

Counting Classes and the Fine Structure between NC^1 and L

Samir Datta^a, Meena Mahajan^b, B V Raghavendra Rao^c, Michael Thomas^d,
Heribert Vollmer^d

^a *Chennai Mathematical Institute, India.*

^b *The Institute of Mathematical Sciences, Chennai, India.*

^c *Universität des Saarlandes, Saarbrücken, Germany.*

^d *Leibniz Universität, Hannover, Germany.*

Abstract

The class NC^1 of problems solvable by bounded fan-in circuit families of logarithmic depth is known to be contained in logarithmic space L , but not much about the converse is known. In this paper we examine the structure of classes in between NC^1 and L based on counting functions or, equivalently, based on arithmetic circuits. The classes PNC^1 and $\text{C}_{=}\text{NC}^1$, defined by a test for positivity and a test for zero, respectively, of arithmetic circuit families of logarithmic depth, sit in this complexity interval. We study the landscape of Boolean hierarchies, constant-depth oracle hierarchies, and logarithmic-depth oracle hierarchies over PNC^1 and $\text{C}_{=}\text{NC}^1$. We provide complete problems, obtain the upper bound L for all these hierarchies, and prove partial hierarchy collapses. In particular, the constant-depth oracle hierarchy over PNC^1 collapses to its first level PNC^1 , and the constant-depth oracle hierarchy over $\text{C}_{=}\text{NC}^1$ collapses to its second level.

1. Introduction

The class NC^1 occupies a special place in the study of complexity classes inside P , owing to its robustness and multiple characterizations. It is defined as the class of languages accepted by families of circuits of polynomial size and logarithmic depth using bounded fan-in Boolean gates. By uniform NC^1 we mean the subclass where the circuit families have succinct descriptions: given a size parameter in unary, the circuit for that size from the family can be “easily” computed. Various notions of uniformity are known to give rise to the same class of languages, also coinciding with the class of languages accepted by logarithmic-time alternating machines ALOGTIME . Other characterizations of NC^1 include

Email addresses: sdatta@cmi.ac.in (Samir Datta), meena@imsc.res.in (Meena Mahajan), bvrr@cs.uni-sb.de (B V Raghavendra Rao), thomas@thi.uni-hannover.de (Michael Thomas), vollmer@thi.uni-hannover.de (Heribert Vollmer)

Supported in part by the Indian DST and the German DAAD.

polynomial-sized formulas, bounded-width branching programs, bounded-width circuits and programs over finite monoids.

It is known that all NC^1 languages can be accepted in logarithmic space L , but it is not known whether this containment is strict. All L -complete languages are candidates for membership in L but not in NC^1 , and most of these candidates lie in classes defined using the natural counting classes associated with NC^1 , namely, $\#\text{NC}^1$ and GapNC^1 . The former counts “proving sub-circuits” in an NC^1 circuit (see Section 2 for formal definitions); the latter is its closure under subtraction. It is not yet known whether these functions can be evaluated in NC^1 , although the best upper bound is very very close (an $O(\log^*)$ factor in depth). It is known that functions in $\#\text{NC}^1$ and GapNC^1 can be evaluated in function logarithmic space FL ; thus languages definable by applying simple predicates to such functions are also in L . The natural choices of predicates are a test for zero and a test for positivity, giving rise to the language classes $\text{C}=\text{NC}^1$ and PNC^1 sitting between NC^1 and L . (There are also predicates testing for zero modulo a fixed prime; the resulting language classes are already known to coincide with NC^1 .) A nice survey of these classes can be found in [All04].

It is not clear how much structure is there between NC^1 and L in the event that the classes are distinct. In this note, we attempt to explore the structure between NC^1 and L , based on hierarchies of language classes built upon $\text{C}=\text{NC}^1$ and PNC^1 . For a complexity class \mathcal{C} , there are three standard ways of defining the hierarchies above \mathcal{C} : the Boolean hierarchy $\text{BH}(\mathcal{C})$, the constant-depth hierarchy using oracle gates $\text{AC}^0(\mathcal{C})$, and the NC^1 -oracle-gate hierarchy $\text{NC}^1(\mathcal{C})$, with $\text{BH}(\mathcal{C}) \subseteq \text{AC}^0(\mathcal{C}) \subseteq \text{NC}^1(\mathcal{C})$.

Our results. As a first step in our study, we describe the oracle hierarchies in terms of arithmetic circuits augmented with test gates. These are the arithmetic-Boolean circuits defined in [vzGS91]; with size and depth restrictions as in NC^1 , and with test gates for “= 0?” or “> 0?”, we obtain the classes a-NC^1_{\leq} and $\text{a-NC}^1_{>}$. We observe that if each path in the circuit has $O(1)$ test gates, then a-NC^1_{\leq} and $\text{a-NC}^1_{>}$ coincide with $\text{AC}^0(\text{C}=\text{NC}^1)$ and $\text{AC}^0(\text{PNC}^1)$ respectively (Proposition 11). However, there is a subtlety in similarly characterizing $\text{NC}^1(\text{C}=\text{NC}^1)$ and $\text{NC}^1(\text{PNC}^1)$. We introduce a syntactic restriction on the arithmetic-Boolean circuits giving rise to a reasonable definition, and show that (1) the classes so defined coincide with $\text{NC}^1(\text{C}=\text{NC}^1)$ and $\text{NC}^1(\text{PNC}^1)$ (Proposition 13), and (2) as expected, are indeed contained in L (Theorem 17). On the other hand, without this restriction, the best upper bound we can show for the arithmetic circuit hierarchy is the complexity class TC^1 (Theorem 18), which subsumes L and even nondeterministic logspace NL , but is contained in NC^2 .

Next, we show that the constant-depth hierarchy over PNC^1 (and hence also the Boolean hierarchy) collapses to PNC^1 (Theorem 19). We adapt the proof of [Ogi98], where an analogous result for PL is shown. One difficulty in the adaptation is showing the required normal form for GapNC^1 circuits. We use the equivalent characterization of GapNC^1 as arithmetic bounded-width branching programs GapWBWP , and establish the normal form here. Another difficulty is computing an exponential sum; we use the notion of read-once certified circuits

and read-once exponential sums, introduced in [MR09], to carry the proof through.

Finally, we examine the hierarchies over $C=NC^1$. Since $C=NC^1$ is not even known to be closed under complementation, we do not expect a collapse all the way down. Our first result is a characterization of the Boolean hierarchy over $C=NC^1$ as the class of languages described by checking feasibility of small systems of linear equations, where the coefficients themselves are GapNC^1 -computable functions of the input word (Theorem 34). Our second result is that the constant-depth hierarchy over $C=NC^1$ collapses to a class slightly weaker than the second level (Theorem 41). Both these results appear as analogues of known results [ABO99] for the corresponding logarithmic-space class $C=L$, but require substantially different proofs.

Also, unlike in the case of PL and $C=L$, our results do not seem to go through for the NC^1 -hierarchies over PNC^1 and $C=NC^1$.

2. Background

For any language L , χ_L denotes its characteristic function: $\chi_L(x) := 1$ if $x \in L$, $\chi_L(x) := 0$ if $x \notin L$.

Boolean circuits and language classes. We denote by L the class of languages accepted by deterministic logarithmic-space Turing machines.

We consider Boolean circuits with internal gates labelled \vee , \wedge , or \neg . By NC^1 we denote the class of languages which can be accepted by a family $\{C_n\}_{n \geq 0}$ of Boolean circuits of polynomial size whose depth is bounded by $O(\log n)$, with each gate having constant fan-in. The class AC^0 denotes the set of languages accepted by a Boolean circuit family $\{C_n\}_{n \geq 0}$ of polynomial size and constant depth, with unbounded fan-in. Without loss of generality, we can assume that negation gates appear only at the leaves of the circuit. Also, without loss of generality we can assume that AC^0 and NC^1 circuits are actually formulas: every gate has out-degree one. An NC^0 circuit is a Boolean circuit, or formula, of constant size, with each gate having constant fan-in. We denote by AC^0_k (respectively) NC^0_k the polynomial size (respectively, constant size) circuit families of depth at most k .

By TC^0 and TC^1 we denote the class of languages decided by circuit families of polynomial size and constant (respectively, logarithmic) depth, where each gate is either a negation gate or an unbounded fan-in majority gate: it outputs 1 if and only if more than half of its inputs are 1. Integer addition and multiplication are known to be in TC^0 .

A branching program (BP for short) is a layered acyclic graph G with edges labelled by constants (0 or 1) or literals, and with two special vertices s and t . It accepts an input x if there is an $s \rightsquigarrow t$ path where each edge is labelled by a true literal or the constant 1; we call such a path an *accepting path* on input x . BWBP denotes the class of languages that can be accepted by families of polynomial size bounded width branching programs $\{G_n\}_{n \geq 0}$, where the graph G_n considers n variables. It is known that BWBP equals NC^1 ([Bar89]). Restricted to uniform

circuits (with appropriate notions of uniformity, see for instance [Vol99]), it is known that $\text{NC}^1 = \text{BWBP} \subseteq \text{L}$.

Proposition 1 (Known containments).

$$\text{AC}^0 \subseteq \text{TC}^0 \subseteq \text{NC}^1 = \text{BWBP} \subseteq \text{L} \subseteq \text{TC}^1 \subseteq \text{DSPACE}(\log^2 n) \cap \text{P}.$$

Arithmetic circuit classes. For the purposes of this paper, an arithmetic circuit is a circuit where the gates are labelled from the set $\{+, \times, -1, 0, 1, x_1, \dots, x_n\}$. The gates $+$ and \times are the addition and multiplication operations over \mathbb{Z} . Such a circuit computes a function $f : \{0, 1\}^n \rightarrow \mathbb{Z}$.

An a-NC^1 circuit family $\{C_n\}_{n \geq 0}$ is a family of bounded fan-in arithmetic circuits where for each n , C_n is of size polynomial in n , depth logarithmic in n , and computes a function $f_n : \{0, 1\}^n \rightarrow \mathbb{Z}$. The family computes the function $f : \{0, 1\}^* \rightarrow \mathbb{Z}$ where $f(x) := C_{|x|}(x)$. GapNC^1 is the class of functions computed by a-NC^1 circuit families. The analogous arithmetic class for constant-depth unbounded fan-in circuits is denoted by a-AC^0 .

An arithmetic branching program is a BP B where edges are labelled by literals or constants from the set $\{-1, 0, 1\}$. For an $s \rightsquigarrow t$ path P , let $wt(P(a))$ denote the product of all the edge labels in P under the assignment a . Then the function computed by B is defined as follows:

$$\text{for all } a \in \{0, 1\}^n \quad f(a) := \sum_{P \text{ is an } s \rightsquigarrow t \text{ path in } B} wt(P(a))$$

An a-BWBP family $\{B_n\}_{n \geq 0}$ is a family of arithmetic branching programs of polynomial size and bounded width. GapBWBP is the class of functions computed by a-BWBP program families.

For a Boolean (no edge labelled -1) BP B and an input assignment a , let $\#[s \rightsquigarrow t](a)$ denote the the number of $s \rightsquigarrow t$ paths in B under the assignment a . $\#\text{BWBP}$ is the class of functions $\{0, 1\}^* \rightarrow \mathbb{N}$ computed by BWBP. The class DiffBWBP is the closure of $\#\text{BWBP}$ under finite subtractions; For a Boolean (no edge labelled -1) BP B and an input assignment a , let $\#[s \rightsquigarrow t](a)$ denote the the number of $s \rightsquigarrow t$ paths in B under the assignment a . $\#\text{BWBP}$ is the class of functions $\{0, 1\}^* \rightarrow \mathbb{N}$ computed by BWBP. The class DiffBWBP is the closure of $\#\text{BWBP}$ under finite subtractions; $\text{DiffBWBP} = \{f - g \mid f, g \in \#\text{BWBP}\}$.

The above three classes coincide:

Proposition 2 ([CMTV98]). $\text{GapNC}^1 = \text{GapBWBP} = \text{DiffBWBP}$.

We will often use the following equivalent form for GapNC^1 functions: for any GapNC^1 function f , there is a BWBP B with start node s , two target nodes t_1 and t_2 , and $f(a) = \#[s \rightsquigarrow t_1](a) - \#[s \rightsquigarrow t_2](a)$. We say that B *gap-represents* the function f .

It is known that NC^1 circuits can be made unambiguous [Lan93]. In terms of arithmetic circuits, this yields:

Proposition 3. *Let L be any AC^0 (or NC^1) language. Then there is an a-AC^0 (a-NC^1 , respectively) circuit family C that does not use the constant -1 such that for each string w , $C(w) = \chi_L(w)$.*

The classes $C_{=}NC^1$ and PNC^1 which are the central to this paper are defined as follows.

Definition 4.

$$C_{=}NC^1 := \left\{ L \in \{0,1\}^* \mid \begin{array}{l} \text{for some } f \in \text{GapNC}^1, \text{ for all } x \in \{0,1\}^*, \\ x \in L \text{ if and only if } f(x) = 0. \end{array} \right\}$$

$$PNC^1 := \left\{ L \in \{0,1\}^* \mid \begin{array}{l} \text{for some } f \in \text{GapNC}^1, \text{ for all } x \in \{0,1\}^*, \\ x \in L \text{ if and only if } f(x) > 0. \end{array} \right\}$$

Proposition 5 ([CMTV98]). 1. $NC^1 \subseteq C_{=}NC^1 \subseteq PNC^1 \subseteq L$.
 2. $C_{=}NC^1$ is closed under union and intersection.
 3. PNC^1 is closed under union, intersection and complementation.

Arithmetic-Boolean circuits. Let a *test gate* for “=0?” (respectively “>0?”) be a unary gate that outputs 1 if its input is equal to 0 (respectively greater than 0) and 0 otherwise. Define an $a\text{-NC}_{=}^1$ circuit (respectively $a\text{-NC}_{>}^1$ circuit) to be an arithmetic circuit of logarithmic depth and polynomial size over Boolean input gates, binary +- and \times -gates, constants $-1, 0$ and 1 as well as test gates for “=0?” (respectively “>0?”). From the definitions, it follows that

Proposition 6. A language L is in $C_{=}NC^1$ (or PNC^1) if and only if χ_L can be computed by an $a\text{-NC}_{=}^1$ (respectively $a\text{-NC}_{>}^1$) circuit family in which each circuit has exactly one test gate appearing as the output gate.

Read-once certificates. Let B be a branching program on variables $X = \{x_1, \dots, x_n\} \cup Y = \{y_1, \dots, y_m\}$. B is said to be *read-once certified* in Y if there are indices $i_0 = 1 < i_1 < i_2 < \dots < i_m$ such that variable y_j appears only between layer i_{j-1} and i_j . By specialising the arguments in [MR09] to the counting classes, we can compute exponential sums over the variables in Y efficiently.

Proposition 7 (adapted from Theorem 3(2), [MR09]). Let $f(X, Y)$ be a function computed by an $a\text{-BP}$ B of size s and width w , read-once certified in Y . Let g be the function defined as

$$g(X) := \sum_{e \in \{0,1\}^m} f(X, e)$$

Then g can be computed by an $a\text{-BP}$ of size $\text{poly}(s)$ and width $O(w^2)$. Hence, if $f \in \text{GapBWBP}$, then $g \in \text{GapBWBP}$.

Miscellaneous. We denote by $C_1 \cdot C_2$ a circuit which can be split horizontally into two parts, with the top part being a circuit of type C_1 , and all its inputs being either circuit inputs (literals or constants) or circuits of type C_2 . We denote by $[C]$ an oracle gate for a language in C .

Thus $[C] \cdot AC^0$ is the class of all languages accepted by $AC^0(C)$ oracle circuits such that each circuit has a single oracle gate at the output, and each input bit to the oracle gate is the output of an AC^0 sub-circuit.

3. Hierarchies

3.1. Defining the Hierarchies

Among the simplest is the Boolean hierarchy, which characterizes the languages expressible as Boolean combinations of any constant number of languages from respectively $C_{=NC}^1$ or PNC^1 .

Definition 8 (The Boolean Hierarchy). Let \mathcal{C} be a complexity class. The *Boolean hierarchy over \mathcal{C}* is defined as the set of languages L for which there exists an NC^0 circuit C with k inputs and $A_1, \dots, A_k \in \mathcal{C}$ such that for all $x \in \{0, 1\}^*$,

$$x \in L \iff C(\chi_{A_1}(x), \chi_{A_2}(x), \dots, \chi_{A_k}(x)) = 1$$

We denote this class of languages by $NC^0 \cdot \mathcal{C}$ or $BH(\mathcal{C})$.

Remark 9. A perhaps more standard way of defining the Boolean hierarchy is to define the levels $BH_0(\mathcal{C}) := \mathcal{C}$, and $BH_i(\mathcal{C}) := \{L_1 \Delta L_2 \mid L_1, L_2 \in BH_{i-1}(\mathcal{C})\}$, and then take the union $\bigcup_{i \geq 0} BH_i(\mathcal{C})$. If \mathcal{C} is closed under union and intersection, then these definitions coincide with each other and with the definition of $NC^0 \cdot \mathcal{C}$ above ([KSW87]).

The other way of defining hierarchies is via oracle queries. As shown in [AO96] (see also [ABO99]), nesting queries above a base machine is equivalent to adding oracle gates in an AC^0 circuit. And in many cases, it also turns out to be equivalent to adding oracle gates in an NC^1 circuit. We present the oracle-circuit definitions, first introduced by Wilson [Wil85], below.

Let L be any language. An $AC^0(L)$ *circuit family* is a sequence $\{C_n\}_{n \geq 0}$ of AC^0 circuits containing additional oracle gates for L of unbounded fan-in. Similarly, an $NC^1(L)$ *circuit family* is a sequence $\{C_n\}_{n \geq 0}$ of NC^1 circuits with additional oracle gates for L of unbounded fan-in such that oracle gates of fan-in m account for depth $\lceil \log m \rceil$.

Definition 10 (The AC^0 and the NC^1 Hierarchy). Let \mathcal{C} be a complexity class. Then $AC^0(\mathcal{C})$ (respectively $NC^1(\mathcal{C})$) is defined to comprise those problems decidable by an $AC^0(L)$ (respectively $NC^1(L)$) circuit family for some $L \in \mathcal{C}$.

3.2. Characterizing the Hierarchies using Arithmetic-Boolean Circuits

From Proposition 6, we know that $C_{=NC}^1$ and PNC^1 have equivalent Arithmetic-Boolean circuits. It is natural to ask whether there are equivalent such circuits for the hierarchies above these classes. For the AC^0 hierarchy, this is easy to see; we show below that $AC^0(C_{=NC}^1)$ and $AC^0(PNC^1)$ can be characterized using arithmetic-Boolean circuits. We need the notion of nesting depth: in a circuit C , the nesting depth of gates of a type t is the largest number k such that some path from the output to a leaf of C goes through exactly k gates of type t .

Proposition 11. $AC^0(C_{=NC}^1)$ (respectively $AC^0(PNC^1)$) equals the class of languages decidable by $\mathbf{a}\text{-}NC_{=}^1$ (respectively $\mathbf{a}\text{-}NC_{>}^1$) circuit families such that the nesting depth of test gates is bounded by a constant and the output gate of each circuit is a test gate.

Proof. We will consider the case $C=NC^1$ only, the proof for PNC^1 is completely analogous.

For the direction from left to right, let $L \in C=NC^1$ and denote by f the function witnessing this fact (that is, $x \in L \iff f(x) = 0$ for all x). Let C be an unambiguous $AC^0(L)$ circuit with n inputs (without loss of generality we may assume that C is unambiguous). From C , construct an arithmetic circuit C' as follows:

- Replace each \wedge -gate with a \times -gate.
- Replace each \neg -gate with input x with the sub-circuit $1 + (-1 \times x)$.
- Replace each \vee -gate with inputs x_1 and x_2 with the sub-circuit $x_1 + ((1 + (-1 \times x_1)) \times x_2)$.
- Replace each oracle gate with inputs x_1, \dots, x_m with a test gate whose input is the arithmetic circuit that computes $f(x_1 \cdots x_m)$.

It holds that $C'(x) \in \{0, 1\}$ and $C'(x) = C(x)$ for all possible inputs x . The size of C' is clearly polynomial in the size of C . Its depth is $O(\log n)$ owing to the replacement of an oracle gate by a test gate atop an $a-NC^1$ circuit; however, the nesting gate of test gates is at most the depth of C and hence a constant. Thus, the circuit D deciding whether $C'(x) + (-1)$ is equal to zero is the desired $a-NC^1_{\leq}$ circuit.

For the converse direction, let C denote an $a-NC^1_{\leq}$ circuit with a test gate at the output and $O(1)$ nesting depth of test gates. Without loss of generality, we can assume that the circuit is a formula: every gate has out-degree 1. If we cut all the edges leading out of a test gate, the circuit breaks up into blobs, each of which is an arithmetic circuit with a test gate at the output. In particular, let g_1, \dots, g_m enumerate all test gates in C , and denote by S_i the maximal connected sub-circuit of C rooted at g_i that consists of $+$, \times -gates and the constants $-1, 0, 1$. Then each S_i computes some function $f_i: \{0, 1\}^m \rightarrow \{0, 1\}$. By the structure of the circuit, using Proposition 6, we see that these functions f_i are all characteristic functions of $C=NC^1$ languages. Replacing S_i with an oracle gate for the corresponding language yields a Boolean circuit C' of polynomial size comprising input gates and oracle gates for $C=NC^1$ only. The depth of this circuit is the nesting depth of test gates in C and hence a constant. This is the required $AC^0(C=NC^1)$ circuit. \square

Remark 12. A small technicality in the above proof: in going from right to left, different oracle gates appear to query different languages in $C=NC^1$. However, these can all be replaced by queries to any one language that is complete for $C=NC^1$ under projections.

It is tempting to believe that dropping the requirement on nesting depth of test gates will characterize $NC^1(C=NC^1)$ and $NC^1(PNC^1)$. This, however, is not the case. The conversion from left to right ($NC^1(C=NC^1)$ to $a-NC^1_{\leq}$) goes through, but for the converse, the requisite depth bound does not follow. We describe a certain condition under which we can obtain an exact characterization.

Let C be an $\mathbf{a-NC}_{=}^1$ circuit (respectively $\mathbf{a-NC}_{>}^1$ circuit) with n inputs and let g_1, \dots, g_m enumerate all of its test gates. Denote by S_i the maximal connected sub-circuit of C rooted at g_i that consists of $+$, \times -gates and the constants -1 , 0 , 1 ; these are the blobs in the proof of Proposition 11. As the depth of C is logarithmic in the number of its inputs, we may without loss of generality assume that S_1, \dots, S_m induce a partition of the non-input gates of C . Thus any path from the output to a leaf in C goes through a chain of these blobs. There can be $O(\log n)$ blobs on any such chain, and the logarithm of the size of a blob can be as large as $\theta(\log n)$, and this causes the problem in replicating the above proof. We “define away” the problem: We say that C has the *small-blob-chains* property if for every path π from the root of C to an input gate or a constant,

$$\sum_{g_i \text{ occurs in } \pi} \log |S_i| \in O(\log n).$$

Now we can show that the \mathbf{NC}^1 hierarchies above $\mathbf{C}_{=}\mathbf{NC}^1$ and \mathbf{PNC}^1 are characterized exactly.

Proposition 13. $\mathbf{NC}^1(\mathbf{C}_{=}\mathbf{NC}^1)$ (respectively $\mathbf{NC}^1(\mathbf{PNC}^1)$) equals the class of languages decidable by $\mathbf{a-NC}_{=}^1$ (respectively $\mathbf{a-NC}_{>}^1$) circuit families with the small-blob-chains property in which the output gate of each circuit is a test gate.

Proof. We will consider the case $\mathbf{C}_{=}\mathbf{NC}^1$ only, the proof for \mathbf{PNC}^1 is completely analogous. The constructions are exactly as in the proof of Proposition 11, only the analysis is different.

For the direction from left to right, let C' be the circuit obtained from the unambiguous $\mathbf{NC}^1(L)$ circuit C . Without loss of generality assume that C' is a formula. The size of C' is clearly polynomial in the size of C . As in C any oracle gate of fan-in m accounts for depth $\log m$ and the depth of C bounded by $O(\log n)$, we obtain that the depth C' remains bounded by $O(\log n)$ and also satisfies the small-blob-chain property. Thus, the circuit D deciding whether $C'(x) + (-1)$ is equal to zero is the desired $\mathbf{a-NC}_{=}^1$ circuit.

For the converse direction, let C denote an $\mathbf{a-NC}_{=}^1$ circuit and let C' be the Boolean circuit constructed from C as in Proposition 11; it is of polynomial size comprising input gates and oracle gates for $\mathbf{C}_{=}\mathbf{NC}^1$ only. From the small-blob-chains property, we finally obtain that for each path π' starting in the root of C' and leading to an input gate or a constant, $\sum_{g \in \pi'} \log(\text{fan-in}(g)) = \sum_{S \in \pi} O(\log |S|) \in O(\log n)$, where π is the corresponding path in C going through blob S instead of oracle gate g . Thus the depth of C' remains bounded by $O(\log n)$. \square

There exist arithmetic-Boolean circuits violating the small-blob-chains property. Consider, for example, any circuit family $\{C_n\}_{n \geq 0}$ such that C_n contains n input gates, $\log n$ test gates, and the input to each of these test gates is a binary gate g with inputs i_1 and i_2 such that the sub-circuits computing i_1 and i_2 are disjoint, the sub-circuit computing i_1 contains $\geq n$ gates, and the sub-circuit computing i_2 is rooted at a test gate (see Fig. 1). Then the path starting in




Figure 1: A circuit violating the small-blob-chains property.

the root of circuit C_n and proceeding to the right ancestor of every binary gate crosses all maximal sub-circuits consisting of all but test gates rooted at test gates that consist of $+$, \times -gates as well as the constants $-1, 0, 1$. Thus,

$$\sum_{S_i \text{ occurs in } \pi} \log |S_i| \geq \log n \cdot \log n \notin O(\log n).$$

Hence, dropping the small-blob-chains property from the circuits in Proposition 13 leads to presumably different class of languages. We denote these classes by AH , for arithmetic hierarchy, defined analogously to the classes figuring in Proposition 6 and Proposition 13.

Definition 14 (Arithmetic Circuit Hierarchies over C=NC^1 and PNC^1). A language L is said to be in $\text{AH}(\text{C=NC}^1)$ (or $\text{AH}(\text{PNC}^1)$) if and only if χ_L can be computed by an a-NC^1_{\leq} (respectively $\text{a-NC}^1_{>}$) circuit family such that in each circuit, the output gate is a test gate.

The following chain of inclusions holds.

Observation 15.

$$\begin{array}{ccccccccc} \text{C=NC}^1 & \subseteq & \text{BH}(\text{C=NC}^1) & \subseteq & \text{AC}^0(\text{C=NC}^1) & \subseteq & \text{NC}^1(\text{C=NC}^1) & \subseteq & \text{AH}(\text{C=NC}^1) \\ \cap & & \cap & & \cap & & \cap & & \cap \\ \text{PNC}^1 & \subseteq & \text{BH}(\text{PNC}^1) & \subseteq & \text{AC}^0(\text{PNC}^1) & \subseteq & \text{NC}^1(\text{PNC}^1) & \subseteq & \text{AH}(\text{PNC}^1). \end{array}$$

Remark 16. We can also augment the a-NC^1_{\leq} and $\text{a-NC}^1_{>}$ circuits in Definition 14 by allowing oracle gates, with $\lceil \log(\text{fan-in}(g)) \rceil$ charged to the depth of each such gate g . Since, without loss of generality, we deal with languages over a binary alphabet, the inputs to the oracle gate must be Boolean inputs. But the circuit computes arithmetic values, except at test gates. Thus, we will require that all the inputs to an oracle gate are either Boolean circuit inputs (literals or the constants 0,1, but not -1) or the outputs of test gates. It can be shown that

allowing $C=NC^1$ oracle gates in $a-NC^1_{\leq}$ circuits, or PNC^1 oracle gates in $a-NC^1_{>}$ circuits, with this condition, does not add to the power of the circuit families beyond $AH(C=NC^1)$ and $AH(PNC^1)$ respectively.

3.3. Upper bounds

We first show that the AC^0 and the NC^1 hierarchies over $C=NC^1$ and PNC^1 are contained in L . By the containments depicted in Observation 15, it suffices to show this bound for $NC^1(PNC^1)$.

Theorem 17. $NC^1(PNC^1) \subseteq L$.

Proof. Let L be in $NC^1(PNC^1)$ and suppose without loss of generality that $L \subseteq \{0,1\}^*$. Then there exists a language B in PNC^1 and an $NC^1(B)$ circuit family $\{C_n\}_{n \geq 0}$ such that $C_{|x|}(x) = 1$ if and only if $x \in L$. As apparent from the proof of Proposition 13, we may without loss of generality assume that C_n contains oracle gates and input gates only.

Given x , we now simulate the circuit $C_{|x|}$ in a top-down way. At any (oracle) gate for B , we simulate the PNC^1 circuit for B , using the fact that PNC^1 is known to be in L (Proposition 5). If this simulation requests an input g , we recursively simulate the sub-circuit of C rooted at g to obtain this bit.

The amount of space required by this approach is $O(\log n)$ for the paths in C plus the space needed for the simulations. As $PNC^1 \subseteq L$, each simulation requires space logarithmic in the fan-in of the oracle gate. Let $s(g)$ denote the space needed to simulate the circuit associated with the (oracle) gate g , assuming that its inputs are explicitly available. Then, for each path π starting in the root of C and leading to an input gate or a constant, we have $\sum_{g \in \pi} s(g) = \sum_{g \in \pi} O(\log(\text{fan-in}(g))) \in O(\log n)$. Thus the entire recursive simulation requires $O(\log n)$ space, and so we conclude that $NC^1(PNC^1)$ is contained in L . \square

Recall from Proposition 13 that $NC^1(PNC^1)$ can be described in terms of $a-NC^1_{\leq}$ circuits with the small-blob-chains property. The proof above can be restated assuming that the equivalent $a-NC^1_{\leq}$ circuit family is given, rather than the oracle circuit. We describe this restatement here, primarily to highlight why it does not work for $AH(PNC^1)$.

Alternative proof of Theorem 17. Let C be the $a-NC^1_{>}$ circuit with test gates placed so as to satisfy the small-blob-chains property. Let it have n input gates, and let $x \in \{0,1\}^n$. We use the facts that PNC^1 is known to be in L (Proposition 5), and that PNC^1 is characterized by $a-NC^1_{>}$ circuits with a single test gate at the output (Proposition 6).

As in the proof of Proposition 11, we assume that the circuit is a formula, and cut all the edges leading out of test gates, so that the circuit breaks up into blobs, each of which is an arithmetic circuit with a test gate at the output. If g_1, \dots, g_m enumerate all test gates in C , and S_i the maximal connected sub-circuit of C rooted at g_i with no test gate other than g_i , then each S_i computes $f_i = \chi_{L_i}$ for some $L_i \in PNC^1$, and so each f_i is computable in L .

Given x , we evaluate the circuit in a top-down way, starting with the topmost blob and carrying out the logspace evaluation of the associated function. When this evaluation needs a bit that is an input to this blob and is an output of another test gate, we recursively starts a fresh logspace computation of the required function. When this computation terminates, we resume the earlier computation.

At any stage, the computation traces out a path, or a chain, through the blobs in C , with the current computation focussing on the leaf of the chain. The space requirement per blob is logarithmic in the size of the blob. By the small-blob-chains property, we conclude that the overall space requirement is $O(\log n)$. \square

By Proposition 13, the arithmetic hierarchy over PNC^1 differs from the NC^1 hierarchy only in the small-blob-chains property. In the absence of this property, the recursive simulation in the second proof above yields only a $O(\log^2 n)$ space bound. Also, since the log-space evaluation of each blob may not be read-once in its inputs, each blob may have to be evaluated several times. So we cannot obtain a polynomial time bound for the recursive procedure.

However, using a bottom-up evaluation, we can show that $\text{AH}(\text{PNC}^1)$ circuits can be evaluated in TC^1 .

Theorem 18. $\text{AH}(\text{PNC}^1) \subseteq \text{TC}^1$.

Proof. We perform a straightforward bottom-up Boolean evaluation of the arithmetic circuit; that is, we evaluate bit representations of the values carried on the wires. Since C is a formula, all intermediate values have polynomial-sized bit representations, so the $+$ and \times gates can be replaced by TC^0 sub-circuits. The test gates can also be trivially replaced by appropriate TC^0 sub-circuits. This yields a log depth circuit with majority gates, that is, a TC^1 circuit. \square

4. The PNC^1 hierarchy collapses

In this section we show that the constant-depth PNC^1 hierarchy, $\text{AC}^0(\text{PNC}^1)$, collapses to the base level.

Theorem 19. $\text{AC}^0(\text{PNC}^1) = \text{PNC}^1$.

Proof. Since PNC^1 is closed under complementation, and since unbounded fan-in \vee and \wedge functions are in NC^1 and hence in PNC^1 , we can assume without loss of generality that the $\text{AC}^0(\text{PNC}^1)$ circuit has only oracle gates. Theorem 20 below shows how to collapse two adjacent levels of PNC^1 oracle gates into one. Applying this repeatedly gives the desired result. \square

The rest of this section is devoted to proving Theorem 20:

Theorem 20. $[\text{PNC}^1] \cdot [\text{PNC}^1] = \text{PNC}^1$.

We adapt the techniques of [Ogi98] to the case of constant width branching programs. Also, as in [Ogi98], we use the polynomial technique developed earlier ([BRS95, FR96]) to show closure properties of the complexity class PP. A new ingredient we require is the notion of read-once certified circuits and exponential sums from [MR09].

4.1. Overview of the Collapse Argument

Consider a language L in $\text{PNC}^1 \cdot \text{PNC}^1$. Then there is a language $H \in \text{PNC}^1$ and a circuit family $\{C_n\}$ accepting L where each C_n has depth 2 and has only oracle gates for H . That is, the output gate g is an oracle gate whose inputs are themselves oracle gates or literals or constants. Without loss of generality, we can assume that in fact all inputs to g are outputs of oracle gates. Let g have fan-in t . On input x , its inputs are $\chi_H(Y_1), \chi_H(Y_2), \dots, \chi_H(Y_t)$, where each Y_i is a projection (re-ordering of bits) of the input x .

Let f be the GapNC^1 function witnessing that $H \in \text{PNC}^1$. Then there is a a-BWBP family computing f . The idea is to consider the a-BWBP B for inputs of length t , say y_1, \dots, y_t , and try to replace each edge labelled y_i by a copy of the a-BWBP on Y_i . However, since Y_i is the input to an oracle gate, we want a 0-1 value for the sign of $f(Y_i)$, not the value of $f(Y_i)$ itself. If the sign function can be computed by a suitable polynomial function, then we can apply this function to each $f(Y_i)$ to get another GapNC^1 function. Unfortunately, the sign function cannot be represented in this fashion. However, it can be approximated by rational functions (ratios of polynomials); this approximation was first used in [BRS95], and later in [FR96] and [Ogi98]. We follow the presentation from [Ogi98] and describe the polynomials in Section 4.3.

To show that using such approximations is valid, we require that B satisfies a certain condition: All paths should have equal susceptibility to error, so as to not change the overall outcome. In particular, since a y_i edge label corresponds to using the output of an oracle gate, and since different oracle gates can have different errors, we will require that each path has exactly the same multiset of edge labels, independent of the input. This is a strong normal form. Such a normal form was required to collapse $\text{AC}^0(\text{PL})$ to PL, and was shown in [Ogi98]. We show a corresponding normal form for a-BWBPs in Section 4.2.

Finally, we need to show that there is a GapBWBP function h which has the same sign as the value of the a-BWBP B with the rational approximations in place. In [Ogi98], the analogous result is shown by describing an appropriate probabilistic log-space machine. In the GapBWBP setting, things are a bit more complicated since we have only $O(1)$ storage. We get around this by using the notion of exponential sums over read-once certified circuits, introduced in [MR09]. The GapBWBP family computing the desired h is described in Section 4.4, completing the proof of Theorem 20.

4.2. A Normal Form for PNC^1 and GapBWBP

We introduce a notation here. A node v in a BP B is called a *nondeterministic* node if there is an input assignment for which v has two out-edges labelled 1.

We show the following normal form for branching programs computing functions in GapNC^1 ; this is analogous to Lemma 3.1 in [Ogi98] for PL and #L functions.

Lemma 21. *Let f be a function in GapNC^1 . Then there exists a branching program Q of width $O(1)$ such that*

1. Q has a single start node s and two terminal nodes t_1 and t_2 ;
2. every path originating from s ends at either t_1 or t_2 and nowhere else;
3. any path of Q on any given input x contains exactly q nondeterministic nodes, where $q = q(n) \leq \text{poly}(n)$;
4. every edge is labelled by a literal y_i or $\neg y_i$;
5. on any input y , Q has exactly 2^q paths originating from s ;
6. $f = \#[s \rightsquigarrow t_1] - \#[s \rightsquigarrow t_2]$.

Proof. From Proposition 2, there is a BP P of width $w = O(1)$ and size $s = \text{poly}(n)$, with nodes s , t_1 and t_2 such that $f = \#[s \rightsquigarrow t_1] - \#[s \rightsquigarrow t_2]$. The edge labels in P are literals or the constant 1. We modify P so that

- Every node has out-degree 0,1, or 2.
- For each layer k , there is an index i_k such that, edges from layer k to $k+1$ are labelled from the set $\{1, y_{i_k}, \neg y_{i_k}\}$.

This can be achieved by doing necessary staggering of the program P . Copy all the nodes of a layer into new nodes, and then implement the edges according to their labels, taking one variable at a time. Repeat this process for all the layers. This ensures that every edge in a particular layer is labelled by a single variable, its negation or a constant. This will double the width and increase the size by a factor of wn . Now, in a similar way, we can ensure that the out-degree of every node is 0,1 or 2; this will double the width, and increase the size by a factor of w rather than wn , since all outgoing edges at a layer read the same variable. The resulting BP P' will thus have width bounded by $4w$ and size $O(w^2n \cdot s)$, where $s = \text{size}(P)$.

We create a new line (a path) called the “zero-gap” line starting from s . This line remains a single path until the last layer and forks out to both t_1 and t_2 , with all the edges being labelled by 1. Note that this line produces a zero-gap. To meet condition 2 in the lemma, for every node v , if v has an out-edge labelled y_i (respectively $\neg y_i$) and no out-edge labelled $\neg y_i$ (respectively y_i), then add an edge labelled $\neg y_i$ (respectively y_i) to the zero-gap line. Note that this process ensures condition 2 without changing the gap-function.

Recall that a node in P' is nondeterministic if it has two out-going edges labelled 1 on at least one input assignment to y (for example, a node with two outgoing edges labelled y_i and 1 respectively). A layer is called nondeterministic if at least one of the nodes in that layer is nondeterministic. In order to ensure that for every input y , the total number of paths originating from s remains the same, we make all the nodes in every nondeterministic layer nondeterministic, by adding necessary paths to the zero-line as follows: Consider a deterministic node v in a nondeterministic layer. Let y_i be the allowed variable label for this

layer. There are two possibilities: v has two out-edges, one labelled by y_i and the other labelled by $\neg y_i$ or v has a single out-edge labelled by 1. In both the cases, we add an edge to the zero-gap line, with label 1. In the case when v is a node already in the zero-gap line, we just add a parallel edge. Again, this construction does not alter the gap function.

Finally we eliminate the constant 1 on edge labels: replace an edge labelled 1 by parallel edges, one labelled y_i and the other labelled $\neg y_i$, where y_i is the variable at this layer.

Let Q denote the resulting BP. Let q be the number of nondeterministic layers of Q . Since we have two choices at every nondeterministic layer and only one choice at deterministic layers, on any input y , the number of distinct paths originating from s is exactly 2^q . As we have not changed the gap values in the whole process, this proves the lemma. \square

4.3. Rational Approximations for the Sign: Ogihara's Polynomials

In the following, we define the functions that will be needed in the construction. We follow the same notations from [Ogi98].

Definition 22 ([Ogi98]).

$$\begin{aligned}\mathcal{P}_m(z) &:= (z-1) \prod_{i=1}^m (z-2^i)^2, \\ \mathcal{Q}_m(z) &:= -(\mathcal{P}_m(z) + \mathcal{P}_m(-z)), \\ \mathcal{R}_{m,k}(z) &:= \left(\frac{2\mathcal{P}_m(z)}{\mathcal{Q}_m(z)} \right)^{2k}, \\ \mathcal{S}_{m,k}(z) &:= (1 + \mathcal{R}_{m,k}(z))^{-1}, \\ \mathcal{A}_{m,k}(z) &:= \mathcal{Q}_m(z)^{2k}, \text{ and} \\ \mathcal{B}_{m,k}(z) &:= \mathcal{Q}_m(z)^{2k} + (2\mathcal{P}_m(z))^{2k}.\end{aligned}$$

The following properties of $\mathcal{S}_{m,k}$ are proved in [Ogi98]:

Proposition 23 ([Ogi98]). *For $m, k \geq 1$, and for every z , the following holds:*

1. $\mathcal{S}_{m,k}(z) = \frac{\mathcal{A}_{m,k}(z)}{\mathcal{B}_{m,k}(z)}$,
2. $1 \leq z \leq 2^m \Rightarrow 1 - 2^{-k} \leq \mathcal{S}_{m,k} \leq 1$, and
3. $-2^m \leq z \leq -1 \Rightarrow 0 \leq \mathcal{S}_{m,k}(z) \leq 2^{-k}$.

Let f be a function from strings to non-zero integers, and let $\mu = \mu(|x|)$ be such that for all strings x , $|f(x)| \leq 2^{\mu(|x|)}$. Let B' be a function mapping a string x to a sequence Y_1, \dots, Y_p of p Boolean strings for some p , and let $\kappa = 2p + 1$. For $i \in [1, p]$, define the following functions:

$$\begin{aligned}S(x, i, 1) &:= \mathcal{S}_{\mu, \kappa}(f(Y_i)), \\ S(x, i, 0) &:= 1 - \mathcal{S}_{\mu, \kappa}(f(Y_i)) \\ \alpha(x, i, 1) &:= \mathcal{A}_{\mu, \kappa}(f(Y_i)), \\ \alpha(x, i, 0) &:= \mathcal{B}_{\mu, \kappa}(f(Y_i)) - \mathcal{A}_{\mu, \kappa}(f(Y_i)), \text{ and} \\ \beta(x, i) &:= \mathcal{B}_{\mu, \kappa}(f(Y_i)).\end{aligned}$$

For $w \in \{0, 1\}^p$, define

$$\begin{aligned}\tilde{S}(x, w) &:= \prod_{i=1}^p S(x, i, w_i), \\ \tilde{\alpha}(x, w) &:= \prod_{i=1}^p \alpha(x, i, w_i), \text{ and} \\ \tilde{\beta}(x) &:= \prod_{i=1}^p \beta(x, i).\end{aligned}$$

Let H be the language defined as $\{y \mid f(y) > 0\}$. Then

- Lemma 24** ([Ogi98]). 1. If $w_i = \chi_H(Y_i)$ for $1 \leq i \leq p$, then $1 - p2^{-\kappa} \leq \tilde{S}(x, w) \leq 1$. Otherwise, $0 \leq \tilde{S}(x, w) \leq 2^{-\kappa}$.
2. $\tilde{S}(x, w) = \frac{\tilde{\alpha}(x, w)}{\tilde{\beta}(x)}$.

4.4. The Collapse: Proof of Theorem 20

We use the following characterization of the class PNC^1 .

Proposition 25. A language L belongs to PNC^1 if and only if there is a function $f \in \text{GapNC}^1$ such that if $x \in L$ then $f(x) \geq 1$ and if $x \notin L$ then $f(x) \leq -1$.

We now complete the proof of Theorem 20. Let $L \in [\text{PNC}^1] \cdot [\text{PNC}^1]$. As described in Section 4.1, there is a GapNC^1 function f and a circuit family accepting L such that for every word x , $f(x) \neq 0$, and furthermore, $x \in L \Leftrightarrow f(b_1, \dots, b_t) > 0$, where $b_i = \chi_H(Y_i)$ and so $b_i = 1$ if $f(Y_i) > 0$; $b_i = 0$ otherwise. Each query string Y_i is obtained from x by a projection and is an oracle query at the bottom layer; b_i is the oracle reply.

Replace each b_i by a variable y_i and apply Lemma 21 to get a polynomial size branching program Q , with three special nodes s, t_1 , and t_2 , computing $f(Y)$ on t -bit inputs via the gap $f = \#[s \rightsquigarrow t_1] - \#[s \rightsquigarrow t_2]$. Note that for every layer k of Q , there is a variable $u_k \in Y$ such that the edges from layer k to layer $k + 1$ are labelled from the set $\{u_k, \neg u_k\}$. Note that all the u_k need not be distinct. Henceforth, we denote by y_k and Y_k the variable at layer k of Q and the corresponding query string, respectively. Without loss of generality we can assume that every layer is a nondeterministic layer. Let Q have p layers. Then every pair of bit-strings w, u , each of length p , uniquely represents a path in the BP Q , by considering the i th bit w_i of w as the query answer at the i th layer and i th bit u_i of u as the nondeterministic choice. For w, u , with $|w| = |u| = p$, define the Boolean function $e(w, u)$ as follows: $e(w, u) = 1$ if and only if the path of Q represented by the strings w and u is an accepting path (that is, it

terminates at t_1). Now define the following functions:

$$\begin{aligned} T(x) &:= \sum_{u,w \in \{0,1\}^p} e(w,u) \tilde{S}(x,w), \\ a(x) &:= \sum_{u,w \in \{0,1\}^p} e(w,u) \tilde{\alpha}(x,w), \text{ and} \\ h(x) &:= a(x) - 2^{p-1} \tilde{\beta}(x). \end{aligned}$$

By Lemma 24, we have $T(x) = a(x)/\beta(x)$. Using the properties of \tilde{S} we have:

Lemma 26 ([Ogi98]). *If $x \in L$ then $T(x) > 2^{p-1}$, and if $x \notin L$ then $T(x) < 2^{p-1}$. Hence, $x \in L$ if and only if $h(x) \geq 0$.*

Hence it is sufficient to prove that $h(x) \in \text{GapBWBP}$.

Since GapBWBP is closed under taking polynomially bounded sums and products, it follows easily that $\alpha(x, i, w_i)$ and $\beta(x, i)$ are in GapBWBP . We now show that $a(x) \in \text{GapBWBP}$. At this point, it is convenient to revert to the GapBWBP formulation rather than the DiffBWBP formulation used until now. We modify Q by adding one layer at the end with one node t , edge $t_1 \rightarrow t$ with weight 1, and edge $t_2 \rightarrow t$ with weight -1 ; call this a-BWBP Q' . It computes the same function gap represented by Q .

First we show that $e(W, U)$ can be computed by a constant width branching program which is read-once certified in the variables W and U . Let r be the width of Q' ; note that $r = O(1)$. We build a Boolean circuit C which is “read-once certified” in W and U (that is, there exists a partition of C into sub-circuits $C_1, \dots, C_{|U|+|W|}$ with the following property: For each input from $U \cup W$, there exists an index i such that all wires from the input gate lead to the sub-circuit C_i), has width $O(r)$, and computes $e(W, U)$. We proceed layer by layer. At the i th stage, C computes the index of the node v at layer i which is a part of the unique path in Q' represented by W and U . Note that this index requires $O(\log r)$ bits and hence C has those many gates as the output of this stage. Now the $(i+1)$ th stage computes the index of the node v' at layer $i+1$, which is uniquely defined given the index of v and bits w_i and u_i . Given the bit representation of the index of v , the index of v' depends only on w_i and u_i hence can be computed by a circuitry of size (hence width) $O(r)$. We can further split this circuitry into a substage that reads only w_i , followed by a substage that reads only u_i . After the final stage, the circuit outputs 1 if and only if the index corresponds to the node t_1 at the last layer. Clearly C computes $e(W, U)$ and is of size $O(r \cdot \text{size}(Q'))$ and width $O(r)$. Also from the description of C , it is clear that C is read-once certified in $W \cup U$, with the variables being read in the order $w_1, u_1, w_2, u_2, \dots, w_p, u_p$.

Using a standard subset construction (see [Bar89], Section 5), we obtain a BP B computing $e(W, U)$ so that:

- $\text{width}(B) = 2^{O(r)}$ and $\text{size}(B) = 2^{O(r)} \text{size}(C)$; and




Figure 2: Composing B_i , P_i , and B'_i

- B is read-once certified in $W \cup U$, with the variables being read in the order $w_1, u_1, w_2, u_2, \dots, w_p, u_p$.

Let B_i be the part that depends only on w_i , and let B'_i be the part that depends only on u_i . Now we construct an a-BWBP that represents the product $e(W, U)\tilde{\alpha}(x, W)$ and is read-once certified in the variables $W \cup U$. Let P_i be a constant width a-BWBP computing $\alpha(x, i, W_i)$; by Proposition 2 we may assume that each P_i is of width 3 and size $\text{poly}(p)$. We interleave the programs B_i , B'_i and P_i as follows. Starting with B , split each node layer common to B_i and B'_i into two node layers connected by a perfect matching. Replace each edge of the matching by a copy of P_i . See Figure 2. Let B' be the BP thus obtained. Clearly B' computes $e(W, U)\tilde{\alpha}(x, W)$ and is read-once certified in $W \cup U$. The size of B' is bounded by $\text{size}(B) + \text{width}(B) \sum \text{size}(P_i)$, and its width is $\text{width}(B) \max\{\text{width}(P_i)\}$; these are bounded by $2^{O(r)}\text{size}(Q')\text{poly}(p)$ and $2^{O(r)} \times 3$ respectively.

Notice that $a(x)$ is exactly the read-once exponential sum of $e(x, W, U)\tilde{\alpha}(x, W)$ over the variables $W \cup U$. Thus, applying Proposition 7 to B' , we get an a-BWBP B'' of size $\text{poly}(\text{size}(B'))$ and width $(\text{width}(B'))^2$ computing $a(x)$. Hence we obtain a GapBWBP bound for $a(x)$.

5. The Hierarchy above $C_{=NC}^1$

Since we do not even know if $C_{=NC}^1$ is closed under complementation, we cannot hope for a direct collapse of the hierarchies above $C_{=NC}^1$ all the way down to $C_{=NC}^1$. However, we show here two partial collapses. For the analogous class $C_{=L}$, it has been shown in [ABO99] that the hierarchy collapses to $L_{C_{=L}}$, and that testing feasibility of systems of linear equations FSLE is complete for this class. At the level of NC^1 , we show that the analogous situation splits into two counterparts. We define an appropriate non-trivial notion of constant-dimension FSLE and show that it is complete for the Boolean hierarchy over $C_{=NC}^1$, $BH(C_{=NC}^1)$. We then show that the constant-depth hierarchy over $C_{=NC}^1$, $AC^0(C_{=NC}^1)$, collapses to a certain level within the hierarchy that we denote $AC^0 \cdot C_{=NC}^1$; this is contained in the second level of the hierarchy.

5.1. The Boolean Hierarchy above $C_{=NC}^1$

Definition 27. For any $k \in \mathbb{N}$, and any class \mathcal{C} of functions from words to integers, the language class $FSLE^k[\mathcal{C}]$ is defined as follows: A language L belongs to the class $FSLE^k[\mathcal{C}]$ if there are functions $A_{ij} \in \mathcal{C}$ for $1 \leq i, j \leq k$ and a vector

$b \in \mathbb{Z}^k$ such that for each $w \in \{0, 1\}^*$, $w \in L$ if and only if the system $Az = b$ of linear equations in k variables z_j , where $A_{ij} = A_{ij}(w)$, has a feasible solution over the rationals. The class $\text{FSLE}^{bdd}[\mathcal{C}]$ is the union of $\text{FSLE}^k[\mathcal{C}]$ over all k .

Proposition 28. *The following two containments hold:*

$$\begin{aligned} \text{coC=NC}^1 &\subseteq \text{FSLE}^1[\text{GapNC}^1], \\ \text{C=NC}^1 &\subseteq \text{FSLE}^2[\text{GapNC}^1]. \end{aligned}$$

Proof. Given $g = g(x_1, \dots, x_n) \in \text{GapNC}^1$, and a word $a \in \{0, 1\}^n$, consider the system of equations $Az = b$, where

1. $A := (g(a))$, $b := (1)$, and z consists of a single variable z_1 . Clearly, the system is feasible if and only if $g(a) \neq 0$.
2. $A := \begin{pmatrix} g(a) & 0 \\ 1 & 0 \end{pmatrix}$, $b := \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, and z consists of two variables. Clearly, the system is feasible if and only if $g(a) = 0$. \square

In fact, we can prove something slightly stronger:

Lemma 29. $\text{NC}_d^0 \cdot \text{C=NC}^1 \subseteq \text{FSLE}^{3(2^{d+1}-1)}[\text{GapNC}^1]$.

Proof. We use the following constructions:

Claim 30. *If $(A_1, b_1), (A_2, b_2)$ are $\text{FSLE}^{k_1}[\text{GapNC}^1]$ - and $\text{FSLE}^{k_2}[\text{GapNC}^1]$ -instances respectively, then we can construct an instance of $\text{FSLE}^{k_1+k_2}[\text{GapNC}^1]$ -instance (A, b) which is feasible if and only if both $(A_1, b_1), (A_2, b_2)$ are feasible.*

Proof of Claim. Let

$$A := \begin{pmatrix} A_1 & O_1 \\ O_2 & A_2 \end{pmatrix}, \quad b := \begin{pmatrix} b_1 \\ b_2 \end{pmatrix},$$

where O_1, O_2 are all zeroes matrices of appropriate dimensions ($k_1 \times k_2$ and $k_2 \times k_1$ respectively). \diamond

Claim 31. $\text{FSLE}^k[\text{GapNC}^1] \subseteq \text{co-FSLE}^{k+1}[\text{GapNC}^1]$

Proof of Claim. We need to show that given an $\text{FSLE}^k[\text{GapNC}^1]$ -instance (A, b) , we can construct an $\text{FSLE}^{k+1}[\text{GapNC}^1]$ -instance (A', b') such that (A, b) is feasible if and only if (A', b') is infeasible. We essentially mimic the construction in [ABO99] to achieve this. We set

$$A' := \begin{pmatrix} A^T & O \\ b^T & 0 \end{pmatrix}, \quad b' := \begin{pmatrix} O \\ 1 \end{pmatrix},$$

where O is a $k \times 1$ column vector of all zeroes. For a proof of correctness see [ABO99]. \diamond

Assume that the NC^0 circuit consists of \wedge - and \vee -gates, with \neg -gates occurring only at the base level. This is to fix the notion of depth for such circuits.

We proceed by induction on the depth d of the NC^0 circuit. For the purposes of the construction, we eliminate all internal \vee gates, replacing them by equivalent sub-circuits using \wedge and \neg .

The base case $d = 0$ is considered in Proposition 28—notice that we do not include negation gates in our depth computation.

Depending on what the top gate is, we use one or both of the preceding claims. Notice that the dimension of the $\text{FSLE}^{bdd}[\text{GapNC}^1]$ -instance constructed at most doubles at an \wedge -gate. To simulate an \vee -gate, we need a negation, an \wedge , and then one more negation, so the the dimension of the $\text{FSLE}^{bdd}[\text{GapNC}^1]$ -instance goes from k to $2k + 3$.

Hence, we conclude that a depth- d NC^0 circuit over $\text{C}=\text{NC}^1$ or $\text{coC}=\text{NC}^1$ translates to an $\text{FSLE}^{bdd}[\text{GapNC}^1]$ -instance of dimension at most $2 \times 3(2^d - 1) + 3 = 3(2^{d+1} - 1)$. \square

Now we will show the converse. Let us fix some notation first. Given a $k \times k$ square matrix A , for $S, T \subseteq [k]$, let $A_{S,T}$ denote the square sub-matrix of A with exactly the rows in S and columns in T . Also denote by b_S the column vector of b with only the entries indexed in S . Also denote by $[A : b]$ the square matrix formed by adding the column b to A .

Lemma 32. $\text{FSLE}^k[\text{GapNC}^1] \subseteq \text{NC}^0_{3k} \cdot \text{C}=\text{NC}^1$.

Proof. Let $Az = b$ describe the system of linear equations on an input word, and let $r = \text{rank}(A)$. We observe the following.

Proposition 33. *The following are equivalent.*

1. For every $S, T \subseteq [k]$ such that $|S| = |T| = r$ and $\text{rank}(A_{S,T}) = r$, and for every $j \notin S$, $\det([A_{S',T} : b_{S'}]) = 0$, where $S' = S \cup \{j\}$.
2. For some $S, T \subseteq [k]$ such that $|S| = |T| = r$ and $\text{rank}(A_{S,T}) = r$, and for every $j \notin S$, $\det([A_{S',T} : b_{S'}]) = 0$, where $S' = S \cup \{j\}$.
3. The system $Az = b$ is feasible.

Proof. $1 \Rightarrow 2$: Obvious.

$2 \Rightarrow 3$: Let S, T be the subsets with the given property. Since $\text{rank}(A_{S,T}) = r$, $A_{S,T}$ is full-rank, and so the sub-system $A_{S,T}z_T = b_S$ has a unique solution z_T . Now extend this to a solution by setting $z = 0$ outside T . Clearly, this satisfies the equations indexed by S . By the condition $\det([A_{S',T} : b_{S'}]) = 0$ for $j \notin S$ and $S' = S \cup \{j\}$, it also satisfies every equation $j \notin S$. So it is a feasible solution.

$3 \Rightarrow 1$: Let z be a feasible solution, and let S, T be any pair of subsets of $[k]$ of size r such that $A_{S,T}$ is non-singular. If z is 0 outside T , then $A_{S,T}z_T = b_S$ and for each $j \notin S$, $A_{\{j\},T}z_T = b_j$. But $A_{\{j\},T}$ is expressible uniquely as a linear combination of the rows of $A_{S,T}$. It follows that b_j is the same combination of b_S , and so $\det([A_{S',T} : b_{S'}]) = 0$.

But we can assume without loss of generality that z is 0 outside T . This is because the columns indexed by T span all the columns of A . So if for $j \notin T$, $z_j \neq 0$, then we can set it to 0 and adjust the values of z in the T part to account for the contribution of the j th column. \square

So now, to check if the system $Az = b$ is feasible, we check condition (2) of Proposition 33. We claim that this can be checked in $\text{NC}^0_{k+1} \cdot \text{C}_{=} \text{NC}^1$. To see this, note that A is a matrix of $O(1)$ size. So determinants of all its submatrices can be computed by $O(1)$ -size arithmetic formulas taking the GapNC^1 function values as inputs. In terms of the input word, these are themselves GapNC^1 functions. Thus, in particular, using the closure properties of $\text{C}_{=} \text{NC}^1$ (Proposition 5), computing the rank of A or a sub-matrix of A is in $\text{C}_{=} \text{NC}^1$.

Now, the condition (2) is a disjunction over $\sum_{r=0}^k \binom{k}{r}^2 < 2^{2k}$ conditions, one for each $S, T \subseteq [k]$. Each condition is of the form

$$[|S| = |T| = \text{rank}(A) = \text{rank}(A_{S,T})] \Rightarrow \bigwedge_{j \notin S; S' = S \cup \{j\}} [\det([A_{S',T} : b_{S'}]) = 0]$$

The whole condition can thus be expressed as an NC^0 circuit of depth at most $2k + O(\log k) \leq 3k$. \square

From Lemmas 29 and 32, we have shown the following:

Theorem 34. $\text{NC}^0 \cdot \text{C}_{=} \text{NC}^1 = \text{FSLE}^{bdd}[\text{GapNC}^1]$.

5.2. The AC^0 Hierarchy above $\text{C}_{=} \text{NC}^1$

We now show the collapse of the constant-depth hierarchy over $\text{C}_{=} \text{NC}^1$, that is, we prove $\text{AC}^0(\text{C}_{=} \text{NC}^1) = \text{AC}^0_3 \cdot [\text{C}_{=} \text{NC}^1]$. First we set up some notation.

Let $\text{AC}^0_k(\mathcal{C})$ denote the class of languages accepted by AC^0 oracle circuits, where the oracle gates are for a language in \mathcal{C} , and where on any root-to-leaf path, the number of oracle gates encountered is at most k . (This is in analogy with AC^0_k denoting depth- k AC^0 circuits.) Then, $\text{AC}^0_k(\mathcal{C})$ is exactly $\text{AC}^0 \cdot [\mathcal{C}] \cdot \text{AC}^0 \dots (k \text{ times}) \dots [\mathcal{C}] \cdot \text{AC}^0$. In particular, when $\mathcal{C} = \text{C}_{=} \text{NC}^1$, using notation from Proposition 11 we can see that $\text{AC}^0_k(\text{C}_{=} \text{NC}^1)$ equals $\text{a-NC}^1_{\leq k}$ circuits where the nesting depth of the test gates is at most k .

Proposition 35.

$$\begin{aligned} [\text{C}_{=} \text{NC}^1] \cdot \text{AC}^0 &= \text{C}_{=} \text{NC}^1 \\ [\text{coC}_{=} \text{NC}^1] \cdot \text{AC}^0 &= \text{coC}_{=} \text{NC}^1 \end{aligned}$$

Proof. The inclusion from right to left is obvious. We prove the left-to-right inclusion for $\text{C}_{=} \text{NC}^1$; the proof for $\text{coC}_{=} \text{NC}^1$ is identical. So let A be a language in $[\text{C}_{=} \text{NC}^1] \cdot \text{AC}^0$. At each length n , there is an AC^0 oracle circuit C with a single oracle gate g at the top. Let r be the number of input wires to g . On input $x = x_1 x_2 \dots x_n$, let these wires carry the values $z_i(x)$ for $i = 1, \dots, r$. By Proposition 3, there are arithmetic a-AC^0 circuits and hence a-NC^1 circuits Z_i

such that for all x , $Z_i(x) = z_i(x)$. By Proposition 6, there is an $\mathbf{a}\text{-NC}_{=}^1$ circuit D_g with a single test gate at the output such that $D_g(z_1(x), \dots, z_r(x))$ is the bit computed by the oracle gate g . Replacing the inputs $z_i(x)$ in D_g by the circuits Z_i gives an $\mathbf{a}\text{-NC}_{=}^1$ circuit D'_g computing the same function as C . Since D'_g has a single test gate at the top, by Proposition 6, it computes the characteristic function of a language in $\mathbf{C}_{=} \text{NC}^1$. \square

The heart of our collapse result is the following lemma, stating that two adjacent levels of $\mathbf{coC}_{=} \text{NC}^1$ oracle gates can be combined into one.

Lemma 36. $[\mathbf{coC}_{=} \text{NC}^1] \cdot [\mathbf{coC}_{=} \text{NC}^1] \subseteq \text{AC}^0 \cdot [\mathbf{coC}_{=} \text{NC}^1]$. *In particular, the AC^0 circuitry is of depth 3, consisting of an OR of ANDs and some negations at the leaves.*

Proof. The result follows immediately from Lemma 37 below. \square

Lemma 37. *Let $h : \{0, 1\}^t \rightarrow \{0, 1\}$, $f_1, f_2, \dots, f_t : \{0, 1\}^n \rightarrow \{0, 1\}$ be functions in GapNC^1 , where for all w , $f_i(w) \geq 0$. Then for some $T \in t^{O(1)}$, there exist GapNC^1 functions $g_1, g_2, \dots, g_T : \{0, 1\}^n \rightarrow \{0, 1\}$ and an AC^0 circuit H on T inputs such that, for all $w \in \{0, 1\}^n$,*

$$h(b_1, b_2, \dots, b_t) \neq 0 \iff H(d_1, d_2, \dots, d_T) = 1$$

$$\text{where } b_i := \begin{cases} 1 & \text{if } f_i(w) \neq 0, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad d_j := \begin{cases} 1 & \text{if } g_j(w) \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Let C be an $\mathbf{a}\text{-NC}^1$ circuit computing h . Without loss of generality, assume that C is a formula (all gates have out-degree 1), that is layered with alternating $+$ - and \times -layers, and that the underlying graph is a complete binary tree where every root-to-leaf path is of length exactly $2d$.

Without loss of generality, assume that each f_i is non-negative everywhere. (If this not the case, use the function f_i^2 instead of f_i ; this will not change the zero-test and hence will not change b_i . And f_i^2 is also in GapNC^1 .)

Let \mathcal{F} denote the set of functions $\{f_1, \dots, f_t\}$. For each $i \in [t]$, let \mathcal{F}_i denote the set of functions $\mathcal{F} \setminus \{f_i\}$. Also, for each $w \in \{0, 1\}^n$, let $\mathcal{F}(w)$ and $\mathcal{F}_i(w)$ denote the set (or possibly multiset) of values taken by the functions in \mathcal{F} and \mathcal{F}_i on the input w .

Over any set Y of l variables, the symmetric polynomials S_l^k for $0 \leq k \leq l$ are defined as follows.

$$S_l^k(Y) := \sum_{S \subseteq Y; |S|=k} \prod_{y \in S} y$$

For each $i, k \in [t]$, define the function

$$F_i^k(w) := f_i(w) S_{k-1}^{t-1}(\mathcal{F}_i(w)).$$

Let C_k be the circuit obtained from C by replacing the constant 1 with the value $S_k^t(\mathcal{F})$, and replacing the constant -1 with the value $-S_k^t(\mathcal{F})$. Now define the predicate Π_k as follows:

$$\Pi_k(w) \equiv [S_t^k(\mathcal{F}(w)) \neq 0 \wedge S_t^{k+1}(\mathcal{F}(w)) = 0 \wedge C_k(F_1^k(w), \dots, F_t^k(w)) \neq 0]$$

Each entry of B^{-1} can be written as a rational function with $\det(B) = \prod_{0 \leq i < j \leq l} (j - i)$ in the denominator. So, as long as the variables in Y take integer values, the function $\det(B)S_t^k(Y)$ is integral.

Thus to compute $S_t^k(\mathcal{F}(w))$, set $l = t$ in the argument above. Then $D_t^k(w) = \det(B)$, and $N_t^k(w) = \sum_{j=0}^l \det(B)B_{kj}^{-1} \prod_{i=1}^l (f_i(w) + j)$. The terms $\det(B)$ and $\det(B)B_{kj}^{-1}$ are constants that can be precomputed and can be expressed via a-NC¹ circuits. Since each f_i is in GapNC¹, the computation of N_t^k is also in GapNC¹. \diamond

Claim 40. *For each k , there is a GapNC¹ circuit G_k such that $G_k(w) = 0$ if and only if $C_k(F_1^k(w), \dots, F_t^k(w)) = 0$.*

Proof of Claim. We do local surgery on the circuit C_k to carry at every gate g a rational value as a pair of values corresponding to the numerator N_g and the denominator D_g . At a leaf labelled by a constant c , set $N_g = c$, $D_g = 1$. At a leaf labelled $F_1^k(w)$, set N_g and D_g as the output gates of the corresponding circuits described in the previous claims. The gate $g = g_1 + g_2$ has the obvious implementation: $D_g = D_{g_1} \times D_{g_2}$, and $N_g = N_{g_2} \times D_{g_1} + N_{g_1} \times D_{g_2}$. Similarly at gate $g = g_1 \times g_2$, set $D_g = D_{g_1} \times D_{g_2}$, and $N_g = N_{g_1} \times N_{g_2}$. In the resulting circuit, label the wire carrying N_g for the output gate of C_k as the output gate. \diamond

Now we can complete the proof of the lemma. The required GapNC¹ functions are those computed by the circuits N_t^k and G_k for each k . The AC⁰ circuit H implements the check

$$\bigvee_{k=0}^t [N_t^k(\mathcal{F}(w)) \neq 0 \wedge N_t^{k+1}(\mathcal{F}(w)) = 0 \wedge G_k(w) \neq 0]. \quad \square$$

Using Proposition 35 and Lemma 36, we can now establish our collapse result.

Theorem 41. *The AC⁰ hierarchy over $C=NC^1$ collapses to its first level, requiring a single layer of oracle gates and a depth-3 circuit above it,*

$$AC^0(C=NC^1) = AC^0 \cdot [C=NC^1] = AC^0 \cdot [coC=NC^1] = AC^0_3 \cdot [C=NC^1].$$

Proof. Let A be a language in $AC^0(C=NC^1)$; since the AC⁰ circuitry is allowed to use negation gates, equivalently A is in $AC^0(coC=NC^1)$. Then there is a constant k such that $A \in AC^0_k(coC=NC^1)$. The circuit for A thus has the form

$$AC^0 \cdot [coC=NC^1] \cdot AC^0 \cdot \dots \cdot [coC=NC^1] \cdot AC^0.$$

Using Proposition 35, we can absorb all except the topmost AC⁰ circuitry into the oracle gates, giving a circuit of the form

$$AC^0 \cdot \underbrace{[coC=NC^1] \cdot [coC=NC^1] \cdot \dots \cdot [coC=NC^1]}_{(k \text{ times})}.$$

Using Lemma 36, we replace the bottom two oracle layers by a sub-circuit of the form $AC^0 \cdot coC=NC^1$. Then using Proposition 35 again, we absorb this new AC^0 circuitry into the oracle gate layer above it to get a circuit of the form

$$AC^0 \cdot \underbrace{[coC=NC^1] \cdot [coC=NC^1] \cdot \dots \cdot [coC=NC^1]}_{(k-1 \text{ times})}.$$

Repeating this process another $k-2$ times gives the desired circuit of the form $AC^0 \cdot [coC=NC^1]$. Since $coC=NC^1$ contains AC^0 , it can be written in the form $[coC=NC^1] \cdot [coC=NC^1]$. Now using Lemma 36 again, we can replace the top oracle gate by a depth-3 AC^0 circuit. \square

Acknowledgments

The authors thank the anonymous referees of the MFCS 2010 Conference (where a preliminary version of this paper appeared, [DMR⁺10]), and especially the anonymous referees of this journal, for their comments and suggestions; these significantly helped improve the presentation of the results.

References

- [ABO99] E. Allender, R. Beals, and M. Ogihara. The complexity of matrix rank and feasible systems of linear equations. *Computational Complexity*, 8(2):99–126, 1999.
- [All04] E. Allender. Arithmetic circuits and counting complexity classes. In Jan Krajíček, editor, *Complexity of Computations and Proofs*, Quaderni di Matematica Vol. 13, pages 33–72. Seconda Università di Napoli, 2004. An earlier version appeared in the Complexity Theory Column, SIGACT News 28, 4 (Dec. 1997) pp. 2–15.
- [AO96] E. Allender and M. Ogihara. Relationships among PL, #L, and the determinant. *RAIRO Theoretical Information and Applications*, 30:1–21, 1996. Conference version in *Proc. 9th IEEE Structure in Complexity Theory Conference* (1994), 267–278.
- [Bar89] D. A. Barrington. Bounded-width polynomial size branching programs recognize exactly those languages in NC^1 . *Journal of Computer and System Sciences*, 38:150–164, 1989.
- [BRS95] R. Beigel, N. Reingold, and D. A. Spielman. PP is closed under intersection. *Journal of Computer and System Sciences*, 50(2):191–202, 1995.
- [CMTV98] H. Caussinus, P. McKenzie, D. Thérien, and H. Vollmer. Non-deterministic NC^1 computation. *Journal of Computer and System Sciences*, 57:200–212, 1998. Preliminary version in Proceedings of the 11th IEEE Conference on Computational Complexity, 1996, 12–21.

- [DMR⁺10] Samir Datta, Meena Mahajan, B. V. Raghavendra Rao, Michael Thomas, and Heribert Vollmer. Counting classes and the fine structure between nc^1 and l . In *Proceedings of the 35th international conference on Mathematical foundations of computer science*, MFCS'10, pages 306–317, 2010.
- [FR96] L. Fortnow and N. Reingold. PP is closed under truth-table reductions. *Inf. Comput.*, 124(1):1–6, 1996.
- [KSW87] J. Köbler, U. Schöning, and K. W. Wagner. The difference and truth-table hierarchies for NP. *Theoretical Informatics and Applications*, 21(4):419–435, 1987.
- [Lan93] K.-J. Lange. Unambiguity of circuits. *Theor. Comput. Sci.*, 107(1):77–94, 1993.
- [MR09] M. Mahajan and B. V. Raghavendra Rao. Small-space analogues of Valiant’s classes. In *FCT, LNCS vol. 5699*, pages 250–261, 2009.
- [Ogi98] M. Ogiwara. The PL hierarchy collapses. *SIAM J. Comput.*, 27(5):1430–1437, 1998.
- [Tza08] I. Tzamaret. *Studies in Algebraic and Propositional Proof Complexity*. PhD thesis, Tel Aviv University, 2008.
- [Vol99] H. Vollmer. *Introduction to Circuit Complexity: A Uniform Approach*. Springer-Verlag New York Inc., 1999.
- [vzGS91] J. von zur Gathen and G. Seroussi. Boolean circuits versus arithmetic circuits. *Information and Computation*, 91(1):142–154, 1991.
- [Wil85] C. B. Wilson. Relativized circuit complexity. *J. Comput. Syst. Sci.*, 31(2):169–181, 1985.