with indegree 2, this natural refutation can be formalized as a linear-size tree-like resolution refutation.

However, we show here that if the graph is sufficiently complicated (it has high pebbling number), then any linear or negative resolution refutation of $Imp(G)$ must have large width.

We also define a more general formula, $Imp^*(G)$, based on G as follows. Now there are two variables x_i and y_i associated with a vertex i in G . The formula is the following conjunction of

clauses: (1) for each source vertex i in G , $(x_i \vee y_i)$; (2) for the sink vertex s in G , $(\neg x_s)$ and $(\neg y_s)$; (3) for every i, j , and k such that (i, k) and (j, k) are edges in G , we have the following clauses stating that if one of the variables associated with i is true, and one of the variables associated with j is true, then one of the variables associated with k is also true: $(\neg x_i \vee \neg x_j \vee x_k \vee y_k)$, $(\neg x_i \vee \neg y_j \vee x_k \vee y_k)$, $(\neg y_i \vee \neg x_j \vee x_k \vee y_k)$, and $(\neg y_i \vee \neg y_j \vee x_k \vee y_k)$.

We will show that any linear or negative resolution refutation of $\text{Imp}^*(G)$ requires exponential size for certain G .

2.3 Graphs with high pebbling

We will show that a negative or linear resolution proof of $\text{IMP}(G)$ of small width can be converted into an efficient pebbling strategy for the corresponding graph, G . Interesting connections between pebbling and propositional proofs were made previously in [ET99, BSW99].

DEFINITION 2.1: Let $D = (V, E)$ be a directed, acyclic graph. A configuration is a subset of V . A legal pebbling of a vertex v in D is a sequence of configurations, the first being the empty set and the last being $\{v\}$ and in which each configuration C' follows from the previous configuration C by one of the following rules:

1. v can be added to C to get C' if all immediate predecessors of v are in C .
2. Any vertex can be removed from C to obtain C' .

The *complexity* of a legal pebbling of v is the maximal size of any configuration in the sequence. The *pebbling number* of a graph D with a single sink vertex s is the minimal number n such that there exists a legal pebbling of s with complexity n .

Cook [Coo73] showed that the pyramid graphs, Pyramid_n , with $n = m + (m - 1) + \dots + 1$ underlying vertices have pebbling measure $\Omega(\sqrt{n})$. These are layered graphs, consisting of m layers, with m source vertices at layer 1, labelled $x_1^1, x_2^1, \dots, x_m^1$, $m - 1$ vertices at layer 2, labelled x_1^2, \dots, x_{m-1}^2 , and so on with one sink vertex, x_1^m at layer m . All nonsource vertices have indegree 2, and in general x_j^{i+1} has parents x_j^i and x_{j+1}^i .

[PTC77] exhibits a sequence of graphs, G_n , based on a construction by Valiant that have n nodes and in-degree 2, but with pebbling measure $\Omega(n/\log n)$. This is an optimal lower bound, since [HPV75] shows that any graph has pebbling number $O(n/\log n)$

3 Lower Bounds

The negation-width of a clause C is the number of negative literals occurring in C . The negation width of a resolution refutation P is the maximum negation-width of all clauses in P . The lower bounds for both linear and negative resolution will follow the same strategy. We will begin with an




Figure 1: A $Pyramid_{10}$ graph

alleged small linear (or negative) resolution refutation of $IMP^*(G)$, where G is a graph with high pebbling number. The following lemma shows that we can always find a restriction ρ such that after applying ρ to P , what remains is a linear (or negative) refutation of $IMP(G)$, but now with small negation-width. Then we will use particular properties of linear and negative resolution to argue that that any linear (negative) refutation of $IMP(G)$ requires large negation-width, thus reaching a contradiction.

Lemma 2: For any graph G_n , if there is a linear (negative) resolution refutation of $IMP^*(G_n)$ of size at most S , then there is a linear (negative) resolution refutation of $IMP(G_n)$ of negation-width at most w , where $w > \log S$.

Proof: Let P be a linear (negative) resolution refutation of $IMP^*(G_n)$ of size at most S . Call a clause of P negation-wide if its negation-width is at least w . Let C_1, \dots, C_m be the set of negation-wide clauses in P , and for each C_j , let s_j be a set of w negative literals occurring in C_j . Clearly m (the number of negation-wide clauses in P) is at most S .

We will define a restriction ρ probabilistically as follows. For every $i \in \{1, \dots, n\}$, choose x_i with probability $1/2$. Choose y_i if and only if x_i is not chosen. The assignment associated with ρ will set $x_i = 0$ if and only if x_i is chosen, and otherwise, sets $y_i = 0$.

We want to upper bound the probability that ρ is bad, where a restriction ρ is bad if not all negation-wide clauses in P are set to 1 by ρ .

A restriction ρ is good for a particular negation-wide clause C_j if some element in s_j was chosen by ρ . The probability that this does not happen is at most $(1/2)^w$. Therefore the overall probability that ρ is bad is at most $S(1/2)^w$. Since $\log S < w$, this overall probability is less than 1, and therefore there must exist at least one good ρ .

Fix a good such ρ and apply the restriction ρ to the entire proof P . What remains will be a linear (negative) resolution refutation of $IMP(G_n)$, of negation width at most w . \square

3.1 Lower bounds for negative resolution

Lemma 3: Any negative resolution refutation of $IMP(G_n)$ has negation-width at least $\Omega(q)$, where q is the pebbling number of G_n .

Proof: Let P be a negative resolution refutation of $IMP(G_n)$ of negation-width w . We will show how to use P to pebble G_n with at most w pebbles.

Any negative resolution refutation of $IMP(G_n)$ must begin by resolving on the clause $(\neg x_s)$, where s is the sink of G_n , since this is the only all negative clause occurring in $IMP(G_n)$. Furthermore, all new clauses produced by negative resolution are all-negative clauses. This is because all other initial clauses have exactly one positive literal, which gets resolved away. Thus, every resolution step in a negative resolution proof must involve resolving an initial clause with an all-negative clause.

If D_1, \dots, D_m is the sequence of all-negative clauses generated by the proof P , then this sequence in reverse, D_m, \dots, D_1 will be the sequence of configurations in our pebbling strategy for G_n . More precisely, if $\neg x_i$ occurs in D_j , then the configuration corresponding to D_j will include vertex i of G_n . It is clear that the pebbling number of our sequence of configurations corresponds to the negation-width of P . $D_m \dots D_1$ is a valid sequence of configurations since for any i , D_{i+1} must be the same as D_i except that one node in D_i is replaced by its parents in D_{i+1} . Therefore, we can go from pebbling configuration D_{i+1} to D_i in two moves of the pebbling game and with no extra pebbles.

Because G_n has pebbling number q , it follows that the negation-width of P must be at least q . \square

Theorem 4: For any graph G_n with pebbling measure q , any negative resolution refutation of $IMP^*(G_n)$ requires size $2^{\Omega(q)}$. In particular, there exists an infinite sequence of graphs G_n such that any negative resolution refutation of $IMP^*(G_n)$ requires size $2^{\Omega(n/\log n)}$.

Proof: The above theorem follows from Lemma 2 and Lemma 3. \square

3.2 Lower bounds for linear resolution

We begin by analyzing the structure of linear resolution refutations. Let $P = C_1, \dots, C_m$ be a linear resolution refutation of $IMP(G)$, for some G . Let $P' = C'_1, \dots, C'_r$, for $r < m$, be the subsequence of P constructed by removing all initial clauses (i.e. clauses in $IMP(G)$). First note that every clause in P must be a horn clause; that is, each clause in P involves at most one positive literal. This follows from the fact that all clauses in $IMP(G)$ are horn and that any resolvent of two horn clauses is horn. Let $p = p_1, p_2, \dots, p_\ell$ be the ordered sequence of vertices of G such that $p_k = j$ if and only if x_j is the k th distinct variable occurring positively in the sequence P' . We claim that p must be a simple path in G : consider some C'_i in P' such that C'_{i+1} adds a new vertex to p_k to p . C'_{i+1} must be the resolvent of C'_i and some other clause D , since the proof is linear. If C'_{i+1} adds a new vertex to p , however, then D must be an initial clause and the resolution must be on the variable which appears positively in C'_i , since otherwise p_k would already appear in p . This means that the variable appearing positively in C'_i must be an immediate predecessor of p_k . The path is simple because


the graph is acyclic. The path p will be called the path associated with P . Note finally that p may originate at any vertex in G , but that it must reach the sink since we must at some point use the sole initial clause mentioning the sink ($IMP(G)$ without this clause would be satisfiable).

This structure gives us immediate insight into any refutation on the pyramid graph:

Lemma 5: Let $Pyramid_n$ be the pyramid graphs with $n = m + (m - 1) + \dots + 1$ vertices. Then any linear resolution refutation of $IMP(Pyramid_n)$ has negation-width at least $\Omega(\sqrt{n}/2)$.

Lemma 6: For any path p in $Pyramid_n$, there is a pyramid graph, $Pyramid_{n'}$ embedded in $Pyramid_n$ such that no vertices of $Pyramid_{n'}$ intersect p , and such that $n' \geq (m/2 - 1) + (m/2 - 2) + \dots + 1$.

Proof: If path p does not begin at a source of $Pyramid_n$, consider any path p' from a source of $Pyramid_n$ (say, x_k^1) to the start of p . Assume without loss of generality that $k \geq \lceil m/2 \rceil$. Consider the pyramid with sources $x_1^1, \dots, x_{\lceil m/2 \rceil}^1$. Clearly there is no path starting at x_k^1 that intersects this pyramid, so the path $p'p$ cannot intersect it and, in particular, p cannot intersect it.



□

Lemma 7: Let P be a linear resolution refutation of negation-width w of $IMP(Pyramid_n)$. Let p be the path associated with P , and let $Pyramid_{n'}$ be the subgraph of $Pyramid_n$ given by the above claim. Then $Pyramid_{n'}$ can be pebbled with w pebbles.

Proof: Let $P = C_1, \dots, C_m, P' = C'_1, \dots, C'_r$. We want to obtain a sequence of configurations (sets of vertices of $Pyramid_{n'}$), where the sequence in reverse will be a pebbling strategy for $Pyramid_{n'}$. The first configuration in the sequence is the configuration consisting of just the sink vertex of $Pyramid_{n'}$. Consider the first clause in P' where the variable corresponding to the sink vertex t of $Pyramid_{n'}$ appears negatively. At some point, this clause is resolved with an initial clause of $IMP(Pyramid_n)$ on the variable t , so t is replaced by its parents. This is because we must use every clause in $IMP(Pyramid_n)$ and because our copy of $Pyramid_{n'}$ with sink t is independent of p .

The next configuration in the sequence are these two parent vertices. Following through P' , we will come to some first place where one of the vertices, v' , in the current configuration is resolved upon, and replaced by its two parents. The next configuration in the sequence will be obtained from the current configuration by replacing v' with its two parents.

This sequence of configurations in reverse will be a valid pebbling strategy for $Pyramid_{n'}$ since each pebbling configuration follows from the previous one by two moves of the pebbling game. \square

Proof: (of Lemma 5) We will now complete the proof of Lemma 5. Since $n' = \Omega(n/4)$, $Pyramid_{n'}$ requires $\Omega(\sqrt{n})$ pebbles. Assume for sake of contradiction that P has negation width $o(\sqrt{n})$. By the above lemma, this gives us a pebbling strategy for $G_{n'}$ violating the known pebbling measure for $G_{n'}$. \square

Theorem 8: Any linear resolution refutation of $IMP^*(Pyramid_n)$ requires size $2^{\Omega(\sqrt{n})}$.

Proof: The above theorem follows by Lemma 2 and Lemma 5. \square

We will now present a better lower bound for linear resolution, by utilizing graphs with maximal pebbling measure. Let G_n be the graphs of in-degree 2 given by [PTC77] with n underlying vertices. We modify these graphs slightly so that the multiple sinks are identified to one using a binary tree.

Lemma 9: Any linear resolution refutation of $IMP(G_n)$ has negation-width at least $\Omega(n/\log n)$.

Let P be a linear resolution refutation of $IMP(G_n)$. As before, corresponding to P is a path $p = p_1, \dots, p_\ell$ in G_n . The path p has an origin, and must end at the sink of G_n . Let $anc(p_i)$ be the direct ancestors of p_i that lie outside of p (if there are any), and let $anc(p) = \cup_{i=1}^{\ell} anc(p_i)$.

Lemma 10: Consider p as a induced subgraph of G_n . Call this graph G_p . G_p has pebbling number $O(\log n)$.

Proof: In the construction used by [PTC77], $G_n = G(i)$, for some $i \leq \log n$, where $G(i)$ is constructed recursively from two copies of $G(i-1)$ and two copies of $C(i-1)$ (the basic building block is $C(8) = G(8)$). The graphs are combined in series so that the sinks of $C_1(i-1)$ are connected to the sources of $G_1(i-1)$, whose sinks are connected to the sources of $G_2(i-1)$, whose sinks are connected to the sources of $C_2(i-1)$. Sources for $G(i)$ are added before the sources of $C_1(i-1)$ and sinks are added after the sinks of $C_2(i-1)$. The only other edges that are added go directly from the sources of $G(i)$ to its sinks.

Now consider any path p in $G(i)$. For a vertex v on p , let $G^v(i-1)$ be the subgraph (i.e. $G_1(i-1)$, $G_2(i-1)$, $C_1(i-1)$ or $C_2(i-1)$) that contains v . Let $C^v(8)$ be the copy of $C(8)$ that contains v . $G(i)$ can contribute only one edge not in $G^v(i-1)$ that goes from an ancestor of v in P to a descendant of v in P . Hence there are only $i-8$ such edges in $G(i)$ (except those contributed by $C^v(8)$).




Figure 2: A $G(i+1)$ graph

We now use the following pebbling strategy to pebble G_p : let c be the length of the longest path in $C(8)$ (a constant number). Starting at the beginning of p , lay down c pebbles and subsequently always leave a trail of c pebbles behind the current node. This will take care of pebbling any ancestors of the next node that are contained in the same copy of $C(8)$. Whenever we reach a node that is the parent of a future node in the path, leave a pebble on it until its descendent is pebbled. This is guaranteed to add at most $i \leq \log n$ pebbles. \square

Corollary 11: The set $anc(p)$ is not empty.

Proof: Otherwise G_n would have pebbling number $O(\log n)$. \square

We now introduce a slight variation on the pebbling game. Consider a dag containing nodes u_1, \dots, u_k and v . The pebbling number of v from u_1, \dots, u_k is the minimal number of pebbles needed to pebble v where we add the rule that during any step a pebble may be placed on u_i , for $1 \leq i \leq k$.

Lemma 12: There exists an i , such that $1 \leq i \leq \ell$ and some pebble in $anc(p_i)$ has pebbling number $\Omega(n/\log n)$ from the sources of G_n and $\{p_1, \dots, p_{i-1}\}$.

Proof: We first show that there must be a node in $anc(p)$ that has pebbling number $\Omega(n/\log n)$ from the sources of G_n . If this were not the case, we could use the following strategy to pebble the sink of G_n in $o(n/\log n)$ pebbles: use the $O(\log n)$ pebbling strategy for G_p , except that whenever we are about to place a pebble on p_i , pebble the (at most 2) pebbles in $anc(p_i)$ using $o(n/\log n) + O(\log n) + 1 = o(n/\log n)$ pebbles, place the pebble on p_i and remove all the pebbles used to pebble $anc(p_i)$.

Now, consider the smallest i such that $anc(p_i)$ contains a node with pebbling number $\Omega(n/\log n)$. Call this node v . Note that any node in $\{p_1, \dots, p_{i-1}\}$ can be pebbled from the sources of G_n using $o(n/\log n)$ pebbles using the $\log n$ pebbling strategy for G_p and the fact that any node in $anc(p_j)$ for $j < i$ can be pebbled in $o(n/\log n)$ pebbles. If v could be pebbled from the sources of G_n and $\{p_1, \dots, p_{i-1}\}$ using $o(n/\log n)$ pebbles, then we could pebble v from the sources of G_n using that strategy and spending $o(n/\log n)$ pebbles whenever a pebble needs to be placed on a node in $\{p_1, \dots, p_{i-1}\}$. Therefore we can pebble v in $o(n/\log n)$ pebbles—a contradiction. \square

Lemma 13: Let P be a linear resolution refutation of negation-width w of $IMP(G_n)$. Let p be the path associated with P . Any vertex in G_n not lying on p can be pebbled from the sources of G_n and the vertices of p using w pebbles.

Proof: Analogously to lemma 7, start with the first clause of P^l that mentions a given node v negatively. The set $\{v\}$ is our initial configuration. Since v is not on p , it must be resolved upon using an initial clause and therefore replaced by its parents in the resolvent. Likewise, v is replaced by its parents in our next configuration. If u appears in any configuration and u is either a source of G_n or lies on p , then u remains in all the subsequent configurations and is not replaced by any of its ancestors. Again, this sequence of configurations forms a pebbling strategy in reverse. \square

Proof: (of lemma 9) Immediate from lemmas 13 and 12. \square

Theorem 14: Any linear resolution refutation of $IMP^*(G_n)$ requires size $2^{\Omega(n/\log n)}$.

Proof: Let P be a linear resolution refutation of $IMP^*(G_n)$ of size $S = 2^{o(n/\log n)}$. Applying Lemma 2, it follows that there is also a linear resolution refutation of $IMP(G_n)$ of width $o(n/\log n)$. But this violates Lemma 9. \square

3.3 Size versus width

It is a simple corollary that there are no size-width relationships for either negative or linear resolution.

Corollary 15: There is no size-width tradeoff for linear (negative) resolution. More specifically, the formulas $IMP(G_n)$ have polynomial-size linear (negative) resolution refutations but any linear (negative) resolution refutation requires width $\Omega(n/\log n)$.

Proof: We will describe a resolution proof that is both negative and linear. Start from the root of the graph, and work up towards the sources, finally deriving that one of the (variables associated with the) sources has to be 0. Resolving this clauses with the initial clauses expressing that each source is 1 produces the empty clause. By the above wide clause lemmas for linear and negative resolution, any such refutation of $IMP(G_n)$ requires large negation-width, and thus also requires large width. \square

Acknowledgements

We are very grateful to Alasdair Urquhart for helpful discussions and generous feedback on this work.

References

- [BEGJ98] M.L. Bonet, J. L. Esteban, N. Galesi, and J. Johannsen. Exponential separations between restricted resolution and cutting planes proof systems. pages 638–647, 1998.
- [BOCIP00] J. Buresh-Oppenheim, M. Clegg, R. Impagliazzo, and T. Pitassi. Homogenization and the polynomial calculus. 2000.
- [BSW99] E. Ben-Sasson and A. Wigderson. Short proofs are narrow – resolution made simple. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 517–526, Atlanta, GA, May 1999.
- [Coo73] Stephen A. Cook. An observation on time-storage trade off. In *Conference Record of Fifth Annual ACM Symposium on Theory of Computing*, pages 29–33, Austin, TX, April-May 1973.
- [ET99] J. L. Esteban and Jacobo Toran. Space bounds for resolution. In *Proc. 16th STACS*, pages 551–561. Springer-Verlag LNCS, 1999.
- [Goe92] A. Goerdt. Unrestricted resolution versus n-resolution. *Theoretical Computer Science*, 93:159–167, 1992.
- [HPV75] John E. Hopcroft, Wolfgang J. Paul, and Leslie G. Valiant. On time versus space and related problems. In *16th Annual Symposium on Foundations of Computer Science*, pages 57–64, Berkeley, CA, October 1975. IEEE.
- [PTC77] Wolfgang J. Paul, Robert E. Tarjan, and J. R. Celoni. Space bounds for a game on graphs. *Mathematical Systems Theory*, 10(3):239–251, 1977. Correction, *ibid.* 11(1):85, 1977.
- [RM97] R. Raz and P. McKenzie. Separation of the monotone nc hierarchy. 1997.