

Approximate counting and NP search problems

Leszek Aleksander Kołodziejczyk* and Neil Thapen†

December 22, 2018

Abstract

We study a new class of NP search problems, those which can be proved total in the theory APC_2 of [Jeřábek 2009]. This is an axiomatic theory in bounded arithmetic which can formalize standard combinatorial arguments based on approximate counting. In particular, the Ramsey and weak pigeonhole search problems lie in the class. We give a purely computational characterization of this class and show that, relative to an oracle, it does not contain the problem CPLS, a strengthening of PLS.

As CPLS is provably total in the theory T_2^2 , this shows that APC_2 does not prove every $\forall\Sigma_1^b$ sentence which is provable in bounded arithmetic. This answers the question posed in [Buss, Kołodziejczyk, Thapen 2014] and represents some progress in the programme of separating the levels of the bounded arithmetic hierarchy by low-complexity sentences.

Our main technical tool is an extension of the “fixing lemma” from [Pudlák, Thapen 2017], a form of switching lemma, which we use to show that a random partial oracle from a certain distribution will, with high probability, determine an entire computation of a P^{NP} oracle machine. The paper is intended to be accessible to someone unfamiliar with NP search problems or with bounded arithmetic.

1 Introduction

An *NP search problem* is specified by a polynomial-time relation $R(x, y)$ and a polynomial $p(x)$. Given an input x , a solution to the problem is any y such that $R(x, y)$ holds and $|y| < p(|x|)$ (where $|x|$ is the length of a string x). We only consider *total* problems, where a solution is guaranteed to exist for all x . The class of all such problems is called TFNP, standing for *total functional* NP [30].

Subclasses of TFNP are sometimes described as consisting of all search problems which can be proved to be total by some particular combinatorial lemma or style of argument [30, 32]. For example, the class PPA “is based on the lemma that every graph has an even number of odd-degree nodes” [3]. Often, the particular lemma or argument can be represented by a specific axiomatic theory.

*Institute of Mathematics, University of Warsaw, lak@mimuw.edu.pl

†Institute of Mathematics, Czech Academy of Sciences, thapen@math.cas.cz

In this paper we study the class, which we call APPROX, of problems proved total by the theory APC_2 of approximate counting developed by Jeřábek in [19, 20]. APC_2 provides machinery to count the size of a set well enough to distinguish between sets of size a and $(1 + \varepsilon)a$, for a given in binary (but not between sets of size a and $a + 1$), and to formalize a certain amount of induction in this language. In this way it can carry out the standard proofs of, for example, the finite Ramsey theorem and the tournament principle [20]. Our main result is that – in the relativized setting – a search problem known as CPLS [27], which is a natural strengthening of PLS, is not in APPROX.

This answers a question about a hierarchy of theories collectively known as bounded arithmetic. For each $i \in \mathbb{N}$, the theory T_2^i is axiomatized by induction for formulas at level i in the polynomial hierarchy. It is a long-standing open problem whether the theories T_2^i can be separated by sentences expressing that various NP search problems are total – known as $\forall\Sigma_1^b$ sentences. In other words: does the class of provably total NP search problems get strictly bigger as i increases? This is open for $i \geq 2$.

APC_2 lies between T_2^1 and T_2^3 . In [11] we pointed out that the NP search problems typically used in arguments separating T_2^1 from T_2^i for $i \geq 2$ could be proved total using approximate counting. This led us to state the following open problem, which is an important special case of the more general one: is there any i such that T_2^i proves the totality of more NP search problems than APC_2 does?

Our result here implies that the totality of CPLS is not provable in APC_2 . Since the same statement is provable from T_2^2 , this makes APC_2 one of the strongest natural theories that has been separated from theories higher up in the bounded arithmetic hierarchy – in fact, from T_2^i for the lowest possible i – in terms of NP search problems. Intuitively speaking, the conclusion is that the power of T_2^2, T_2^3, \dots to prove many NP search problems total is based on more than just a limited ability to count.

Our main technical tool is the “fixing lemma” from [39]. This is a simplified, but more widely applicable, version of the switching lemma. It shows that a random partial assignment can, with high probability, determine the value of a CNF. We strengthen it slightly, to show that a random partial oracle can determine an entire computation of a P^{NP} oracle machine. The proof of our version is almost identical, and many definitions are identical, to what appears in [39]. We will assume the reader has access to that paper.

In the rest of this introductory section, we give an overview of bounded arithmetic and the theory APC_2 , describe the structure of TFNP from this point of view, and outline how we handle relativization and reductions. In Section 2 we define the problem CPLS and our search-problem class APPROX, formally state our main Theorems 10 and 11, obtain some corollaries, and give an outline of the proofs. In Section 3 we prove Theorem 10, that the class APPROX captures the $\forall\Sigma_1^b$ consequences of APC_2 . In Section 4 we prove our version of the fixing lemma. In Section 5 we use this to prove Theorem 11, that CPLS is not in APPROX. In Section 6 we briefly sketch an alternative way to

prove our main result about bounded arithmetic, Corollary 12, by going through propositional proof complexity rather than NP search problems. In Section 7 we mention some open problems.

Acknowledgements. The first author was partially supported first by grant 2013/09/B/ST1/04390 and then by grant 2017/27/B/ST1/01951 of the National Science Centre, Poland. The second author was partially supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement 339691. Part of this research was carried out during the first author’s visit to Prague, funded by the same ERC grant. The Institute of Mathematics of the Czech Academy of Sciences is supported by RVO:67985840.

1.1 Bounded arithmetic

Fix a language L_{PV} containing a symbol for every function or relation computed by a polynomial-time machine. Then an NP search problem naturally corresponds to a true L_{PV} sentence of the form $\forall x \exists y < t(x) R(x, y)$, where R is a polynomial-time relation, t is a polynomial-time function, and x and y range over natural numbers written in binary notation. Let T be any sound theory. The set of such sentences provable in T then defines a class of search problems. For the class to have some reasonable properties, T should not be too weak; and to get classes of the kind usually studied in complexity theory, it should not be too strong.

Natural theories T come from bounded arithmetic, which has close ties to computational complexity. For the purposes of this paper, we will take such theories to be given by a *base theory* fixing some basic properties of the symbols of L_{PV} , together with one or more axiom schemes that allow us to do stronger kinds of reasoning, typically induction. All axioms are universal closures of *bounded formulas*, that is, formulas in which all quantifiers appear in the form $\forall x < t$ or $\exists x < t$.

In more detail, a PV *formula* is a quantifier-free formula of L_{PV} . A Σ_i^b *formula* is one of the form

$$\exists x_1 < t_1(\bar{z}) \forall x_2 < t_2(\bar{z}, x_1) \dots \varphi(\bar{z}, \bar{x})$$

where φ is a PV formula, the bounds t_i are L_{PV} -terms, quantifiers may appear in alternating \exists and \forall blocks, and there are at most i blocks. Such formulas define precisely the Σ_i^p relations, that is, those at level i in the polynomial hierarchy. The Π_i^b formulas are defined dually. The *universal closure* of a formula $\varphi(\bar{z})$ is the sentence $\forall \bar{z} \varphi(\bar{z})$. Given a class of formulas Γ , we write $\forall \Gamma$ for the set of universal closures of formulas from Γ .

We will consider two base theories, both containing only $\forall PV$ sentences. The first and more usual one is the theory PV, which we will not define here but which comes from Cobham’s characterization of the polynomial-time functions as a function algebra [16, 15]. The second, which we denote $\forall PV(\mathbb{N})$, consists

simply of all \forall PV sentences which are true under the standard interpretation in \mathbb{N} . This is simpler to understand than PV and works more naturally for defining NP search problems. Our results translate easily between the two, and a reader unfamiliar with bounded arithmetic will not go very wrong by reading PV as \forall PV(\mathbb{N}) throughout – see Subsection 1.3.

The important family of theories T_2^i , for $i \geq 0$, is defined as

$$T_2^i := \text{PV} + \Sigma_i^b\text{-IND}$$

where $\Sigma_i^b\text{-IND}$ is the usual induction scheme for Σ_i^b formulas with parameters. PV already proves induction for quantifier-free formulas (as that sort of induction can be witnessed by polynomial-time binary search) so in this setting T_2^0 is the same as PV. We write T_2 for the union of this family.

We can now state a fundamental theorem. By the Σ_i^b -definable functions of a theory we mean the provably total functions with Σ_i^b graphs.

Theorem 1 ([9]). *For $i \geq 0$, the Σ_{i+1}^b -definable functions of T_2^i are precisely the $P^{\Sigma_i^p}$ functions, that is, those that are polynomial-time computable with an oracle from level i of the polynomial hierarchy.*

We will also use a related family of theories S_2^i , for $i \geq 1$, which is defined by replacing the $\Sigma_i^b\text{-IND}$ scheme in T_2^i with the apparently weaker scheme $\Sigma_i^b\text{-LIND}$ in which inductions can only run for polynomially many steps (in the binary length of a parameter). We have $T_2^i \subseteq S_2^{i+1} \subseteq T_2^{i+1}$ [9]. Theorem 1 remains true if T_2^i is replaced by S_2^{i+1} and/or the base theory PV is replaced by \forall PV(\mathbb{N}).

In addition to being related to computational complexity by Theorem 1, bounded arithmetic is a natural environment in which to ask questions about the provability or consistency of theorems or conjectures from complexity theory. For recent examples see [35, 31].

We now make the above definitions slightly more complicated. As in complexity theory, we typically cannot expect to show that two theories of bounded arithmetic are distinct without either making some extra assumption or working relative to some oracle. We will use oracles. We redefine L_{PV} to include a unary relation symbol α standing for “an arbitrary oracle”, and function and relation symbols for all polynomial-time machines with oracle access to α . Other formula classes and theories are redefined to use this extended language. In particular, \forall PV(\mathbb{N}) becomes the set of \forall PV sentences which are true in $\langle \mathbb{N}, A \rangle$ for every oracle A interpreting the symbol α . Strictly speaking, we should change the names to $\text{PV}(\alpha)$, $\Sigma_i^b(\alpha)$, $T_2^i(\alpha)$ etc. However, since we never use the unrelativized versions, we simplify notation by keeping the old names. The results mentioned above still hold.

An open problem. By adapting oracle separation results for the polynomial hierarchy, it has been shown that the strength of the (relativized) theories T_2^i increases strictly with i : for each i there is a $\forall \Sigma_{i+1}^b$ sentence provable in T_2^{i+1}

but not in T_2^i [13]. A pressing open question in proof complexity¹ is whether this remains true if we measure the strength of theories only by their $\forall\Sigma_k^b$ consequences for some fixed k , in particular for $k = 1$.

We write $\forall\Sigma_k^b(T)$ for the $\forall\Sigma_k^b$ consequences of a theory T . From [24, 14] we know that

$$\forall\Sigma_1^b(\text{PV}) \neq \forall\Sigma_1^b(T_2^1) \neq \forall\Sigma_1^b(T_2^2)$$

and from [14, 40] we know that for any $i, k \geq 1$,

$$\text{if } \forall\Sigma_k^b(T_2^i) = \forall\Sigma_k^b(T_2^{i+1}) \text{ then } \forall\Sigma_k^b(T_2^i) = \forall\Sigma_k^b(T_2).$$

The following is open for $k \leq 2$:

$$\text{does } \forall\Sigma_k^b(T_2^2) = \forall\Sigma_k^b(T_2) ? \tag{1}$$

The answer is expected to be negative, even for $k = 0$, by analogy with the Π_1 separation between IS_i and IS_{i+1} given by the second incompleteness theorem. The case $k = 1$ seems to be particularly approachable, as classes $\forall\Sigma_1^b(T)$ have a natural computational interpretation in terms of NP search problems.

Approximate counting. Jeřábek [19, 20] developed a bounded arithmetic theory for approximate counting. Following [11] we call this theory² APC_2 and define it as T_2^1 together with the *surjective weak pigeonhole principle* (sWPHP) for P^{NP} functions, which asserts that no such function can be a surjection from n to $2n$, for any $n > 0$. APC_2 can formalize many arguments in finite combinatorics that use approximate counting, such as the standard proofs of the finite Ramsey theorem and the tournament principle, as well as some probabilistic reasoning. It lies between T_2^1 and T_2^3 in strength, as this instance of the weak pigeonhole principle is provable in T_2^3 .

In [11] we asked the analogue of question (1) for APC_2 in place of T_2^2 . More specifically:

$$\text{does } \forall\Sigma_1^b(\text{APC}_2) = \forall\Sigma_1^b(T_2) ? \tag{2}$$

We expected the answer to be “no”, but the opposite did not seem completely implausible. Approximate counting is a powerful tool in finite combinatorics, and typical combinatorially natural examples of hard $\forall\Sigma_1^b$ statements used to separate T_2^1 from T_2 were known to be provable in APC_2 [11]. Moreover it was shown in [12], by formalizing Toda’s theorem, that all of bounded arithmetic collapses to the analogue of APC_2 if we add a parity quantifier to the language.

Both [11] and later [2] showed unprovability results for various natural subtheories of APC_2 , but these fell well short of answering (2). In fact, they were obtained using a $\forall\Sigma_1^b$ sentence that is actually provable in APC_2 .

¹This is essentially equivalent to a question in propositional proof complexity about separating bounded-depth Frege systems by formulas of fixed depth, and in particular finding a family of small-width CNF’s which have short refutations in bounded-depth Frege but require long refutations in $\text{Res}(\log)$.

²Our definition is slightly different from Jeřábek’s in [20], which uses a variant of the surjective weak pigeonhole principle with a smaller difference between domain and range. However, the theories prove the same $\forall\Sigma_2^b$ statements, which is all that matters for this paper.

1.2 TFNP

In our language, a total NP search problem is simply a true $\forall\Sigma_1^b$ sentence, that is, one of the form $\forall x \exists y < t(x) R(x, y)$ where $R(x, y)$ is a PV formula and t is a L_{PV} -term. This represents the search-task of finding a witness y , given x . We will often assume that the bound $y < t(x)$ is implicit in $R(x, y)$, and we will usually write $R(x, y)$ or just R as a name for the search problem.

As before, polynomial-time is defined relative to an oracle symbol α , and we will occasionally use notation like $R(x, y; \alpha)$ to emphasize the specific oracle being used. The oracle leads to a slight complication in what we mean when we say a search problem is total: $\forall x \exists y < t(x) R(x, y; \alpha)$ must be true in $\langle \mathbb{N}, A \rangle$ for every oracle A interpreting α . This behaviour is essentially the same as what is called a total type-2 NP search problem in e.g. [3, 10].

We define two important notions of reducibility between search problems Q and R . We will introduce one more in Subsection 2.2.

Definition 2. $Q(x, y)$ is *polynomial-time many-one reducible*, or simply *reducible*, to $R(x', y')$, written $Q \leq R$, if there are polynomial-time functions f and g and a polynomial-time relation P (all of which may query the oracle α) such that

$$R(f(x), y'; P(x, \cdot)) \rightarrow Q(x, g(x, y'); \alpha)$$

holds for all x, y' and α , where $P(x, \cdot)$ represents the oracle $\{z : P(x, z)\}$. Two problems are *equivalent* if they are reducible to one another.

Definition 3. $Q(x, y)$ is *polynomial-time Turing-reducible* to $R(x', y')$ if there is a polynomial-time relation P and a polynomial-time oracle machine M which, on input x , makes a series of (adaptive) queries to $R(x', y'; P(\langle x, x' \rangle, \cdot))$. If all replies are correct, then M outputs some y such that $Q(x, y; \alpha)$.

We are interested in search problem classes corresponding to bounded arithmetic theories. There has been a research programme, motivated partly by the logical separation question discussed above, to characterize these classes.

- PV corresponds to FP, the class of search problems which can be solved in deterministic polynomial time [16, 9].
- T_2^1 corresponds to PLS [22, 13]. A PLS problem is given by polynomial-time *neighbourhood* and *cost* functions N_x and C_x and *domain* predicate F_x , such that $0 \in F_x$ and if $y \in F_x$, then $|y| \leq |x|^k$ for some fixed k . A solution to an instance x is any $y \in F_x$ such that either $N_x(y) \notin F_x$ or $C_x(N_x(y)) \geq C_x(y)$. Such a y exists because costs cannot decrease indefinitely. A complete problem for the class is to find a local minimum for a function on a bounded-degree graph.
- T_2^2 corresponds to CPLS, a generalization of PLS described below [27].
- For $k \geq 1$, T_2^k corresponds to a class GI_k defined by the k -turn game induction principle [40] (see also [5, 6]). Equivalent search problems include

further generalizations of PLS and principles about feasible Nash equilibria [38], and LLI_k , the k -round linear local improvement principle [23].

The theory T_2^i is equivalent to a natural formalization of “every $P^{\Sigma_i^p}$ machine has a computation on every input”, essentially by Theorem 1. The search problems above can thus be thought of as the projections onto TFNP of increasingly strong computation models. This can be taken further: there are two “second-order” bounded arithmetic theories, U_2^1 and V_2^1 , which are equivalent (with respect to their $\forall\Sigma_1^b$ consequences) to similar statements about computations of, respectively, PSPACE and EXPTIME machines [9, 23]. In terms of NP search problems, from [23, 7] we have:

- U_2^1 corresponds to LLI_{\log} , the linear local improvement principle with polynomially many rounds.³
- V_2^1 corresponds to LI, the local improvement principle.

One can think of the classes $FP \subseteq GI_1 \subseteq GI_2 \subseteq \dots \subseteq LLI_{\log} \subseteq LI$ as forming a backbone for TFNP. This could be extended even beyond bounded arithmetic, to say $\forall\Sigma_1^b(\text{PA})$ or $\forall\Sigma_1^b(\text{ZFC})$ [4, 41], or by using potentially stronger systems of reasoning, as in [17]. However, experience suggests⁴ that it is difficult to find any natural “combinatorial” NP search problem that is not already provably total in U_2^1 , and thus reducible to LLI_{\log} . In particular, U_2^1 is strong enough to formalize the counting arguments needed to prove the totality of complete problems for the well-known classes PPA and PPP introduced in [32]. Thus, all problems in those classes are reducible to LLI_{\log} .

On the other hand the bijective pigeonhole principle, called *OntoPIGEON* in the search problem literature, is a complete problem for the class PPAD [8] which is contained in both PPA and PPP. Standard proof-complexity lower bound arguments for the pigeonhole principle [26, 36] show that this problem is not provably total in any theory T_2^k and not reducible to any GI_k .

Finally let us mention the search problem class PWPP [21], based on the injective weak pigeonhole principle, which is contained in GI_2 and PPP but not in PLS [29, 24]. We will discuss two other search problems, RAMSEY and HOP, when we state our main results in Subsection 2.3.

It is open whether the hierarchy $GI_2 \subseteq GI_3 \subseteq \dots$ is strict. This is essentially the same problem as the separation of $\forall\Sigma_1^b(T_2^{i+1})$ from $\forall\Sigma_1^b(T_2^i)$ discussed above.

1.3 True and provable reductions

In the previous section, we did not explicitly say what it means for a search problem class to correspond to a theory T . The obvious meaning, that the class

³Unfortunately LLI_{\log} is rather complicated to describe or use. The authors believe that it is equivalent to the simpler game induction principle GI, which is like the principle GI_k of [40] but with polynomially many rounds.

⁴A related issue in propositional proof complexity is the difficulty of finding candidates for separating the Frege and extended Frege proof systems [1].

is precisely $\forall\Sigma_1^b(T)$, potentially has a problem. Namely, such a class does not have the desirable property of being closed under many-one reductions, unless the reductions work provably in T . There may even be two PV formulas R_1 and R_2 which “semantically” define the same relation on \mathbb{N} , and thus the same search problem by the usual complexity-theoretic definition, but are such that T proves that one is total but not the other.

There are two natural ways around this issue. One is to define our class as the closure of $\forall\Sigma_1^b(T)$ under many-one reductions. The other is to stick to theories T that contain the set $\forall\text{PV}(\mathbb{N})$ of all true $\forall\text{PV}$ sentences, and exploit the fact that the statement that a reduction works is such a sentence.

The next lemma shows that, for theories of the kind we consider, these two approaches have the same result. In this paper we prefer the second one.

Lemma 4 (folklore, see also [18]). *Let $Q(x, y)$ be an NP search problem. Let T be a bounded arithmetic theory containing PV and with axioms closed under substituting polynomial-time relations for oracles. Then (1) and (2) below are equivalent. If T contains S_2^1 , then (3) is also equivalent.*

1. Q is provably total in $T + \forall\text{PV}(\mathbb{N})$.
2. $Q \leq R$ for some TFNP problem R provably total in T .
3. Q is Turing reducible to a TFNP problem R provably total in T .

Proof. Suppose (1) holds. We have that $T + \forall z \varphi(z) \vdash \forall x \exists y Q(x, y)$ for some PV formula φ such that $\forall z \varphi(z) \in \forall\text{PV}(\mathbb{N})$. Hence $T \vdash \forall x [\exists z \neg\varphi(z) \vee \exists y Q(x, y)]$. Since T is a bounded theory, by Parikh’s theorem [33] we may add some term $t(x)$ bounding both existential quantifiers. Therefore $T \vdash \forall x \exists y < t(x) R(x, y)$ where $R(x, y)$ is the formula $\neg\varphi(y) \vee Q(x, y)$. Now R is an NP search problem, provably total in T , and Q is reducible to R in \mathbb{N} using the identity function, since $\varphi(y)$ is true for every y and every oracle. Hence (2) holds.

Now suppose (2) holds. Then (3) is immediate. For (1), from the definition of a reduction, there are PV function symbols f, g and a relation symbol P such that, for every oracle A ,

$$\langle \mathbb{N}, A \rangle \models \forall x \forall y' [R(f(x), y'; P(x, \cdot)) \rightarrow Q(x, g(x, y'); A)].$$

Now define the search problem $R^*(x', y'; \alpha) := R(x', y'; P(x, \cdot))$. Then $\forall\text{PV}(\mathbb{N})$ proves $\forall x' \exists y' R^*(x', y') \rightarrow \forall x \exists y Q(x, y)$, and by the property of closure under substitution for the oracle, T proves that R^* is total. Hence we have (1).

Lastly we show that (3) implies (1) under the stronger assumption. Turing reducibility means that there is a polynomial-time oracle machine M which, on input x , makes oracle queries to R and, if the replies are correct, outputs y such that $Q(x, y)$. Formally, for every A , $\langle \mathbb{N}, A \rangle \models \forall x \forall w \varphi(x, w)$ for a PV relation $\varphi(x, w)$ expressing that: if w is a computation of M on input x , and every oracle query x' in w has a reply y' in w such that $P(x', y')$, then $Q(x, \text{output}(w))$. But T proves that, for all x , such a w exists, since Σ_1^b -LIND is enough to construct w query by query. Hence $T + \forall\text{PV}(\mathbb{N}) \vdash \forall x \exists y Q(x, y)$. \square

2 Main definitions and results

2.1 Coloured polynomial local search

We study a search problem introduced in [27]. We will need several results about it from [39], so we take the definition verbatim from there.

Let a, b, c be parameters. Consider a levelled directed graph whose nodes consist of all pairs (i, x) from $[0, a) \times [0, b)$. We refer to (i, x) as *node x on level i* . If $i < a - 1$, this node has a single neighbour in the graph, node $f_i(x)$ on level $i + 1$. Every node in the graph is coloured with some set of colours from $[0, c)$. The principle CPLS, *coloured polynomial local search*, says that the following three statements cannot all be true:

- (i) Node 0 on level 0 has no colours.
- (ii) For every node x on every level $i < a - 1$, and for every colour y , if the neighbour $f_i(x)$ of x on level $i + 1$ has colour y , then x also has colour y .
- (iii) Every node x on the bottom level $a - 1$ has at least one colour, $u(x)$.

When the parameters a, b, c are universally quantified, CPLS is expressed as a $\forall\Sigma_1^b$ sentence about an oracle α encoding the functions f_i and u and a predicate G , where $G_i(x, y)$ means “node x on level i has colour y ”.

To describe it explicitly as a search problem: the inputs are the parameters a, b, c and a solution is a witness that one of items (i)-(iii) above fails. That is, a colour y such that $G_0(0, y)$; or a node (i, x) and a colour y such that $G_{i+1}(f_i(x), y) \wedge \neg G_i(x, y)$; or a node $(a - 1, x)$ such that $\neg G_{a-1}(x, u(x))$.

To see that the principle is true, or equivalently that the search problem is total, suppose that (i)-(iii) hold simultaneously. Then we can reach a contradiction by arguing inductively on i that for all i , some node on level i has no colours. This argument can be formalized as a proof of CPLS in T_2^2 . Moreover, this has a kind of converse, in that it is shown in [27] that CPLS is complete for the search-problem class $\forall\Sigma_1^b(T_2^2)$ with respect to many-one reductions. CPLS simplifies to a PLS problem if we fix the number of colours c to 1.

2.2 Retraction WPHP and Σ_2^p search problems

The *retraction weak pigeonhole principle* [20] asserts that given two functions $f: n \rightarrow 2n$ and $g: 2n \rightarrow n$, there must be some $v < 2n$ such that $f(g(v)) \neq v$. It is true, because otherwise simultaneously f would be a surjection and g an injection. If f and g are polynomial time, this principle naturally gives rise to a problem in TFNP. We will be in a situation where f and g are P^{NP} , and for this we will define a more complex kind of search problem.

Definition 5. A Σ_2^p search problem is specified by a coNP relation $R(x, y)$ and a polynomial bound q such that $\forall x \exists y < 2^{q(|x|)} R(x, y)$. We will often assume that the bound q is implicit in R and will not write it. The problem represents the search-task of finding such a y , given x .

As this definition makes sense outside the context of bounded arithmetic, we have written it in standard complexity-theory notation. But we could alternatively define a Σ_2^P search problem as a true $\forall\Sigma_2^b$ sentence, in the style of our syntactical definition of TFNP problems.

We will need a precise notion of a computation w of a P^{NP} machine M on input v , as follows.

Definition 6. We define a Π_1^b formula “ w is a computation of M on input v ”. The formula interprets w as a sequence $\bar{q}, \bar{r}, \bar{y}$ of respectively NP queries, YES/NO replies and witnesses to replies. It expresses that

1. For each i , q_i is the i -th query asked by M in a computation on input v , assuming the previous replies were r_1, \dots, r_{i-1} ,
2. For each i , if reply r_i is YES then y_i witnesses this,
3. For all sequences \bar{z} of possible counterexamples, for each i , if reply r_i is NO then z_i is not a counterexample to this.

The machine only accesses the oracle α via the NP queries. We say that w is a *precomputation of M on input v* if it satisfies 1. and 2. above.

Note that being a precomputation of M on a given input is a PV formula, so it makes sense to speak of precomputations also when α is only partially defined (as long as the defined part is large enough to verify 2. above). Note also that it is implicit in clause 1. of the definition that each query asked in a computation of M depends only on the input and the previous YES/NO replies to queries, not on the witnesses to the previous replies.

Definition 7. rWPHP_2 is a class of Σ_2^P search problems. A problem in the class is specified by P^{NP} machines for functions $f_x(u)$ and $g_x(v)$, where we treat one argument x as a parameter. The functions f_x and g_x are constrained to take values less than $2x$ and x respectively. An input to the problem is a size parameter x . A solution is a pair (v, w) such that $v < 2x$, w is a computation of $f_x(g_x(v))$ in the sense of Definition 6, and the output of w is not v .

Definition 8. An NP search problem $Q(x, y)$ is PLS *counterexample reducible* to a Σ_2^P search problem $R(x', y')$ if there is a PLS problem $P(x'', y'')$ and polynomial time functions d and e with the following property: for any x, y', y'' such that $P(\langle x, y' \rangle, y'')$, either

1. $d(x, y'')$ witnesses that $R(e(x), y')$ is false, or
2. $Q(x, d(x, y''))$.

This definition should be understood as follows. We are given x and want to find y such that $Q(x, y)$. We create an input $e(x)$ to R and are given a purported solution y' for which it is claimed that $R(e(x), y')$ – this is a coNP claim which we cannot check directly. We then use $\langle x, y' \rangle$ as input for our PLS problem P , and find a solution y'' . Then either 1. or 2. above holds, that is, either the

coNP claim about R was false and $d(x, y'')$ is a counterexample, or $d(x, y'')$ is a solution to our original problem.

We can now introduce a subclass of TFNP with an unusual definition.

Definition 9. The search problem class APPROX consists of all NP search problems PLS counterexample reducible to an rWPHP₂ problem.

We will show that this class coincides with the class of NP search problems that are provably total using approximate counting.

2.3 Results

Theorem 10. $\forall\Sigma_1^b(\text{APC}_2 + \forall\text{PV}(\mathbb{N})) = \text{APPROX}$.

This is proved in Lemmas 15 and 16 in Section 3, by applying standard witnessing techniques to the definition of the theory APC₂.

Theorem 11. CPLS is not in APPROX.

This is proved in Section 5, using a lemma about random oracles proved in Section 4. We briefly sketch the proof. We first fix an alleged PLS counterexample reduction of CPLS to a problem from rWPHP₂ specified by a pair of P^{NP} functions f and g , then choose a large size parameter n and use it to set suitable values for the parameters a, b, c of CPLS. We define a notion of a “legal” partial oracle, which in particular is one which does not contain any witness to CPLS. We adapt a lemma on random restrictions from [39] to show that with high probability a random partial oracle ρ from a certain distribution will “fix” YES or NO replies to all NP queries made in a P^{NP} computation, in the sense that these replies will never become wrong in any legal extension of ρ (Lemma 25). It follows that most partial oracles ρ from this distribution will fix computations of $(f \circ g)(v)$ in this sense on most inputs v . This is enough for ρ to determine a solution (v, w) to our instance of rWPHP₂ for which it is difficult to find a counterexample (Lemma 26). Finally we again adapt a proof from [39] to show, by an Adversary argument in which the Adversary’s strategy uses only legal extensions of ρ , that our PLS reduction is not able to find a witness to CPLS from $\langle v, w \rangle$.

The main result about bounded arithmetic, answering the question posed in [11], is an immediate consequence of these two theorems:

Corollary 12. *The principle CPLS is not provable in APC₂. Since it is provable in T₂², it follows that, in the relativized setting, APC₂ does not prove all $\forall\Sigma_1^b$ consequences of full bounded arithmetic T₂.*

This naturally also limits the strength of theories that are provable in APC₂, such as the following one based on the ordering principle.

Corollary 13. *Consider the theory consisting of T₂¹ together with axioms stating that for every PV formula $R(x, y)$ and every a , if R is a partial ordering on $[0, a)$ then $[0, a)$ contains an R -minimal element. This theory, in the relativized setting, is strictly weaker than T₂².*

Proof. This theory is provable in T_2^2 by straightforward induction on a . It is also provable in APC_2 by an entirely different proof, as is shown in [11] (by an argument due to Jeřábek). By Corollary 12, the theory does not prove CPLS. \square

Both CPLS and the ordering principle have short proofs in the *resolution* propositional proof system, and the argument above could also be used to show that the ordering principle is not “complete” for resolution, in the sense that there are things with short proofs in resolution which do not follow from it. But it is not clear what the most suitable notion of “follow from” is here.

We get similar corollaries about some previously-studied TFNP problems. The first is RAMSEY, which has an oracle relation R and takes an input x : the task is to find a string encoding a set $y \subseteq [0, x)$ of cardinality $\lfloor \log x/2 \rfloor$ such that $[y]^2$ is homogeneous with respect to R . The second is HOP, or the *Herbrandized ordering principle*, which has an oracle for a relation \preceq and a function h , and takes an input x : the task is to find either a witness to the fact that \preceq restricted to $[0, x)$ is not a total ordering, or an element $y < x$ such that $h(y) \not\preceq y$. The function h is needed to make the ordering principle, which is naturally $\forall\Sigma_2^b$, into an NP search problem. The non-reducibility result below would also hold if we redefined HOP to be about partial orderings.

Corollary 14. *In the relativized setting, CPLS is not polynomial-time Turing reducible to either RAMSEY or HOP.*

Proof. Both RAMSEY [37, 20] and HOP [11] are provably total in APC_2 . If CPLS were polynomial-time Turing reducible to either of these problems, then Lemma 4 would imply that CPLS is provable in $\text{APC}_2 + \forall\text{PV}(\mathbb{N})$, contradicting Theorem 10 and Theorem 11. \square

3 Witnessing and definability

This section contains our main technical work in logic: a proof of Theorem 10 via two lemmas corresponding to the two containments in the statement of the theorem. The proofs assume some familiarity with bounded arithmetic.

Intuitively, APC_2 is a combination of T_2^1 and the weak pigeonhole principle, and what we show is that the NP search problems provably total in APC_2 arise as a combination of PLS (which is known to correspond to T_2^1 [13]) and the weak pigeonhole principle, with an important difference that, while a proof can make many “calls” to WPHP, our reductions only allow one call.

Lemma 15. *Every NP search problem provably total in $\text{APC}_2 + \forall\text{PV}(\mathbb{N})$ is PLS counterexample reducible to an rWPHP_2 problem.*

Proof. Let $Q(x, y)$ be an NP search problem. Assume that

$$\text{APC}_2 + \forall z \varphi(z) \vdash \forall x \exists y Q(x, y),$$

where $\varphi(z)$ is a PV formula such that $\mathbb{N} \models \forall z \varphi(z)$ for all oracles. Thus we have

$$\text{APC}_2 \vdash \exists y Q(x, y) \vee \exists z \neg \varphi(z).$$

Writing out the definition of APC_2 , this means

$$T_2^1 + \forall b, c \exists v < 2b \forall u < b \ e(c, u) \neq v \vdash \exists y Q(x, y) \vee \exists z \neg \varphi(z)$$

where the formula on the left is sWPHP for a universal P^{NP} machine $e(c, u)$ running code c on input u with time bound $|c|$. Replacing T_2^1 with the stronger theory S_2^2 and moving sWPHP to the right hand side gives

$$S_2^2 \vdash [\exists b, c \forall v < 2b \exists u < b \ e(c, u) = v] \vee \exists y Q(x, y) \vee \exists z \neg \varphi(z).$$

By Parikh's theorem we can bound b, c , and z by some term $s'(x)$. The presence of S_2^2 lets us use a WPHP amplification argument, as in Corollary 2.2 and Lemma 2.3 of [42] (but relativized by one extra level of quantifiers), to deduce that there is a term $s(x)$ and a one-argument P^{NP} function F such that failure of sWPHP for e at parameters below s' implies that F is a surjection from s onto $2s$. Without loss of generality we may also make the technical assumption that x can be recovered from $s(x)$ by a polynomial time function.

So we obtain

$$S_2^2 \vdash [\forall v < 2s \exists u < s \ F(u) = v] \vee \exists y Q(x, y) \vee \exists z \neg \varphi(z). \quad (3)$$

To match the definition of rWPHP_2 , we define a P^{NP} function of two arguments a, u by $f_a(u) = \min(F(u), 2a - 1)$. Then (3) is equivalent to

$$S_2^2 \vdash [\forall v < 2s \exists u < s \ f_s(u) = v] \vee \exists y Q(x, y) \vee \exists z \neg \varphi(z). \quad (4)$$

The sentence in (4) is $\forall \Sigma_2^b$, so by Buss' witnessing theorem for S_2^2 ([9]) there is a P^{NP} machine which, provably in T_2^1 , maps the input parameters x, v to a triple $\langle u, y, z \rangle$ witnessing one of the three existential quantifiers. Let g be defined so that $g_{s(x)}(v)$ first computes x from $s(x)$ and then outputs the first component u of this witnessing function applied to $\langle x, v \rangle$, as long as $u < s$; otherwise, g outputs 0. We have

$$T_2^1 \vdash [\forall v < 2s \ f_s(g_s(v)) = v] \vee \exists y Q(x, y) \vee \exists z \neg \varphi(z).$$

Now, $f_s(g_s(v)) = v$ can be written as a Π_2^b formula:

$$\forall w [w \text{ is a computation of } f_s(g_s(v)) \rightarrow \text{output}(w) = v],$$

where w is suitably bounded by a term in x , and “ w is a computation of $f_s(g_s(v))$ ” is a Π_1^b formula as in Definition 6, describing the P^{NP} machine that first computes g and then computes f on the output.

So we have

$$T_2^1 \vdash \forall v < 2s \forall w [w \text{ is not a computation of } f_s(g_s(v)) \vee \text{output}(w) = v] \\ \vee \exists y Q(x, y) \vee \exists z \neg \varphi(z).$$

The formula in square brackets is now Σ_1^b , so by the PLS witnessing theorem for T_2^1 ([13]) there is a PLS problem $P(x'', y'')$ witnessing this whole sentence. That is, if we solve P on input $x'' = \langle x, v, w \rangle$ and find y'' such that $P(x'', y'')$, then one of the following holds:

1. w is not a precomputation of $f_s(g_s(v))$, or has output v
2. y'' is a tuple containing a witness that some NO reply in w is wrong
3. y'' is a tuple containing a witness to $\exists y Q(x, y)$.

We know that y'' cannot contain a witness to the last disjunct $\exists z \neg \varphi(z)$ as by assumption $\mathbb{N} \models \forall z \varphi(z)$.

Using the notation of Definition 8, letting d be the function that outputs the witness of incorrectness in case 2., and the witness to $\exists y Q(x, y)$ in case 3., and setting $e(x) = s(x)$, we see that Q is PLS counterexample reducible to the rWPHP₂ problem given by f and g . \square

Lemma 16. *Conversely, if Q is an NP search problem PLS counterexample reducible to an rWPHP₂ problem, then Q is provably total in $\text{APC}_2 + \forall\text{PV}(\mathbb{N})$.*

Proof. Let Q be an NP search problem that is PLS counterexample reducible to the rWPHP₂ problem given by the functions f and g . Let $P(x'', y'')$, d, e be as in the definition of PLS counterexample reducibility. Consider the PV formula $\xi(x, v, w, y, y'')$ defined by

$$v < 2e(x) \wedge P(\langle x, \langle v, w \rangle \rangle, y'') \wedge y = d(x, y'')$$

$$\wedge y \text{ does not witness that } w \text{ is not a computation of } f_e(g_e(v)) \neq v,$$

with y'' suitably bounded. We view ξ as defining an NP search problem with input x and output $\langle v, w, y, y'' \rangle$. Notice that $\xi(x, v, w, y, y'')$ implies that $Q(x, y)$, so we have $Q \leq \xi$.

We claim that APC_2 proves that ξ is total. To see this, work in APC_2 and consider some input x . By sWPHP(PV₂), there is some $v < 2e(x)$ which is outside of the range of $f_{e(x)}$ on inputs below $e(x)$. By T_2^1 , there is some computation of $f_{e(x)}(g_{e(x)}(v))$, say w , which by the choice of v must produce an output different from v . Again by T_2^1 , there is a solution to P on input $\langle x, \langle v, w \rangle \rangle$, say y'' . Clearly, $y = d(x, y'')$ cannot witness that w is not a computation of $f_{e(x)}(g_{e(x)}(v))$ with output different from v , so $\xi(x, v, w, y, y'')$ holds. This proves the claim.

We have shown that Q is reducible to a problem provably total in APC_2 . It follows from Lemma 4 that Q is provably total in $\text{APC}_2 + \forall\text{PV}(\mathbb{N})$. \square

4 Fixing lemma

This section contains our main technical result in complexity, Lemma 22, which is an extension of the “fixing lemma” from [39]. There, the fixing lemma is a

limited switching lemma which says the following: given suitable parameters a, b, c for CPLS, for a well-chosen probability distribution on partial restrictions to an oracle α encoding $(f_i)_{i < a-1}$, u , $(G_i)_{i < a}$, a random restriction has a relatively high probability of determining the value of a narrow CNF in propositional variables standing for bits of α . Importantly, the restriction does not reveal a witness to CPLS; in particular, the (unsatisfiable) CNF asserting that there is no witness to CPLS has to be determined to be true.

In our application in the proof of Theorem 11, we want to fix answers to the NP queries made in a P^{NP} computation. Each query is (the negation of) a CNF, but now there are many of them, and they are made adaptively depending on earlier replies. So we cannot use the lemma from [39] directly. Instead we adapt the proof to show that given a low-depth decision tree labelled with CNFs, with high probability a random restriction fixes the truth values of all CNFs along some branch. This is the basic content of Lemma 22 below.

Our definitions are essentially as the same as in [39], and so is one proof. We will repeat some definitions almost verbatim, but will only give high-level descriptions of some other definitions and of the unchanged proof details.

We think of the bits of the oracle as propositional variables. So, for example, for each node (i, x) there are $\log b$ variables $(f_i(x))_0, \dots, (f_i(x))_{\log b-1}$ expressing the value of $f_i(x)$. A total oracle is defined by a total assignment to all variables. We will be working with partial oracles, which we will also call *partial assignments* or *restrictions*.

We copy in full the definition of a *random restriction* from [39]. First, a *path* in a partial assignment β is a maximal sequence $(i, x_0), \dots, (i+k, x_k)$ of nodes such that $f_{i+j}(x_j) = x_{j+1}$ in β for each $j \in [0, k)$. A path may consist of only one node. If all functions f_i are partial injections, then every node is on some unique path.

Definition 17. ([39, Definition 5.7]) Fix parameters $0 < p, q < 1$. Let $\mathcal{R}_{p,q}$ be the distribution of random restrictions chosen as follows.

- R1. For each pair $i < a$ and $x < b$, with probability $(1-p)$ include (i, x) in a set Z . For each $i < a$, choose f_i uniformly at random from the partial injections from the domain $\{x < b : (i, x) \in Z\}$ into b .
- R2. Set colours on the path beginning at $(0, 0)$ so that $G_i(x, y) = 0$ for all y for all nodes (i, x) on that path.
- R3. For every other path π , with probability $(1-q)$ colour π randomly with one colour. That is, choose uniformly at random a colour y and, for every node (i, x) on π , set $G_i(x, y) = 1$ and then set $G_i(x, y') = 0$ for all $y' \neq y$.
- R4. Finally consider each node $(a-1, x)$ on the bottom level. It is on some path π . If π was coloured at step R3, then set $u(x) = y$ where y is the unique colour assigned to π (that is, $G_{a-1}(x, y) = 1$). Otherwise leave $u(x)$ undefined.

We will also use $\mathcal{R}_{p,q}$ to denote the support of this distribution.

We take the definitions of *legal restrictions* and *good restrictions* from [39, Definition 5.4 and Lemma 5.8]. *Legal* restrictions are those that meet a minimal standard of “niceness” – on every path either no colour variables are set, or they are all set in one of a few particular ways which do not immediately witness CPLS. We do not limit the size of such restrictions and there is no probability measure on them. Our lower bound in the next section will make use of a game played between a Prover, who is trying to witness CPLS by making oracle queries, and an Adversary who is trying to answer queries in a way that does not witness CPLS. It will turn out that the Adversary can restrict herself to answers that come from legal restrictions. In effect, we do not have to worry about the evaluation of formulas under restrictions which are not legal.

A *good* restriction is a legal one which is of typical size, measured in various ways – in particular no path is very long, and there is a reasonable fraction of variables unset at every level. A *bad* restriction is one which is not good.

It may be useful to keep in mind what the analogous definitions would be if we were dealing with the more familiar pigeonhole principle PHP instead of CPLS. A legal restriction would be any restriction representing a partial injection. With probability parameter p , a random restriction would choose holes independently with probability $1-p$, and then randomly map some pigeons to the chosen holes. A good restriction would be a legal one that leaves at least, say, a fraction $p/2$ of holes unset.

We choose a suitable large n and fix our parameters as $a = b = n$, $c = \lfloor n^{1/7} \rfloor$, $p = n^{-4/7}$ and $q = n^{-2/7}$, where b and c are powers of 2.

Lemma 18. ([39, Lemma 5.8]) *The probability that a random restriction is bad is exponentially small in n .*

Definition 19. Let ρ be a restriction. We say that a CNF B is:

- *fixed to 0* by ρ if ρ falsifies B , that is, if for some conjunct in B each literal in the conjunct is set to 0 by ρ ,
- *fixed to 1* by ρ if it is not fixed to 0 by any legal extension of ρ .

Note that a legal restriction can fix a CNF to at most one truth value. The following proposition is therefore obvious.

Proposition 20. *If ρ fixes a CNF B then every extension of ρ also fixes B to the same value.*

It follows from the proof of the “fixing lemma” [39, Lemma 5.9] that, for k reasonably small compared to n , the probability that a given k -CNF is fixed by a random restriction is relatively high – in fact, the probability that it is *not* fixed is $O(kn^{-1/7})$. We need a slightly more general version that also bounds some conditional probabilities.

Lemma 21 (conditional fixing lemma). *Let A_1, \dots, A_m be a sequence of k -CNFs and e_1, \dots, e_m be a sequence of 0/1 values such that*

$$\Pr[\rho \text{ is bad} \mid \rho \text{ fixes each } A_i \text{ to } e_i] < 1/2.$$

Let B be a k -CNF. Then

$$\Pr[\rho \text{ does not fix } B \mid \rho \text{ is good and } \rho \text{ fixes each } A_i \text{ to } e_i] < 12kn^{-1/7}.$$

Proof. If we remove the CNFs A_i , this is essentially the “fixing lemma” of [39, Lemma 5.9] and our proof is almost identical. Define

$$\begin{aligned} F &= \{\rho \in \mathcal{R}_{p,q} : \rho \text{ fixes each } A_i \text{ to } e_i\} \\ G &= \{\rho \in \mathcal{R}_{p,q} : \rho \text{ is good}\} \\ S &= \{\rho \in F \cap G : \rho \text{ does not fix } B\} \end{aligned}$$

so that our S is the intersection with F of the set S defined in [39]. The assumption gives us that $\Pr[F \cap G] / \Pr[F] \geq 1/2$ and our goal is to show that $\Pr[S] / \Pr[F \cap G]$ is small.

Every $\rho \in S$ does not falsify B but does have some legal extension which falsifies B . Exactly as in [39] we define a function θ on S by $\theta(\rho) = \sigma'$, where σ' is a certain minimal legal extension of ρ . We have, over $\rho \in S$,

1. $\Pr[\theta(\rho)] / \Pr[\rho] \geq \frac{1}{2}n^{1/7}$
2. θ is at most $3k$ -to-one
3. $\theta(\rho) \in F$ (although it may happen that $\theta(\rho) \notin F \cap G$).

Items 1 and 2 are proved as in [39]. Item 3 is immediate from Proposition 20.

Now partition S as S_0, \dots, S_{3k-1} where $S_i = \{\rho \in S : \rho \text{ is the } i\text{th preimage of } \theta(\rho)\}$. Then

$$\begin{aligned} \Pr[S_i] &= \sum_{\rho \in S_i} \Pr[\rho] = \sum_{\rho \in S_i} \Pr[\theta(\rho)] \frac{\Pr[\rho]}{\Pr[\theta(\rho)]} \leq 2n^{-1/7} \sum_{\rho \in S_i} \Pr[\theta(\rho)] \\ &\leq 2n^{-1/7} \Pr[F] \end{aligned}$$

where for the last inequality we use that $\sum_{\rho \in S_i} \Pr[\theta(\rho)] \leq \Pr[F]$, since θ is an injection from S_i to F . This step is the main difference from [39], which uses only that θ is an injection from S_i to $\mathcal{R}_{p,q}$, giving the weaker bound $\sum_{\rho \in S_i} \Pr[\theta(\rho)] \leq \Pr[\mathcal{R}_{p,q}] = 1$.

It follows that $\Pr[S] = \Pr[S_0] + \dots + \Pr[S_{3k-1}] \leq 6kn^{-1/7} \Pr[F]$. Hence $\Pr[S] / \Pr[F \cap G] \leq 12kn^{-1/7}$ as required. \square

Lemma 22. *Consider a complete binary decision tree in which each internal node is labelled with a k -CNF and has outgoing edges for NO and YES answers. A node z and a restriction ρ are compatible if ρ is good and, for every CNF B on the path down from the root to z , ρ fixes B to the value specified by the outgoing edge along the path.*

Let $\varepsilon = \Pr[\rho \text{ is bad}]$. A node z is big if $\Pr[\rho \text{ is compatible with } z] > \varepsilon$. Let S_d be the set of good restrictions ρ which are compatible with some big node at depth d . Then

$$\Pr[S_d] \geq 1 - d \cdot 12kn^{-1/7} - 2^{d+1}\varepsilon.$$

This will be used in the next section, where the decision tree will model a computation of a P^{NP} machine. In particular d and k will be polylogarithmic in n and ε will be exponentially small in n . It follows from the lemma that at least one node on the bottom level, and thus at least one computation of the machine, is compatible with some ρ .

Proof. We use induction on d . For the base case $d = 0$, first observe that every good restriction is compatible with the root. It follows that the root is big, as we may assume that $\varepsilon < 1/2$. Hence S_0 is just the set of good restrictions.

At depth d in the tree, by the definition of compatibility each restriction in S_d is compatible with exactly one big node. Consider any such big node z . It is labelled with a k -CNF B and has a NO child z_0 and a YES child z_1 . Define P_z as

$$\Pr[\rho \text{ is not compatible with either } z_0 \text{ or } z_1 \mid \rho \text{ is compatible with } z]$$

This is equal to the probability that ρ does not fix B , under the condition that ρ is good and correctly fixes all CNFs above z . To apply Lemma 21 we need the probability that ρ is bad, given that ρ correctly fixes all CNFs above z , to be less than $1/2$; but this follows from z being big. So by the lemma, $P_z < 12kn^{-1/7}$.

Hence, summing over big nodes at level d , the probability that ρ is compatible with some (not necessarily big) node at depth $d + 1$ is at least

$$\sum_{z \in \{0,1\}^d, z \text{ big}} (1 - P_z) \Pr[\rho \text{ is compatible with } z] \geq (1 - 12kn^{-1/7}) \Pr[S_d].$$

To obtain S_{d+1} we must finally remove the restrictions which are compatible with non-big nodes at depth $d + 1$. But there are at most 2^{d+1} such nodes, so the probability of being compatible with any of them is at most $2^{d+1}\varepsilon$. A straightforward calculation shows that

$$(1 - 12kn^{-1/7})(1 - d \cdot 12kn^{-1/7} - 2^{d+1}\varepsilon) - 2^{d+1}\varepsilon \geq 1 - (d+1) \cdot 12kn^{-1/7} - 2^{d+2}\varepsilon,$$

which completes the inductive step. \square

5 Non-reducibility

Consider a restriction ρ and a Σ_1^b formula $\exists y < t \theta(a, y)$, where θ is a PV formula and a is some number. We say that this formula *is witnessed* in ρ if there is some $b < t$ such that $\theta(a, b)$ holds in ρ . That is, if you run the computation verifying $\theta(a, b)$ and answer queries to α with values from ρ , these values are all defined and the computation is accepting.

Recall that a precomputation of a P^{NP} machine contains a correct witness for every YES reply, but may be wrong about NO replies.

Definition 23. Let ρ be a restriction. A precomputation w of a P^{NP} machine M is *fixed* by ρ if both of the following hold.

1. For every NP query in w with a YES reply, the witness provided by w is correct in ρ .
2. No NP query in w with a NO reply is witnessed in any legal $\sigma \supseteq \rho$.

We say that ρ *fixes a precomputation of M on input v* if there is some such w . For a function f computed by a P^{NP} machine, we write $\rho \Vdash w: f(x) = y$ if ρ fixes a precomputation w of f on input x that outputs y , and we write $\rho \Vdash f(x) = y$ if ρ fixes some such w .

If w is fixed by ρ then ρ fixes, in the sense of Definition 19, each DNF representing an NP query made in w . Note that w does not have to be a computation of M relative to any complete oracle α extending ρ (in fact, in interesting cases w *cannot* be a computation of M).

Remark. The symbol \Vdash is intentionally chosen to be the same one as in forcing. In fact, one could formulate the concept of fixing in terms of a forcing relation, with the restrictions as forcing conditions. However, attempting to preserve all the trappings of forcing in the context of finite combinatorics leads to some annoying issues, so in this paper we do not explore this possibility further.

Lemma 24. *For a P^{NP} function f , a restriction ρ and an input x , there is at most one y such that $\rho \Vdash f(x) = y$.*

Proof. The progress of a P^{NP} precomputation depends only on the YES/NO replies to NP queries, not on the witnesses chosen. In all precomputations of $f(x)$ fixed by ρ these replies are necessarily the same. \square

Below a “suitable” n is one for which $n^{1/7}$ is a power of two.

Lemma 25. *Let M be a P^{NP} machine, running on inputs x with $|x|$ polylogarithmic in n . For all suitable large n , for every such input x ,*

$$\Pr_{\rho \sim \mathcal{R}_{p,q}}[\rho \text{ fixes a precomputation of } M \text{ on } x] \geq 1 - n^{-1/6}.$$

Proof. We can model a run of M on v as a decision tree \mathcal{T}_M . The height d of \mathcal{T}_M is bounded by the running time of M . At each node the tree makes an NP query; by negating the reply, we can view this as a query to a k -CNF, where k is some obvious syntactic upper bound on the time needed to verify a witness to the query. Since M is a P^{NP} machine, k can be chosen polynomial in the running time of M . So we can apply Lemma 22 with $k = d = |n|^c$ for some $c \in \mathbb{N}$. This gives the lower bound

$$1 - |n|^{2c} n^{-1/7} - 2^{|n|^c + 1} \varepsilon \tag{5}$$

on the probability that ρ is compatible with one of the leaves of \mathcal{T}_M . By Lemma 18 the probability ε that ρ is bad is exponentially small in n , so the bound in (5) is at least $1 - n^{-1/6}$ for n sufficiently large.

Finally, suppose ρ is compatible with a leaf of \mathcal{T}_M . We form a precomputation w by answering queries with the replies given on the path from the root

to the leaf. For each YES reply, it follows from the definition of fixing a DNF to 1 (that is, fixing a CNF to 0) that ρ provides enough information to verify at least one witness to the reply; we make some such witness part of w . \square

Lemma 26. *Let a search problem in rWPHP_2 be given by P^{NP} functions $f_x(u)$ and $g_x(v)$. Let s be quasipolynomial in n . Then for all suitable large n ,*

$$\Pr_{\rho \sim \mathcal{R}_{p,q}}[\text{there exist } v \neq v' \text{ such that } \rho \Vdash f_s(g_s(v)) = v'] \geq 1 - 3n^{-1/6}.$$

Proof. Choose n sufficiently large. Let M be the P^{NP} machine which takes input n, v and computes $f_s(g_s(v))$ by first computing g and then f . As s is quasipolynomial in n , we may assume that M satisfies the assumption of Lemma 25 on input size. We will write just f and g below, suppressing the parameter s .

Consider the machine M running on inputs $v < 2s$. By Lemma 25, for any fixed v , a random ρ fixes a precomputation of M on v with probability at least $1 - n^{-1/6}$. It follows that with probability at least $1 - 3n^{-1/6}$ a random ρ simultaneously fixes precomputations for at least $2/3$ of all inputs v (as otherwise the fraction of pairs (ρ, v) in which ρ does not fix a precomputation on v would be more than $n^{-1/6}$). Fix such a ρ . In particular, there are at least $s + 1$ many distinct inputs v_0, \dots, v_s for which ρ fixes precomputations w_0, \dots, w_s .

The machine M first computes $u = g(v)$, which is necessarily less than s , and then computes $f(u)$. Hence, by the pigeonhole principle, there is some u for which there exist distinct i, j such that $\rho \Vdash w_i : g(v_i) = u$ and $\rho \Vdash w_j : g(v_j) = u$. (Here and below we are abusing our notation slightly, as w_i and w_j are really precomputations of g followed by f .)

But, by Lemma 24, there must be a single v' such that $\rho \Vdash w_i : f(u) = v'$ and $\rho \Vdash w_j : f(u) = v'$. At least one of v_i and v_j is distinct from v' ; without loss of generality suppose v_i is. Let $v = v_i$ and $w = w_i$. Thus we have $\rho \Vdash w : f(g(v)) = v'$ and $v' \neq v$, as required. \square

Now consider the following Prover-Adversary game, given by an NP search problem $Q(x, y)$ and a Σ_2^P search problem $R(x', y')$. At the start of the game, the Prover queries R for some input x' , with $|x'|$ polynomial in $|x|$, and the Adversary gives a reply y' . Then the Prover repeatedly queries bits of the oracle α , and the Adversary replies. The Prover is limited in the number of bits of α he can remember at once, and can also forget bits to save memory. The Prover wins when the partial oracle in his memory either witnesses $Q(x, y)$ for some y , or witnesses that $R(x', y')$ is false. This game models PLS counterexample reducibility, in the following sense.

Lemma 27. *Suppose an NP search problem $Q(x, y)$ is PLS counterexample reducible to a Σ_2^P search problem $R(x', y')$. Then for all inputs x the Prover can win the game using only polynomially many (in $|x|$) bits of memory.*

Proof. This is an immediate consequence of the definitions of PLS counterexample reducibility (Definition 8) and PLS. Using the notation of Definition 8, the Prover first asks for y' such that $R(e(x), y')$. He then sets $x'' = \langle x, y' \rangle$ and

simulates the (exponential time, but polynomial memory) task of solving the PLS problem $P(x'', y'')$ by starting with $y'' = 0$ and then repeatedly setting y'' to $N_{x''}(y'')$, finding domain elements of smaller and smaller cost, until either the costs stop decreasing or y'' leaves the domain $F_{x''}$. He never needs to remember more bits of the oracle than are necessary to fix simultaneously the cost and membership of the domain of one solution, the computation of its neighbour, the cost of the neighbour, and possibly computations of d , e , and the witnessing for Q and R . \square

We can now prove Theorem 11, that CPLS is not in the class APPROX.

Proof of Theorem 11. Assume that CPLS is in APPROX and that it is PLS counterexample reducible to the instance of rWPHP_2 given by functions f and g . This means that, by Lemma 27, the Prover can win the Prover-Adversary game in which Q is CPLS and R is rWPHP_2 , using only polynomially many bits of memory. We obtain a contradiction by describing a strategy for the Adversary that defeats any Prover with small memory.

The Prover first makes his query x' to rWPHP_2 . The Adversary then picks a restriction ρ from $\mathcal{R}_{p,q}$ for which there exist a precomputation w and numbers $v, v' < 2x'$ with $v \neq v'$ such that $\rho \Vdash w : f_{x'}(g_{x'}(v)) = v'$. By Lemma 26 such a ρ exists, and by Lemma 18 we may further assume that it is good. The Adversary replies with $\langle v, w \rangle$.

Then, using the limited size of ρ and of the Prover's memory, the Adversary is able to have in hand throughout the game a legal $\sigma \supseteq \rho$ which contains all bits in the Prover's current memory. Such a σ can never witness CPLS, because it is legal. However, a legal σ also cannot witness that $\langle v, w \rangle$ is not a solution to rWPHP_2 , because the only way to do this would be to witness that one of the NO replies in w is wrong, which is impossible by the choice of ρ . The details of the strategy are as in the proof of [39, Theorem 5.10]. \square

6 Reformulation in propositional logic

In this section we sketch another way of presenting our main result about bounded arithmetic, that CPLS, considered as a $\forall\Sigma_1^b$ principle, is not provable in APC_2 . We will use propositional proof complexity and in particular the well-known Paris-Wilkie translation of relativized bounded arithmetic into propositional logic [34].

Suppose φ is bounded formula of L_{PV} , and that we have specified values \bar{n} for all free variables in φ . We can write a propositional formula $\langle \varphi \rangle$ with the same semantics as φ , if we interpret propositional variables x_n as bits $\alpha(n)$ of the oracle. Below we will use *narrow* to mean “of width polylogarithmic in \bar{n} ”.

If φ does not mention the oracle α , then its translation $\langle \varphi \rangle$ is the propositional constant \top or \perp , depending on whether φ is true or false in \mathbb{N} . If φ is $\alpha(n)$, then $\langle \varphi \rangle$ is the propositional variable x_n . If φ is a PV formula, then $\langle \varphi \rangle$ is a narrow CNF — we can take it to be the conjunction of clauses expressing “some oracle reply in w is false” over all possible rejecting computations w

of the polynomial-time machine deciding φ . If φ is a Π_1^b formula $\forall x < n \theta(x)$, then again $\langle \varphi \rangle$ is a narrow CNF, namely the conjunction, over $m < n$, of the translations $\langle \theta(x) \rangle$ with $x \mapsto m$.

The translation theorem we will use follows from the translation of T_2^1 into treelike Res(log) refutations from [25] and the connection between treelike Res(log) and narrow resolution [28]. It can also be shown via PLS witnessing, as described in [11].

Theorem 28. *Let $\varphi(\bar{n})$ be a Π_1^b formula and suppose $T_2^1 \vdash \forall \bar{n} \neg \varphi(\bar{n})$. Then the translations $\langle \varphi \rangle$ have narrow resolution refutations.*

Now suppose for a contradiction that $\text{APC}_2 \vdash \text{CPLS}$. Consider the instances of CPLS described in Section 4, with parameters $a = b = n$ and $c = \lfloor n^{1/7} \rfloor$ and the structure of the problem given entirely by the oracle. Let $Q(n, y)$ assert that y is a solution to such an instance. We may bound y by some term $t(n)$, such that $\text{APC}_2 \vdash \forall n \exists y < t Q(n, y)$. By the proof of Lemma 15, there exist P^{NP} machines f, g defining an instance of rWPHP_2 , and a term $s(n)$, such that

$$T_2^1 \vdash \forall n \forall v < 2s \forall w [w \text{ is not a computation of } f_s(g_s(v)) \vee \text{output}(w) = v \vee \exists y < t Q(n, y)].$$

Let M be the P^{NP} machine which takes input n, v and computes $f_s(g_s(v))$ by first computing g and then f . We think of v as the “real input” to M and of n as a parameter, and write $\text{Comp}_M(v, w)$ for the Π_1^b formula from Definition 6 expressing that w is a computation of M on input v . Noting that the expression on the right above is $\forall \Sigma_1^b$, we can apply Theorem 28 to conclude that the family of narrow CNFs

$$\Phi_{n,v,w} := \langle v < 2s \rangle \wedge \langle \text{Comp}_M(v, w) \rangle \wedge \langle \text{output}(w) \neq v \rangle \wedge \bigwedge_{y < t} \langle \neg Q(n, y) \rangle$$

has narrow resolution refutations, that is, of width polylogarithmic in n, v, w .

Fix a suitable large n . By definition, no legal restriction σ can falsify any clause in the last conjunct $\bigwedge_{y < t} \langle \neg Q(n, y) \rangle$, as otherwise for some y there is an accepting computation of $Q(n, y)$ over σ , so σ witnesses CPLS.

By Lemma 26, with high probability for a random ρ from $\mathcal{R}_{p,q}$ there exist $v < 2s$ and a precomputation w of M on s with $\text{output}(w) \neq v$ such that w is fixed by ρ , meaning that all witnesses in w to YES answers are correct in ρ and no query with a NO answer has a witness in any legal extension of ρ . It follows that no clause in the first three conjuncts is false in any legal extension of ρ . By Lemma 18 we can pick a good ρ for which such v, w exist.

By the Prover-Adversary construction in the proof of [39, Theorem 5.10], we can exploit the limited width of the refutation of $\Phi_{n,v,w}$ to find a legal extension of ρ which falsifies one of the conjuncts of $\Phi_{n,v,w}$. This is a contradiction.

7 Open problems

The *random resolution* propositional proof system was introduced in [11]. Very roughly speaking, a refutation of a CNF F in this system is a refutation of $F \wedge A$, where A is any CNF which is true with high probability.

Suppose a sentence $\forall n \varphi(n)$, with φ a Σ_1^b formula, is provable in the subtheory of APC_2 consisting of T_2^1 together with the surjective WPHP only for polynomial time functions. It was shown in [11] that this implies that the translations $\langle \neg\varphi(n) \rangle$ have narrow refutations in random resolution.

Open Problem 1. Is there a natural propositional proof system which captures, in a similar way, the $\forall\Sigma_1^b$ consequences of full APC_2 ?

Ideally, one would want to show not only that APC_2 proofs translate into the system, but also something in the opposite direction, for example, that if $\langle \neg\varphi(n) \rangle$ has small, suitably uniform refutations in the system, then $\forall n \varphi(n)$ is provable in APC_2 . Some system with these properties could be constructed using the Paris-Wilkie translation and our arguments in Section 6, but it would be rather unnatural and awkward.

It is consistent with what we know that narrow random resolution, or possibly random resolution with no width restriction, already provides a positive answer to Open Problem 1. So, we can ask:

Open Problem 2. Is there a $\forall\Sigma_1^b$ sentence which is provable in APC_2 but whose propositional translations do not have narrow random resolution refutations?

A candidate is the Herbrandized ordering principle HOP, which is provable in APC_2 [11] but not in the subtheory mentioned above [2].

What makes this problem interesting is that, so far, our only tool for proving lower bounds on random resolution is the fixing lemma of [39]. For a typical random restriction, it is a small step from proving this to proving our conditional fixing lemma from Section 4, which implies unprovability in APC_2 . But showing a separation seems to require finding a principle and a random restriction for which one lemma holds, but not the other. The restrictions used to show unprovability of HOP in [2] may be useful here.

Finally we mention a rather obvious question: is every problem in APPROX reducible to CPLS? This is subsumed in the old open problem, discussed in the introduction, of separating the classes GI_k or the theories T_2^k : it is possible that every search problem reducible to any GI_k is already reducible to CPLS.

References

- [1] James Aisenberg, Maria Luisa Bonet, Sam Buss, Adrian Crăciun and Gabriel Istrate *Short Proofs of the Kneser-Lovász Coloring Principle*. Information and Computation, to appear. A conference version appeared in ICALP 2015.

- [2] Albert Atserias and Neil Thapen. *The Ordering Principle in a Fragment of Approximate Counting*. ACM Transactions on Computational Logic 15:4, article 29, 2014.
- [3] Paul Beame, Stephen Cook, Jeff Edmonds, Russell Impagliazzo and Toniann Pitassi. *The relative complexity of NP search problems*. Journal of Computer and System Sciences 57:1, pp. 3-19, 1998.
- [4] Arnold Beckmann. *A Characterisation of Definable NP Search Problems in Peano Arithmetic*. Proceedings of WoLLIC 2009, pp. 1-12, 2009.
- [5] Arnold Beckmann and Samuel Buss. *Polynomial Local Search in the Polynomial Hierarchy and Witnessing in Fragments of Bounded Arithmetic*. Journal of Mathematical Logic 9:1, pp. 103-138, 2009.
- [6] Arnold Beckmann and Samuel Buss. *Characterizing Definable Search Problems in Bounded Arithmetic via Proof Notations*. In Ways of Proof Theory, ONTOS Series in Mathematical Logic, pp. 65-134, 2010.
- [7] Arnold Beckmann and Samuel Buss. *Improved Witnessing and Local Improvement Principle for Second-Order Bounded Arithmetic*. ACM Transactions on Computational Logic 15:1, article 2, 2014.
- [8] Joshua Buresh-Oppenheim and Tsuyoshi Morioka. *Relativized NP search problems and propositional proof systems*. Proceedings of the 19th IEEE Conference on Computational Complexity (CCC), pp. 54-67, 2004.
- [9] Samuel Buss. *Bounded Arithmetic*. Bibliopolis, 1986.
- [10] Samuel Buss and Alan Johnson. *Propositional Proofs and Reductions between NP Search Problems*. Annals of Pure and Applied Logic 163:9, pp. 1163-1182, 2012.
- [11] Samuel Buss, Leszek Aleksander Kołodziejczyk and Neil Thapen. *Fragments of approximate counting*. Journal of Symbolic Logic 79:2, pp. 496-525, 2014.
- [12] Samuel Buss, Leszek Aleksander Kołodziejczyk and Konrad Zdanowski. *Collapsing modular counting in bounded arithmetic and constant depth propositional proofs*. Transactions of the AMS 367, pp. 7517-7563, 2015.
- [13] Samuel Buss and Jan Krajíček. *An application of Boolean complexity to separation problems in bounded arithmetic*. Proceedings of the London Mathematical Society 3:1, pp. 1-21, 1994.
- [14] Mario Chiari and Jan Krajíček. *Witnessing functions in bounded arithmetic and search problems*. Journal of Symbolic Logic 63:3, pp. 1095-1115, 1998.
- [15] Alan Cobham. *The intrinsic computational difficulty of functions*. In Logic, Methodology and Philosophy of Science, Proceedings of the Second International Congress, Y. Bar-Hillel, ed., pp. 24-30, 1965.

- [16] Stephen Cook. *Feasibly Constructive Proofs and the Propositional Calculus*. Proceedings of the Seventh Annual ACM Symposium on Theory of Computing, pp. 83-97, 1975.
- [17] Paul Goldberg and Christos Papadimitriou. *Towards a Unified Complexity Theory of Total Functions*. Journal of Computer and System Sciences 94, pp. 167-192, 2018.
- [18] Jiří Hanika. *Herbrandizing search problems in Bounded Arithmetic*. Mathematical Logic Quarterly 50:6, pp. 577-586, 2004.
- [19] Emil Jeřábek. *Approximate counting in bounded arithmetic*. Journal of Symbolic Logic 72:3, pp. 959-993, 2007.
- [20] Emil Jeřábek. *Approximate counting by hashing in bounded arithmetic*. Journal of Symbolic Logic 74:3, pp. 829-860, 2009.
- [21] Emil Jeřábek. *Integer factoring and modular square roots*. Journal of Computer and System Sciences 82:2, pp. 380-394, 2016.
- [22] David Johnson, Christos Papadimitriou and Mihalis Yannakakis. *How easy is local search?* Journal of Computer and System Sciences 37:1, pp. 79-100, 1988.
- [23] Leszek Aleksander Kołodziejczyk, Phuong Nguyen and Neil Thapen. *The provably total NP search problems of weak second-order bounded arithmetic*. Annals of Pure and Applied Logic 162:2, pp. 419-446, 2011.
- [24] Jan Krajíček. *No Counter-Example Interpretation and Interactive Computation*. Logic from Computer Science, MSRI volume 21, pp. 287-293, 1992.
- [25] Jan Krajíček. *On the weak pigeonhole principle*. Fundamenta Mathematicae 170, pp. 123-140, 2001.
- [26] Jan Krajíček, Pavel Pudlák and Alan Woods. *An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle*. Random Structures and Algorithms 7:1, pp. 15-39, 1995.
- [27] Jan Krajíček, Alan Skelley and Neil Thapen. *NP search problems in low fragments of bounded arithmetic*. Journal of Symbolic Logic 72:2, pp. 649-672, 2007.
- [28] Massimo Lauria. *A note about k-DNF resolution*. Information Processing Letters 137, pp. 33-39, 2018.
- [29] Alexis Maciel, Toniann Pitassi and Alan Woods. *A new proof of the weak pigeonhole principle*. Journal of Computer and System Sciences 64:4, pp. 843-872, 2002.

- [30] Nimrod Megiddo and Christos Papadimitriou. *On total functions, existence theorems and computational complexity*. Theoretical Computer Science 81:2, pp. 317-324, 1991.
- [31] Moritz Müller and Ján Pich. *Feasibly constructive proofs of succinct weak circuit lower bounds*. ECCC technical report TR17-144, 2017.
- [32] Christos Papadimitriou. *On the complexity of the parity argument and other inefficient proofs of existence*. Journal of Computer and system Sciences 48:3, pp. 498-532, 1994.
- [33] Rohit Parikh. *Existence and Feasibility in Arithmetic*. Journal of Symbolic Logic 36:3, pp. 494-508, 1971.
- [34] Jeff Paris and Alex Wilkie. *Counting problems in bounded arithmetic*. Methods in mathematical logic, Springer Lecture Notes in Mathematics 1130, pp. 317-340, 1985.
- [35] Ján Pich. *Logical strength of complexity theory and a formalization of the PCP theorem in bounded arithmetic*. Logical Methods in Computer Science 11:2, 2015.
- [36] Toniann Pitassi, Paul Beame and Russell Impagliazzo. *Exponential lower bounds for the pigeonhole principle*. Computational complexity 3:2, pp. 97-140, 1993.
- [37] Pavel Pudlák. *Ramsey's theorem in bounded arithmetic*. Proceedings of Computer Science Logic '90, Lecture Notes in Computer Science vol. 533, Springer, 1991, pp. 308-317.
- [38] Pavel Pudlák and Neil Thapen. *Alternating Minima and Maxima, Nash Equilibria and Bounded Arithmetic*. Annals of Pure and Applied Logic 163, pp. 604-614, 2012.
- [39] Pavel Pudlák and Neil Thapen. *Random resolution refutations*. To appear in Computational Complexity. Available online as ECCC technical report TR16-175, 2018.
- [40] Alan Skelley and Neil Thapen. *The provably total search problems of bounded arithmetic*. Proceedings of the London Mathematical Society 103:1, pp. 106-138, 2011.
- [41] Amirhossein Tabatabai. *Computational Flows in Arithmetic*. arXiv preprint 1711.01735, 2017.
- [42] Neil Thapen. *A model-theoretic characterization of the weak pigeonhole principle*. Annals of Pure and Applied Logic, vol 118, pages 175-195, 2002.