

Hardness magnification near state-of-the-art lower bounds

Igor C. Oliveira* Ján Pich† Rahul Santhanam‡
 Department of Computer Science Kurt Gödel Research Center Department of Computer Science
 University of Oxford University of Vienna University of Oxford

September 10, 2018

Abstract

This work continues the development of hardness magnification. The latter proposes a strategy for showing strong complexity lower bounds by reducing them to a refined analysis of weaker models, where combinatorial techniques might be successful.

We consider gap versions of the meta-computational problems MKtP and MCSP, where one needs to distinguish instances (strings or truth-tables) of complexity $\leq s_1(N)$ from instances of complexity $\geq s_2(N)$, and $N = 2^n$ denotes the input length. In MCSP, complexity is measured by circuit size, while in MKtP one considers Levin's notion of time-bounded Kolmogorov complexity. (In our results, the parameters $s_1(N)$ and $s_2(N)$ are asymptotically quite close, and the problems almost coincide with their standard formulations without a gap.) We establish that for $\text{Gap-MKtP}[s_1, s_2]$ and $\text{Gap-MCSP}[s_1, s_2]$, a *marginal improvement* over the state-of-the-art in unconditional lower bounds in a variety of computational models would imply explicit *super-polynomial* lower bounds.

Theorem. There exists a universal constant $c \geq 1$ for which the following hold. If there exists $\varepsilon > 0$ such that for every small enough $\beta > 0$

- (1) $\text{Gap-MCSP}[2^{\beta n}/cn, 2^{\beta n}] \notin \text{Circuit}[N^{1+\varepsilon}]$, then $\text{NP} \not\subseteq \text{Circuit}[\text{poly}]$.
- (2) $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn] \notin \text{TC}^0[N^{1+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{TC}^0[\text{poly}]$.
- (3) $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn] \notin B_2\text{-Formula}[N^{2+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{Formula}[\text{poly}]$.
- (4) $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn] \notin U_2\text{-Formula}[N^{3+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{Formula}[\text{poly}]$.
- (5) $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn] \notin \text{BP}[N^{2+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{BP}[\text{poly}]$.
- (6) $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn] \notin (\text{AC}^0[6])[N^{1+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{AC}^0[6]$.

These results are complemented by lower bounds for Gap-MCSP and Gap-MKtP against different models. For instance, the lower bound assumed in (1) holds for U_2 -formulas of near-quadratic size, and lower bounds similar to (3)-(5) hold for various regimes of parameters.

Going beyond the standard boolean devices, we identify a natural computational model under which the hardness magnification threshold for Gap-MKtP lies *below* existing lower bounds: U_2 -formulas that can compute parity functions at the leaves (instead of just literals). As a consequence, if one managed to adapt the existing lower bound techniques against such formulas to work with Gap-MKtP , then $\text{EXP} \not\subseteq \text{NC}^1$ would follow via hardness magnification.

*igor.carboni.oliveira@cs.ox.ac.uk

†jan.pich@univie.ac.at

‡rahul.santhanam@cs.ox.ac.uk

Contents

1	Introduction	3
1.1	Context	3
1.2	Results	4
1.3	Discussion	7
2	Preliminaries	8
3	Hardness Magnification via Error-Correcting Codes	10
3.1	Proof of Theorem 1 Case 2 (Magnification for formulas with parities)	10
3.2	Completing the proof of Theorem 1	13
4	Hardness Magnification via Anti-Checkers	14
4.1	Proof of Theorem 4 (Magnification for MCSP)	14
4.2	Proof of Lemma 17 (Anti-Checker Lemma)	15
4.3	The Anti-Checker Hypothesis	20
A	Unconditional Lower Bounds for Gap-MKtP and Gap-MCSP	26
A.1	MKtP – A near-quadratic lower bound against U_2 -formulas	26
A.2	MKtP – Stronger lower bounds for large parameters	29
A.3	MCSP – A similar near-quadratic lower bound against U_2 -formulas	30
B	Hardness Magnification and Proof Complexity	30

1 Introduction

1.1 Context

Strong lower bounds are known in many restricted computational settings (see e.g. [BS90, Juk12]), but progress in understanding the limitations of more expressive devices has been slow and incremental (cf. [Aar17] for a recent survey and references). Table 1 summarizes the current landscape of unconditional lower bounds with respect to general circuits, formulas, branching programs, bounded-depth threshold circuits, and bounded-depth circuits with modular gates. These constitute some of the most widely investigated models extending the weak computational settings for which we already have explicit super-polynomial lower bounds.

Computational Model	Unconditional Lower Bounds	Reference(s)
Boolean Circuits; w.r.t. different forms of explicitness	$P \not\subseteq \text{Circuit}[cN]$, $\text{MA}/1 \not\subseteq \text{Circuit}[N^k]$ $\text{MA}_{\text{EXP}} \not\subseteq \text{Circuit}[\text{poly}]$	[IM02, FGHK16] [BFT98, San09]
Formulas over B_2	$P \not\subseteq B_2\text{-Formula}[N^{2-o(1)}]$	[Neč66]
Formulas over U_2	$P \not\subseteq U_2\text{-Formula}[N^{3-o(1)}]$	[Hås98, Tal14, DM16]
Branching programs	$P \not\subseteq \text{BP}[N^{2-o(1)}]$	[Neč66]
Low-depth threshold circuits	$P \not\subseteq \text{MAJ} \circ \text{THR} \circ \text{THR}[N^{3/2-o(1)}]$	[KW16]
Depth- d threshold circuits	$P \not\subseteq \text{TC}_d^0[N^{1+\exp(-d)}]$ (wires)	[IPS97]
Depth- d circuits with mod gates	$\text{quasi-NP} \not\subseteq \text{ACC}_d^0[\text{poly}]$	[MW18]

Table 1: A summary of several state-of-the-art lower bounds in circuit complexity theory. In our notation, N denotes input length, and $\mathfrak{C}[s]$ refers to \mathfrak{C} -circuits of size $\leq s$. Establishing stronger lower bounds in these different models is open (or non-trivial lower bounds for a function in $\mathbf{E} = \text{DTIME}[2^{O(N)}]$ in the case of ACC_d^0).

A conditional explanation has been proposed to address the difficulty of establishing strong lower bounds in most of these computational settings. The theory of natural proofs [RR97] shows that if a computational device can compute pseudorandom functions, then sufficiently constructive techniques (such as those that have been successful against weaker models) cannot show lower bounds of the form N^k if k is sufficiently large. This connection has been quite influential, and subsequent works (see e.g. [MV15, Bog18]) have further investigated the limitations of lower bound techniques from this perspective.

The Razborov-Rudich framework suggests that proving unconditional lower bounds in stronger computational models might be tightly related to the investigation of meta-computational problems of a particular form: *those referring to the computational complexity of strings or truth-tables*. Indeed, it has been subsequently proved that the existence of a natural property for a class of circuits yields explicit lower bounds against the same class [IKW02, Wil16, OS17, IKV18].

Our results describe a different phenomenon associated to such problems. They show that in several scenarios, if we could establish slightly stronger lower bounds for them, i.e., lower bounds that marginally improve the size bounds described in Table 1, then *super-polynomial* lower bounds for explicit problems would follow. More specifically, this phenomenon concerns computational problems where the complexity of strings are measured according to circuit complexity (often re-

ferred to as MCSP; see [KC00]) or Levin’s time-bounded Kolmogorov complexity [Lev84] (a problem known as MKtP; see [ABK⁺06]). MCSP and MKtP are fundamental meta-computational problems with connections to areas such as learning theory, cryptography, proof complexity, pseudorandomness and circuit complexity (see e.g. [All17] and references therein). We refer to [All01] for more discussion about the importance of these and related complexity measures.

The new results are part of an emerging theory of *hardness magnification* showing that weak lower bounds for some problems imply much stronger lower bounds. Several results of this form have been obtained in different contexts [Sri03, AK10, LW13, MP17, OS18], and we refer to [OS18] for further discussion. Other forms of hardness magnification are known in settings such as communication complexity and arithmetic circuit complexity. A recent example phrased in a way that is closer to our results appears in [CILM18].

As explained in [AK10, OS18], hardness magnification avoids the natural proofs barrier of [RR97]. It is therefore important to understand the role of magnification in connection to super-polynomial lower bounds, and this work takes another step in this direction. Our main contributions can be summarised as follows:

- (i) We employ new techniques to obtain the first magnification theorem for the usual (worst-case) formulation of the MCSP problem.
- (ii) Our results establish hardness magnification for a natural meta-computational problem (MKtP) near the lower bound frontiers in several standard circuit models. In addition, we identify a computational model where hardness magnification for MKtP lies below existing lower bounds.
- (iii) Our hardness magnification theorems hold for problems for which it is possible to establish a variety of non-trivial lower bounds.

1.2 Results

In this section, we formally state our results. We also briefly discuss some of our techniques, which are explained in more detail in the main body of the paper. We defer a more elaborate discussion of some results to Section 1.3.

Notation. We consider formulas over the bases U_2 (fan-in two ANDs and ORs), B_2 (all boolean functions over two input bits), and extended U_2 -formulas where the input leaves are labelled by literals, constants, or parity functions over the input bits of arbitrary arity. The corresponding classes of formulas of size at most s (measured by the number of leaves) will be denoted by $U_2\text{-Formula}[s]$, $B_2\text{-Formula}[s]$, and $U_2\text{-Formula-}\oplus[s]$, respectively. If we do not specify the type of formulas, we are referring to De Morgan formulas (i.e., formulas over U_2). We also consider bounded-depth majority circuits, where each internal gate computes a boolean-valued majority function (MAJ) of the form $\sum_{i \in S} y_i \geq^? t$ (the circuit has access to input literals $x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n$). We measure the size of such circuits by the number of *wires* in the circuit. Depth- d majority circuits of size s will be denoted by $\text{MAJ}_d^0[s]$, where $d \geq 1$ is fixed. We also consider threshold circuits whose internal gates compute a threshold function (THR) of the form $\sum_{i \in S} w_i \cdot y_i \geq^? t$, for $w_i, t \in \mathbb{R}$. We count number of gates in this case, and let $\text{TC}_d^0[s]$ denote the corresponding class of circuits. $\text{Circuit}[s]$ denotes fan-in two boolean circuits of size s and of unbounded depth (gate types do not matter in our results). More generally, for a circuit class \mathfrak{C} , we use $\mathfrak{C}[s]$ to denote \mathfrak{C} -circuits of size $\leq s$, where size is measured by number of gates. Finally, $\text{BP}[s]$ denotes deterministic branching programs of

size at most s . We refer to a standard textbook (see e.g. [Juk12]) for more information about these boolean devices.

Gap-MKtP and lower bounds for EXP. We use N to denote the input length of an instance of Gap-MKtP $_{[s_1, s_2]}$ (see Definition 7 below), where we need to distinguish strings of Kt complexity [Lev84] (a certain time-bounded variant of Kolmogorov complexity) at most $s_1(N)$ from strings of Kt complexity at least $s_2(N)$. It is not hard to see that for constructive bounds $s_1 < s_2$, Gap-MKtP $_{[s_1, s_2]} \in \text{EXP}$.

We establish a hardness magnification theorem for Gap-MKtP. (In Section 2, we review some relations between the complexity classes and boolean devices appearing below.) Let $n = \log N$.

Theorem 1 (Hardness magnification for MKtP). *There is a universal constant $c \geq 1$ for which the following hold. If there exists $\varepsilon > 0$ such that for every small enough $\beta > 0$*

1. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin \text{Circuit}[N^{1+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{Circuit}[\text{poly}]$.
2. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin U_2\text{-Formula-}\oplus[N^{1+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{Formula}[\text{poly}]$.
3. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin \text{AND-THR-THR-XOR}[N^{1+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{TC}_2^0[\text{poly}]$.
4. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin \text{MAJ}_{2^{d'}+d'+1}^0[N^{1+(2/d')+\varepsilon}]$ for some $d' \geq 1$, then $\text{EXP} \not\subseteq \text{MAJ}_d^0[\text{poly}]$.
5. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin B_2\text{-Formula}[N^{2+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{Formula}[\text{poly}]$.
6. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin U_2\text{-Formula}[N^{3+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{Formula}[\text{poly}]$.
7. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin \text{BP}[N^{2+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{BP}[\text{poly}]$.
8. Gap-MKtP $[2^{\beta n}, 2^{\beta n} + cn] \notin (\text{AC}^0[6])[N^{1+\varepsilon}]$, then $\text{EXP} \not\subseteq \text{AC}^0[6]$.

Interestingly, this result shows the existence of a single meta-computational problem that is connected to several frontiers in complexity theory.

The proof of Theorem 1 relies on a refinement of some ideas from [OS18, Section 3.2]. Our contribution here is mostly conceptual. For a sketch of the argument and its underlying techniques, we refer to the discussion in Section 3. We mention that crucial in the proof is the use of error-correcting codes, and that the complexity of computing these objects using different boolean devices gives rise to the distinct magnification thresholds observed in Theorem 1. The formal proof of Theorem 1 appears in Sections 3.1 and 3.2.

In contrast, we observe the following unconditional lower bounds.

Theorem 2 (Strong lower bounds for large parameters). *For every $\varepsilon > 0$ there exists $\delta > 0$ for which the following results hold:*

1. Gap-MKtP $[2^{(1-\delta)n}, 2^{n-1}] \notin U_2\text{-Formula}[N^{3-\varepsilon}]$.
2. Gap-MKtP $[2^{(1-\delta)n}, 2^{n-1}] \notin B_2\text{-Formula}[N^{2-\varepsilon}]$.
3. Gap-MKtP $[2^{(1-\delta)n}, 2^{n-1}] \notin \text{BP}[N^{2-\varepsilon}]$.

The proof of Theorem 2 is simple, assuming certain results. It relies on the existence of pseudorandom generators against small formulas and small branching programs [IMZ12], together with an observation from [All01]. The argument appears in Appendix A.2.

Note the different regime of parameters for $\text{Gap-MkTtP}[s_1, s_2]$ in Theorems 1 and 2. In order to magnify a weak lower bound using Theorem 1, we need that it holds for $s_1 = 2^{o(n)} = N^{o(1)}$. The next result shows that non-trivial unconditional lower bounds can be obtained in this regime.

Theorem 3 (A near-quadratic formula lower bound). *For every constant $0 < \alpha < 2$ there exists $C > 1$ such that $\text{Gap-MkTtP}[Cn^2, 2^{(\alpha/2)n-2}] \notin U_2\text{-Formula}[N^{2-\alpha}]$.*¹

The proof of Theorem 3 adapts ideas from [HS17, Section 4] (see also the exposition in [OS18, Appendix C.1]) employed in the context of MCSP for larger parameters. A sketch of the argument followed by a proof can be found in Appendix A.1.

Gap-MCSP and lower bounds for NP. We use $N = 2^n$ to denote the input length of an instance of $\text{Gap-MCSP}[s_1, s_2]$ (see Definition 9 below), where one needs to distinguish functions of circuit complexity at most s_1 from functions of circuit complexity at least s_2 . It is not hard to see that for constructive bounds $s_1 < s_2$, $\text{Gap-MCSP}[s_1, s_2] \in \text{NP}$.

We establish the following magnification theorem for Gap-MCSP .

Theorem 4 (Hardness magnification for MCSP). *There is a universal constant $c \geq 1$ for which the following holds. If there exists $\varepsilon > 0$ such that for every small enough $\beta > 0$*

1. $\text{Gap-MCSP}[2^{\beta n}/cn, 2^{\beta n}] \notin \text{Circuit}[N^{1+\varepsilon}]$, then $\text{NP} \not\subseteq \text{Circuit}[\text{poly}]$.

MCSP and MkTtP are quite different problems. In our results, an important distinction is that applying a polynomial-time function to an input of MkTtP does not substantially increase its Kt complexity (cf. Proposition 8), but this is not necessarily true in the context of circuit complexity, where the input string represents an entire truth-table. For this reason, the proof of Theorem 4 is completely different from the proof of Theorem 1.

Theorem 4 is our main technical contribution. The argument relies on the notion of *anti-checkers*. Roughly speaking, an anti-checker is a bounded collection \mathcal{S} of inputs associated with a hard function f such that any small circuit C differs from f on some input in \mathcal{S} . More precisely, it was established in [LY94] that any function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ that requires circuits of size s admits a collection \mathcal{S}_f containing $O(s)$ strings that is an anti-checker against circuits of size roughly s/n . Our argument makes crucial use of anti-checkers, and en route to Theorem 4 we give a new, more constructive proof of their existence. (While the proof in [LY94] uses min-max theory, our proof is combinatorial and self-contained.)

We remark that anti-checkers were first employed for hardness magnification in the context of *proof complexity* [MP17]. However, while the existential result from [LY94] was sufficient in that context, this is not the case in *circuit complexity*, and our argument needs to be more sophisticated. For the reader interested in learning more about hardness magnification in proof complexity, how it relates to meta-computational problems such as MCSP, and how the new results compare with our previous work, we refer to Appendix B.

The proof of Theorem 4 is not difficult given a certain lemma about the construction of anti-checkers (see Section 4.1). The crucial Anti-Checker Lemma (see Lemma 17) says that $\text{NP} \subseteq$

¹The constant C has an exponential dependence on $1/\alpha$.

$\text{Circuit}[\text{poly}]$ implies the existence of circuits of almost linear size which given the truth table of a Boolean function f print a corresponding set \mathcal{S}_f . The circuits provided by the Anti-Checker Lemma simulate our new proof of the existence of anti-checkers, but make the involved argument constructive by using approximate counting and the assumption $\text{NP} \subseteq \text{Circuit}[\text{poly}]$. A high-level exposition and the complete proof are described in Section 4.

We note that the argument is different from the proof of Theorem 1 not only on a technical level, but also conceptually. Both proofs rely on a reduction to a certain succinct problem. In the case of Theorem 1, this reduction is implemented by a concrete algorithm that can be explicitly described. However, in Theorem 4 the assumption that the conclusion does not hold (i.e., $\text{NP} \subseteq \text{Circuit}[\text{poly}]$) and its consequences are invoked multiple times to show that the required reduction exists. More generally, this deviates from the compression/kernelization approach present in certain proofs from [AK10] and [OS18].²

We are able to show the following related unconditional lower bound against *formulas*.

Theorem 5. *For every constant $0 < \alpha < 2$ there exists $d > 1$ such that $\text{Gap-MCSP}[n^d, 2^{(\alpha/2 - o(1))n}] \notin U_2\text{-Formula}[N^{2-\alpha}]$.*

Consequently, if one could establish an analogue of Theorem 4 for sub-quadratic formulas, then $\text{NP} \not\subseteq \text{Formula}[\text{poly}]$. We explain why the argument behind the proof of Theorem 4 fails in the case of formulas in Section 4.2.³ The proof of Theorem 5 is similar to the proof of Theorem 3, and we sketch the necessary modifications in Appendix A.3.

Finally, in Section 4.3 we discuss a certain combinatorial hypothesis (“The Anti-Checker Hypothesis”) connected to the techniques behind the proof of Theorem 4. If this hypothesis holds, then $\text{NP} \not\subseteq \text{Formula}[\text{poly}]$. We observe that the hypothesis does hold in the average-case, but we are unsure about its plausibility in the worst-case context that is sufficient for super-polynomial formula lower bounds.

1.3 Discussion

As alluded to above, the thresholds for magnification in our results appear slightly above existing size lower bounds in the case of standard boolean devices, such as branching programs, bounded-depth threshold circuits, formulas over the bases U_2 and B_2 , general circuits, and bounded-depth circuits with modular gates (w.r.t. a lower bound for a problem in $\text{E} = \text{DTIME}[2^{O(N)}]$). We have also observed that some existing lower bounds can be adapted to Gap-MKtP and Gap-MCSP , under different regimes of parameters. Naturally, one is tempted to ask whether hardness magnification can explain the precise borders where progress on lower bounds has halted.

There are natural problems for which showing lower bounds that are *weaker* than the current state-of-the-art size bounds would also imply super-polynomial lower bounds [OS18]. A representative example presented in [OS18] concerns an average-case version of MCSP, where the problem refers to the average-case circuit complexity of the input function, while the magnification result holds with respect to worst-case formula lower bounds. The reason that work does not imply super-polynomial formula lower bounds via magnification is that the corresponding unconditional lower bounds and magnification theorems hold for a different regime of the average-case complexity parameter.

²In particular, we are not aware of results indicating that $\text{MCSP}[s]$ might be kernelizable.

³Note that Theorem 4 implies lower bounds for a problem in NP. Theorem 1 only gives lower bounds in EXP, but its proof extends to several low-complexity settings.

Our results and techniques were motivated in part by the desire to address this gap. On the one hand, it seems to be easier to analyse problems that refer to the worst-case complexity of the input. But on the other hand, shifting from average-case to worst-case complexity often increases the magnification threshold to size bounds that are beyond existing techniques. Another glimpse of the subtle transition between worst-case and average-case complexity and its role in magnification appears in the discussion of the Anti-Checker Hypothesis in Section 4.3.

We identify a computational model that has not received much attention in the literature, and for which the magnification threshold for Gap-MKtP lies below existing lower bounds. This corresponds to Theorem 4 Item 2, i.e., U_2 -formulas augmented with parities in the leaves (our exposition in Section 3 focuses on this model). Note that, by a straightforward simulation, before breaking the cubic barrier for U_2 -formulas or the quadratic barrier for B_2 -formulas, one needs to show super-linear lower bounds against $U_2\text{-Formula-}\oplus$. But a recent result of Tal [Tal16] implies exactly that: the inner product function over N input bits is not in $U_2\text{-Formula-}\oplus[N^{1.99}]$.

This makes this computational model particularly attractive in connection to hardness magnification and lower bounds. Indeed, it seems “obvious” that $\text{Gap-MKtP}[2^{\delta n}, 2^{\delta n} + cn] \notin U_2\text{-Formula-}\oplus[N^{1.01}]$, given that such formulas cannot compute the much simpler inner product function, and that standard formulas require at least near-quadratic size (Theorem 3). Our work shows that if this is the case, then $\text{EXP} \not\subseteq \text{NC}^1$. Needless to say, establishing the required lower bound is our main open problem.

2 Preliminaries

For $\ell \in \mathbb{N}$, we use $[\ell]$ to denote the set $\{1, \dots, \ell\}$. The length of a string w will be denoted by $|w|$. Our logarithms are in base 2, and we use $\exp(x)$ to denote e^x . We use boldface symbols such as \mathbf{i} and $\mathbf{\rho}$ to denote random variables, and $\mathbf{x} \in_R S$ to denote that \mathbf{x} is a uniformly random element from a set S . We often identify n with $\log N$ or N with 2^n , depending on the context.

For concreteness, we employ a random-access model to formalize uniform algorithms. The details of the model are not crucial in our results, and only mildly affect the gap parameters s_1 and s_2 . We fix some standard encoding of algorithms as strings, and use $\langle M \rangle$ to denote the string encoding the algorithm M . Moreover, we assume for simplicity the following property of this encoding: if an algorithm C is obtained via the composition of the computations of algorithms A and B , then $|\langle C \rangle| \leq |\langle A \rangle| + |\langle B \rangle| + O(1)$. (Roughly speaking, composing two codes gives a new valid code.⁴) The running time of M on x is denoted by $t_M(x)$.

We introduce next the notion of Kt complexity. We adopt a formulation that is more convenient for our purposes. In particular, we avoid the use of universal machines in the definition given below.⁵ (Our definition is easily seen to be within at most a logarithmic additive term of the formulation using universal machines. We stress that our proofs can be adapted to work with any reasonable definition.)

⁴While this holds for instance for programs with relative jump instructions (i.e., goto instructions where the new line is encoded relative to the number of the current line), we remark this is not true in general. For instance, composing two Turing Machines might require renaming all states of one machine, which could result in a new encoding of length $(1 + o(1))(|\langle A \rangle| + |\langle B \rangle|)$. Depending on the computational model, the results in Theorem 1 might need parameters $s_2 = (1 + o(1))s_1$.

⁵Universal machines are still needed to upper bound the time complexity of computing Kt complexity. Moreover, the exact Kt complexity of a string depends on the choice of encoding for algorithms/machines.

Definition 6 (Kt Complexity ([Lev84]; see also [All01])). For a string $x \in \{0, 1\}^*$, $\text{Kt}(x)$ denotes the minimum of $|\langle M \rangle| + |a| + \lceil \log t_M(a) \rceil$ over pairs (M, a) such that the machine M outputs x when it computes on the input string a .

Definition 7 (The Gap-MkTP Problem). We consider the promise problem $\text{Gap-MkTP}[s_1, s_2]$, where $s_1, s_2: \mathbb{N} \rightarrow \mathbb{N}$ and $s_1(N) < s_2(N)$ for all $N \in \mathbb{N}$. For each $N \geq 1$, $\text{Gap-MkTP}[s_1, s_2]$ is defined by the following sets of instances:

$$\begin{aligned} \mathcal{YES}_N &\stackrel{\text{def}}{=} \{x \in \{0, 1\}^N \mid \text{Kt}(x) \leq s_1(N)\}, \text{ and} \\ \mathcal{NO}_N &\stackrel{\text{def}}{=} \{x \in \{0, 1\}^N \mid \text{Kt}(x) > s_2(N)\}. \end{aligned}$$

We will need the following simple result.

Proposition 8 (Kt complexity and composition). Let B be an algorithm that runs in time at most $T_B(N)$ over inputs of length N . Then, for every input $w \in \{0, 1\}^N$, as N grows we have

$$\text{Kt}(B(w)) \leq \text{Kt}(w) + \log(T_B(N)) + O(1).$$

Proof. Let A be a machine and a be a string such that the pair (A, a) witnesses the value $\text{Kt}(w)$. Let C be the composition of machines A and B , i.e., $C(y) = B(A(y))$. We claim that the pair (C, a) witnesses the inequality in the conclusion of the proposition. Indeed, since $C(a) = B(A(a)) = B(w)$, we get

$$\begin{aligned} \text{Kt}(B(w)) &\leq |\langle C \rangle| + |a| + \lceil \log t_C(a) \rceil \\ &\leq |\langle A \rangle| + |\langle B \rangle| + O(1) + |a| + \log(t_A(a) + t_B(w)) \\ &\leq |\langle A \rangle| + |a| + \log(t_A(a)) + \log(t_B(w)) + |\langle B \rangle| + O(1) \\ &\leq \text{Kt}(w) + \log(T_B(N)) + O(1), \end{aligned}$$

where we have used that $|\langle B \rangle|$ is constant as N grows. □

We also consider a natural formulation of the gap version of the Minimum Circuit Size Problem (MCSP). The circuit complexity of a boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is denoted by $\text{Size}(f)$. We use the same notation to represent the circuit complexity of the function encoded by a string $x \in \{0, 1\}^{2^n}$.

Definition 9 (The Gap-MCSP Problem). We consider the promise problem $\text{Gap-MCSP}[s_1, s_2]$, where $s_1, s_2: \mathbb{N} \rightarrow \mathbb{N}$ and $s_1(n) < s_2(n)$ for all $n \in \mathbb{N}$. For each $n \geq 1$, $\text{Gap-MCSP}[s_1(n), s_2(n)]$ is defined by the following sets of instances:

$$\begin{aligned} \mathcal{YES}_n &\stackrel{\text{def}}{=} \{x \in \{0, 1\}^{2^n} \mid \text{Size}(x) \leq s_1(n)\}, \text{ and} \\ \mathcal{NO}_n &\stackrel{\text{def}}{=} \{x \in \{0, 1\}^{2^n} \mid \text{Size}(x) > s_2(n)\}. \end{aligned}$$

A brief review of uniform complexity classes and connections to non-uniform devices.

To provide some context for Theorem 1, we remind the reader about the following relations involving boolean devices and complexity classes. Under an appropriate uniform formulation of circuit classes, we have the inclusions:

$$\text{(uniform classes)} \quad \text{AC}^0 \subseteq \text{ACC}^0 \subseteq \text{MAJ}^0 \subseteq \text{TC}^0 \subseteq \text{NC}^1 \subseteq \text{L} \subseteq \text{P}.$$

Some of these classes are related in the non-uniform case as follows: $\text{NC}^1 = \text{U}_2\text{-Formula}[\text{poly}] = \text{B}_2\text{-Formula}[\text{poly}] = (\text{width } 5) \text{BP}[\text{poly}]$, $\text{L}/\text{poly} = \text{BP}[\text{poly}]$, $\text{P}/\text{poly} = \text{Circuit}[\text{poly}]$, and $\text{MAJ}^0[\text{poly}] = \text{TC}^0[\text{poly}]$. These equivalences might require a complexity overhead in size or depth. We refer to [Raz91, Juk12] for these and other related results.

3 Hardness Magnification via Error-Correcting Codes

In this section, we prove Theorem 1. First, we provide a high-level exposition of the argument.

Proof Idea. The result is established in the contrapositive. The idea is to reduce $\text{Gap-MKtP}[s_1, s_2]$ to a problem in EXP over instances of size $\text{poly}(s_1, s_2) \ll N$, and to invoke the assumed complexity collapse to solve Gap-MKtP using very efficient circuits (or other boolean devices). First, we apply an error-correcting code (ECC) to the input string $w \in \{0, 1\}^N$. Since this can be done by a uniform polynomial time computation, we are able to show that $\text{ECC}(w) \in \{0, 1\}^{O(N)}$ is a string of Kt complexity $\ell < s_2$ if w has Kt complexity $\leq s_1$. On the other hand, using an efficient decoder for the ECC, we can show that if w has Kt complexity $\geq s_2$, then any string of Kt complexity $> \ell$ differs from $\text{ECC}(w)$ on a constant fraction of coordinates. Let $z = \text{ECC}(w)$. Given the gap in the input instances of Gap-MKtP , our task now is to distinguish strings z that have Kt complexity at most ℓ from strings that cannot be approximated by strings of Kt complexity at most ℓ , where $s_1 < \ell < s_2$.

We achieve this by using a random projection of the input z to a string y of size roughly $\ell \ll N$. The intuition is that if z has Kt complexity at most ℓ , then every projection of z also agrees with some string (i.e., z) of Kt complexity at most ℓ . However, it is possible to argue that if z cannot be approximated by a string of Kt complexity at most ℓ , then with high probability no string of Kt complexity at most ℓ agrees with the randomly projected coordinates of z . Checking which case holds when we are given the string y can be done by an exponential time algorithm. Under the assumption that EXP admits small circuits, we are able to solve this problem in complexity $\text{poly}(\ell) \ll N$.

The reduction sketched above requires (1) the computation of an appropriate ECC, and (2) is randomized. A careful derandomization and the computation of the ECC in different models of computation provide the size bounds corresponding to the magnification thresholds appearing in the statement of Theorem 1.

We start with a detailed proof of Item (2), which covers the more interesting scenario of formulas with parity leaves. We then discuss how a simple modification of the argument together with known results imply the other cases.

3.1 Proof of Theorem 1 Case 2 (Magnification for formulas with parities)

We will need the following explicit construction.

Theorem 10 (Explicit linear error-correcting codes (cf. [Jus72, SS96])). *There exists a sequence $\{E_N\}_{N \in \mathbb{N}}$ of error-correcting codes $E_N: \{0, 1\}^N \rightarrow \{0, 1\}^{M(N)}$ with the following properties:*

- $E_N(x)$ can be computed by a uniform deterministic algorithm running in time $\text{poly}(N)$.
- $M(N) = b \cdot N$ for a fixed $b \geq 1$.

- There exists a constant $\delta > 0$ such that any codeword $E_N(x) \in \{0, 1\}^{M(N)}$ that is corrupted on at most a δ -fraction of coordinates can be uniquely decoded to x by a uniform deterministic algorithm D running in time $\text{poly}(M(N))$.
- Each output bit is computed by a parity function: for each input length $N \geq 1$ and for each coordinate $i \in [M(N)]$, there exists a set $S_{N,i} \subseteq [N]$ such that for every $x \in \{0, 1\}^N$,

$$E_N(x)_i = \bigoplus_{j \in S_{N,i}} x_j.$$

We proceed with the proof of Theorem 1 Part (2). We establish the contrapositive. Assume that $\text{EXP} \subseteq \text{Formula}[\text{poly}]$, and recall that $N = 2^n$. For any $\varepsilon > 0$, we prove that $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn] \in U_2\text{-Formula-}\oplus[N^{1+\varepsilon}]$ for a sufficiently small $\beta > 0$ and a universal choice of the constant c . The value of c will be specified later in the proof (see Claim 12 below).

Let $E_N: \{0, 1\}^N \rightarrow \{0, 1\}^M$ be the error-correcting code granted by Theorem 10, where $M(N) = bN$. Given an instance $w \in \{0, 1\}^N$ of $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn]$, we first apply E_N to $w \in \{0, 1\}^N$ to get $z = E_N(w) \in \{0, 1\}^M$.

Claim 11. *There exists $c_0 \geq 1$ such that for every large enough N the following holds. If $\text{Kt}(w) \leq 2^{\beta n}$, then $\text{Kt}(z) \leq 2^{\beta n} + c_0 n$.*

Proof. The claim follows immediately from the upper bound on $\text{Kt}(w)$, the definition of $z = E_N(w)$, the running time of E_N , and Proposition 8. \square

Claim 12. *There exist $c > c_1 > c_0 \geq 1$ such that for every large enough N the following holds. If $\text{Kt}(w) > 2^{\beta n} + cn$, then $\text{Kt}(z') > 2^{\beta n} + c_1 n$ for any $z' \in \{0, 1\}^M$ that disagrees with z on at most a δ -fraction of coordinates.*

Proof. Suppose that a string $z' \in \{0, 1\}^M$ disagrees with z on at most a δ -fraction of coordinates, and that $\text{Kt}(z') \leq 2^{\beta n} + c_1 n$ for some $c_1 > c_0$. We upper bound the Kt complexity of w by combining a description of z' with the decoder D provided by Theorem 10. In more detail, assume the pair (F, a) witnesses $\text{Kt}(z')$. Let B be the machine that first applies the machine F to a (producing z'), then D to z' . It follows from Theorem 10 that $B(a) = D(F(a)) = D(z') = w$. Similarly to the proof of Proposition 8, we also get

$$\begin{aligned} \text{Kt}(w) &\leq |\langle B \rangle| + |a| + \lceil \log t_B(a) \rceil \\ &\leq |\langle F \rangle| + |\langle D \rangle| + O(1) + |a| + \log(t_F(a) + t_D(z')) \\ &\leq \text{Kt}(z') + \log(t_D(z')) + O(1) \\ &\leq (2^{\beta n} + c_1 n) + O(n) + O(1) \\ &\leq 2^{\beta n} + cn, \end{aligned}$$

if n is large enough and we choose c sufficiently large. \square

Next we define an auxiliary language $L \in \text{EXP}$, efficiently reduce Gap-MKtP to L , and use the assumption that EXP has polynomial size formulas to obtain almost-linear size formulas (of the appropriate kind) for Gap-MKtP . Roughly speaking, we are able to obtain a formula of non-trivial size for Gap-MKtP because our reduction maps input instances of length N to instances of L of

length $N^{o(1)}$ (the $o(1)$ term is captured by the parameter β using $n = \log N$). As we will see shortly, the reduction is randomized. In order to get the final U_2 -formula- \oplus computing **Gap-MKtP**, the argument is derandomized in a straightforward but careful way. More details follow.

An input string y encoding a tuple $(a, 1^b, (i_1, \alpha_1), \dots, (i_r, \alpha_r))$ belongs to L (where a and b are positive integers, a is encoded in binary, and $\alpha_j \in \{0, 1\}$) if each i_j (for $1 \leq j \leq r$) is a string of length $\lceil \log a \rceil$ and there is a string z of length a such that $\text{Kt}(z) \leq b$ and for each index j we have $z_{i_j} = \alpha_j$.

Claim 13. $L \in \text{EXP}$.

Proof. L is decidable in exponential time as we can exhaustively search all strings of Kt complexity at most b and length exactly a and check if there is one which has the specified values at the corresponding bit positions. Indeed, using the definition of Kt complexity and an efficient universal machine, a list containing all such strings can be generated in time $\text{poly}(2^b)$, which is at most exponential in the input length 1^b . In turn, checking that a string of length a satisfies the requirement takes time at most exponential in the total input length, since each index i_j is a string of length $\lceil \log a \rceil$. \square

Since $\text{EXP} \subseteq \text{Formula}[\text{poly}]$ by assumption, L has polynomial-size formulas. Assume without loss of generality that L has formulas of size $O(\ell^k)$ for some constant k , where ℓ is its total input length. We choose $\beta = \varepsilon/100k$.

We are ready to describe a low-complexity reduction from **Gap-MKtP** $[2^{\beta n}, 2^{\beta n} + cn]$ to L . First, we use the error-correcting code to compute z from w , as described above. Then we apply the following sampling procedure. We sample uniformly and independently $r = 2^{2\beta n}$ indices $i_1, \dots, i_r \in_R [M]$, where $M = bN$. We then form the string y encoding the tuple

$$(M, 1^{2^{\beta n} + c_1 n}, (i_1, z_{i_1}), \dots, (i_r, z_{i_r})),$$

where $c_1 > c_0 \geq 1$ is provided by Claim 12. Note that this is a string of length $\ell(N) \leq N^{\varepsilon/10k}$.

Claim 14. *The following implications hold:*

- (a) *If $w \in \{0, 1\}^N$ is a positive instance of **Gap-MKtP** $[2^{\beta n}, 2^{\beta n} + cn]$, then $y \in L$ with probability 1.*
- (b) *If $w \in \{0, 1\}^N$ is a negative instance of **Gap-MKtP** $[2^{\beta n}, 2^{\beta n} + cn]$, then $y \notin L$ with probability $> 1/2$.*

Proof. If w is a YES instance, we have by Claim 11 that $\text{Kt}(z) \leq 2^{\beta n} + c_0 n \leq 2^{\beta n} + c_1 n$. In this case, z is a string of length M that has the specified values at the specified bit positions, regardless of the random positions that are sampled by the reduction. Consequently, $y \in L$ with probability 1.

For the claim about NO instances, as previously established in Claim 12, we have that $\text{Kt}(z') > 2^{\beta n} + c_1 n$ for any z' such that $|z'| = |z| = M$ and $\Pr_{i \in_R [M]} [z'_i \neq z_i] \leq \delta$. Now consider any string z'' of length M such that $\text{Kt}(z'') \leq 2^{\beta n} + c_1 n$. For such a string z'' , for each $j \in [r]$, the probability that the random projection satisfies $z''_{i_j} = z_{i_j}$ (where $i_j \in_R [M]$) is at most $1 - \delta$. Hence the probability that z'' agrees with z at all the specified bit positions is at most $(1 - \delta)^r \leq \exp(-\delta r) \leq \exp(-\delta 2^{2\beta n})$. By a union bound over all strings z'' with $\text{Kt}(z'') \leq 2^{\beta n} + c_1 n$, the probability that there exists a string z'' with Kt complexity at most $2^{\beta n} + c_1 n$ which is consistent with the values at the specified bit positions is exponentially small in n . Hence with high probability $y \notin L$. \square

To sum up, there is a randomized reduction from $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn]$ over inputs of length N to instances of L of length $\ell(N) \leq N^{\varepsilon/10k}$. Now let $\{F_{\ell(N)}\}_{N \geq 1}$ be a sequence of U_2 -formulas of size $O(\ell^k)$ for L . Our randomized formulas $\mathbf{G}(\cdot)$ for Gap-MKtP compute as follows.

1. $\mathbf{G}(w) = \bigwedge_{j=1}^N \mathbf{G}^{(j)}(w)$, where each $\mathbf{G}^{(j)}$ is an independent copy.
2. Each $\mathbf{G}^{(j)}(w)$ is a randomized formula of the form $G^{(j)}(w, \mathbf{i}_1, \dots, \mathbf{i}_r)$ that first computes z from w , then computes y from z using the (random) input indices $\mathbf{i}_1, \dots, \mathbf{i}_r \in \{0, 1\}^{\log M}$, and finally applies F_ℓ to y .

It follows from Claim 14 using the independence of each $\mathbf{G}^{(j)}$ that

$$\Pr[\mathbf{G}(w) \text{ is incorrect}] < 2^{-N},$$

where the probability is taken over the choice of the random input of G . Consequently, by a union bound there is a fixed choice $\gamma \in \{0, 1\}^*$ of the randomness of G (corresponding to the positions of the different random projections) such that the *deterministic* formula G_γ obtained from G and γ is correct on *every* input string w .

Claim 15. *Each deterministic sub-formula $G_\gamma^{(j)}(w)$ can be computed by a U_2 -formula extended with parities at the leaves of size at most $O(\ell(N)^k) \leq N^{\varepsilon/2}$.*

Proof. Note that each bit of z can be computed from the input string w using an appropriate parity function (as described in Theorem 10). We argue that the leaves of $G_\gamma^{(j)}$ are precisely the leaves of the U_2 -formula F_ℓ replaced by appropriate literals, constants, or parities. Recall that $G_\gamma^{(j)}$ applies F_ℓ to the string y obtained from z . However, since γ is fixed, the positions of z that are projected in order to compute y are also fixed, and so are the substrings of y describing the corresponding positions. Consequently, the size (i.e. number of leaves) of each $G_\gamma^{(j)}$ is at most the size of F_ℓ , which proves the claim. \square

It follows from this claim that $G_\gamma(w)$ can be computed by a formula containing at most $N^{1+\varepsilon}$ leaves, and hence $\text{Gap-MKtP}[2^{\beta n}, 2^{\beta n} + cn] \in U_2\text{-Formula-}\oplus[N^{1+\varepsilon}]$. (Observe that we have used in a crucial way that the derandomized sub-formulas do not need to compute address functions to generate y from z .) This completes the proof of Theorem 1 Part (2).

3.2 Completing the proof of Theorem 1

In this Section, we discuss how the argument presented in Section 3.1 can be adapted to establish the remaining items of Theorem 1.

First, note that Items (5) and (6) immediately follow from Item (2). This is because a parity gate over at most N input variables can be computed by B_2 -formulas of size $O(N)$ and by U_2 -formulas of size $O(N^2)$. Consequently, using that formula size is measured with respect to the number of leaves, we immediately get $U_2\text{-Formula-}\oplus[s(N)] \subseteq B_2\text{-Formula}[s(N) \cdot N]$ and $U_2\text{-Formula-}\oplus[s(N)] \subseteq U_2\text{-Formula}[s(N) \cdot N^2]$.

In order to get Item (1), it is sufficient to compute an error-correcting code as in Theorem 10 using circuits of (almost) linear size. In other words, we need the entire codeword (and not just each output bit) to be computable from the input message using a circuit of size $O(N)$. The existence

of such codes is well-known [SS96, Spi96]. The rest of the reduction produces an *additive* overhead in circuit size of at most $N^{1+\varepsilon}$ gates.

Finally, to establish Item (4), we use the following construction from [Tel18].

Theorem 16 (Computing ECCs in parallel using majorities and few wires [Tel18]). *For every depth $d' \geq 1$ there are constants $\delta(d') > 0$ and $b(d') \geq 1$ and a sequence $\{E_N\}_{N \in \mathbb{N}}$ of error-correcting codes $E_N: \{0, 1\}^N \rightarrow \{0, 1\}^M$ with the following properties:*

- $E_N(x)$ can be computed by a uniform deterministic algorithm running in time $\text{poly}(N)$.
- $M(N) = b \cdot N$.
- Any codeword $E_N(x) \in \{0, 1\}^M$ that is corrupted on at most a δ -fraction of coordinates can be uniquely decoded to x by a uniform deterministic algorithm D running in time $\text{poly}(M)$.
- $E_N(x) \in \{0, 1\}^M$ can be computed by a multi-output circuit from $\text{MAJ}_{2d'}^0[O(N^{1+(2/d')})]$, where circuit size is measured by number of wires.

Following the steps of the reduction described in Section 3.1, under the assumption that $\text{EXP} \subseteq \text{MAJ}_d^0[\text{poly}]$ the final depth of the circuit solving Gap-MkP is $2d' + d + 1$, where the terms in this sum correspond respectively to the computation of the error-correcting code (for a choice of $d' \geq 1$), each (circuit) $G_\alpha^{(j)}$, and the topmost AND gate in G_α (constant bits can be produced in depth 1 from input literals). Similarly, the overall size (number of wires) of the circuit is $O(N^{1+(2/d')}) + O(N^{1+\varepsilon}) + O(N) \leq N^{1+(2/d')+\varepsilon}$.

Item (3) is established in the obvious way given the previous explanations. Item (8) uses that parity gates can be simulated using mod 6 gates.

Finally, we deal with case (7), which refers to branching program complexity. First, note that the parity of n bits can be computed by a branching program of size $O(n)$. In addition, if $f(x) = g(h_1(x), \dots, h_k(x))$, each h_i has a branching program of size s , and g has a branching program of size t , then f has a branching program of size $\ell = O(t \cdot s)$. Finally, a conjunction of N branching programs of size ℓ has branching program size at most $O(N \cdot \ell)$. Combining these facts in the natural way yields case (7). This completes the proof of all cases in Theorem 1.

4 Hardness Magnification via Anti-Checkers

4.1 Proof of Theorem 4 (Magnification for MCSP)

In this Section, we derive Theorem 4 from Lemma 17, whose proof appears in Section 4.2. Informally, an anti-checker (cf. [LY94]) for a function f is a multi-set of input strings such that any circuit of bounded size that does not compute f is incorrect on at least one of these strings.

Lemma 17 (Anti-Checker Lemma). *If $\text{NP} \subseteq \text{Circuit}[\text{poly}]$ there is a constant $k \in \mathbb{N}$ for which the following hold. For every sufficiently small $\beta > 0$, there is a circuit C of size $\leq 2^{n+k\beta n}$ that when given as input a truth-table $\text{tt}(f) \in \{0, 1\}^N$, where $f: \{0, 1\}^n \rightarrow \{0, 1\}$, outputs $t = 2^{10\beta n}$ strings $y_1, \dots, y_t \in \{0, 1\}^n$ such that if $f \notin \text{Circuit}[2^{\beta n}]$ then every circuit of size $\leq s$ where $s = 2^{\beta n}/10n$ fails to compute f on at least one of these strings.*

The Anti-Checker Lemma is a powerful tool that might be of independent interest. It says that anti-checkers of bounded size for functions requiring circuits of size $2^{o(n)}$ can be produced in time that is almost-linear in the size of the function (viewed as a string), under the assumption that circuit lower bounds do not hold.⁶

Proof of Theorem 4. Assume that $\text{NP} \subseteq \text{Circuit}[\text{poly}]$. We prove that for every given $\varepsilon > 0$ there exists a small enough $\beta > 0$ such that $\text{Gap-MCSP}[2^{\beta n}/10n, 2^{\beta n}] \in \text{Circuit}[N^{1+\varepsilon}]$.

We consider the problem **Succinct-MCSP**, defined next. Its input instances are of the form $\langle 1^n, 1^s, 1^t, (x_1, b_1), \dots, (x_t, b_t) \rangle$, where $x_i \in \{0, 1\}^n$ and $b_i \in \{0, 1\}$, $i \in [t]$. Note that each instance can be encoded by a string of length exactly $m = n + 1 + s + 1 + t + 1 + t \cdot (n + 1)$. An input string is a positive instance if and only if it is in the appropriate format and there exists a circuit D over n input variables and of size at most s such that $D(x_i) = b_i$ for all $i \in [t]$. Note that the problem is in NP as a function of its total input length m . Under the assumption that NP is easy for non-uniform circuits, there exists $\ell \in \mathbb{N}$ such that **Succinct-MCSP** can be solved by circuits E_m of size m^ℓ on every large enough input length m .

Take $\beta = \varepsilon/(100 \cdot \ell \cdot k)$, where k is the constant from Lemma 17. In order to construct a circuit for **Gap-MCSP**, first we reduce this problem to an instance of **Succinct-MCSP** of length m using Lemma 17, then we invoke the m^ℓ -sized circuit for this problem. More precisely, on an input $f: \{0, 1\}^n \rightarrow \{0, 1\}$, we use the circuit C (as in Lemma 17) to produce a list of strings $y_1, \dots, y_t \in \{0, 1\}^n$, generate from this list and f the input instance $z = \langle 1^n, 1^s, 1^t, ((y_1, f(y_1)), \dots, (y_t, f(y_t))) \rangle$, for parameters $s = 2^{\beta n}/10n$, $t = 2^{10\beta n}$, $m = \text{poly}(n) \cdot 2^{10\beta n}$, and output $E_m(z)$.

Correctness follows immediately from Lemma 17 and our choice of parameters. Indeed, if $f \in \text{Circuit}[2^{\beta n}/10n]$ then no matter the choice of y_1, \dots, y_t the circuit E_m accepts z thanks to our choice of $s = 2^{\beta n}/10n$. On the other hand, when $f \notin \text{Circuit}[2^{\beta n}]$ then by Lemma 17 every circuit of size s fails on some string from the list, and consequently $E_m(z) = 0$.

We upper bound the total circuit size using the choice of β . Circuit C has size at most $2^{n+k\beta n} \leq N^{1+\varepsilon}/3$. In addition, producing the input z can be done from f and y_1, \dots, y_t by a circuit of size at most $O(t \cdot N) \leq N^{1+\varepsilon}/3$, since each address function can be computed in linear size $O(N)$ (see e.g. [Weg87]). Finally, E_m has size at most $m^\ell \leq N^{1+\varepsilon}/3$. Overall, it follows that **Gap-MCSP** $[2^{\beta n}/10n, 2^{\beta n}]$ is computable by circuits of size $N^{1+\varepsilon}$. \square

4.2 Proof of Lemma 17 (Anti-Checker Lemma)

This section is dedicated to the proof of Lemma 17. This completes the proof of Theorem 4. We start with a high-level exposition of the argument.

Proof Idea. We take $\beta \rightarrow 0$, for simplicity of the exposition. In principle, the challenge is to *construct* the list of strings from the description of f using a circuit of size $N^{1+o(1)}$, given that the *existence* of such strings is guaranteed by the work of [LY94]. But it is not clear how to use this existential result and the assumption that NP has polynomial size circuits to construct *almost-linear* size circuits for this task. In order to achieve this, we use a *self-contained* argument that produces the strings one by one until very few circuits of bounded size are consistent with the values of f on the partial list of strings. We then find polynomially many *additional* strings that eliminate the

⁶We have made no attempt to optimize the constants in Lemma 17.

remaining circuits, completing the list of strings.⁷

To produce the i -th string $y_i \in \{0, 1\}^n$ given $y_1, \dots, y_{i-1} \in \{0, 1\}^n$ and f , we estimate the number of circuits of size $\leq 2^{\beta n}/10n$ that agree with f over all strings in $\{y_1, \dots, y_i\}$. We show that *some string* y_i will reduce the number of consistent circuits from the previous round by a factor of (roughly) $1 - 1/n$ if there are at least (roughly) n^2 surviving circuits (this is a combinatorial existential proof that relies on the lower bound on the circuit complexity of f). As a consequence, it will be possible to show that at most $2^{O(\beta n)} = N^{o(1)}$ rounds suffice to produce the required set of strings (modulo handling the few surviving circuits). The existence of a *good* string y_i is at the heart of our argument, and we defer the exposition of this result to the formal proof.

In each round, we exhaustively check each of the N candidate strings y_i . As we will explain soon, estimating the number of surviving circuits after picking a new candidate string y_i can be done by a circuit of size $N^{o(1)}$ given access to y_1, \dots, y_i and to the corresponding bits $f(y_1), \dots, f(y_i)$.⁸ In summary, there are $N^{o(1)}$ rounds, and in each one of them we can find a good string y_i using a circuit of size $N^{1+o(1)}$. We remark that it will also be possible to produce the additional strings in circuit complexity $N^{o(1)}$, so that the complete list y_1, \dots, y_t can be computed from f by a circuit of size $N^{1+o(1)}$.

It remains to explain how to fix a good string in each round. We simply pick the most promising string, using that we can upper bound the complexity of estimating the number of surviving circuits. The latter relies on the assumed inclusion $\text{NP} \subseteq \text{Circuit}[\text{poly}]$. Indeed, from this assumption it follows that the polynomial hierarchy $\text{PH} \subseteq \text{Circuit}[\text{poly}]$, and it is known that *relative approximate counting* can be done in the polynomial hierarchy.⁹ Crucially, as described in the paragraph above, the input length of each sub-problem that we need to solve is $\leq N^{o(1)}$ (using that i is at most $N^{o(1)}$), so a polynomial overhead will not be an issue when solving a sub-task of input length $N^{o(1)}$. This completes the sketch of the proof.

We proceed with a formal proof of Lemma 17. Let R be a polynomial-time relation, where $R \subseteq \bigcup_m \{0, 1\}^m \times \{0, 1\}^{q(m)}$ for some polynomial q . For every x , we use $R_{\#}(x)$ to denote $|\{y \in \{0, 1\}^{q(|x|)} : (x, y) \in R\}|$. A randomized algorithm Π is called an (ε, δ) -*approximator* for R if for every input x it holds that

$$\Pr[|\Pi(x) - R_{\#}(x)| \geq \varepsilon(|x|) \cdot R_{\#}(x)] \leq \delta(|x|).$$

Theorem 18 (Relative approximate counting in BPP^{NP} ([Sto83]; see e.g. [Gol08, Section 6.2.2])). *For every polynomial-time relation R and every polynomial p , there exists a probabilistic polynomial-time algorithm \mathbf{A} with access to a SAT oracle that is an $(1/p(m), 2^{-p(m)})$ -approximator for R over inputs x of length m .*

Corollary 19. *Assume $\text{NP} \subseteq \text{Circuit}[\text{poly}]$. For every polynomial-time relation R and for each $m \geq 1$, there is a multi-output circuit $C_R: \{0, 1\}^m \rightarrow \{0, 1\}^{\text{poly}(m)}$ of polynomial size such that on every input $x \in \{0, 1\}^m$,*

$$(1 - 1/m^2) \cdot R_{\#}(x) \leq C_R(x) \leq (1 + 1/m^2) \cdot R_{\#}(x).$$

⁷In particular, our argument implies the worst-case version of the anti-checker result from [LY94] with slightly different parameters.

⁸Technically speaking, projecting $f(y_i) \in \{0, 1\}$ from the input string $f \in \{0, 1\}^N$ and the address $y \in \{0, 1\}^n$ already takes circuit complexity $\Omega(N)$. However, since we are trying all possible strings y_i , the corresponding bit positions of f can be directly hardwired.

⁹In our formal proof, we take a slightly more direct route to compute the relative approximations.

Proof. This follows from Theorem 18 (using $p(m) = m^2$) by non-uniformly fixing the randomness of the algorithm, replacing the SAT oracle using the assumption that NP has small circuits, and translating the resulting deterministic algorithm into a boolean circuit. \square

We define a relation Q . The first input x is of the form $\langle 1^n, 1^s, 1^i, 1^t, (z_1, b_1), \dots, (z_i, b_i) \rangle$, where $z_j \in \{0, 1\}^n$ and $b_j \in \{0, 1\}$ for $1 \leq j \leq i$, and $t = 2^{10\beta n}$ (t is used here to pad the input appropriately). The second input is a string w of length $m^{1/5}$ (for $m = |x|$) that is interpreted as a boolean circuit C_w over n input variables and of size at most s . We let $(x, w) \in Q$ if and only if $C_w(z_j) = b_j$ for all $j \in [i]$. Note that Q is a polynomial-time relation.

We employ circuits obtained from Corollary 19 using parameters $s = 2^{\beta n}/10n$ and $1 \leq i \leq t$, where $t = 2^{10\beta n}$. The following result is immediate from Corollary 19 given that for our choice of parameters $m = \text{poly}(2^{\beta n})$.

Proposition 20 (Circuits for approximate counting). *There is a constant $k_1 \in \mathbb{N}$ for which the following holds. For every $n \geq 1$, let $s = 2^{\beta n}/10n$, $t = 2^{10\beta n}$, $1 \leq i \leq t$. Then there is a multi-output circuit $C_{n,i}$ of size $\leq 2^{k_1\beta n}$ that outputs $\leq 2^{k_1\beta n}$ bits such that on every input $a = ((z_1, b_1), \dots, (z_i, b_i)) \in \{0, 1\}^{i \cdot (n+1)}$,*

$$(1 - 1/n^{10}) \cdot Q_{\#}(x) \leq C_{n,i}(a) \leq (1 + 1/n^{10}) \cdot Q_{\#}(x),$$

where $x = x(a)$ is defined from the string a and from our choice of parameters in the obvious way.

The next step is to guarantee that once just a few circuits remain consistent with f over our partial list of strings (as described in the proof sketch above), we can efficiently find a small number of strings to eliminate all of them.

Lemma 21 (Listing the remaining circuits). *Assume $\text{NP} \subseteq \text{Circuit}[\text{poly}]$. There exists a constant $k_2 \in \mathbb{N}$ for which the following holds. Let $a = ((z_1, b_1), \dots, (z_{t'}, b_{t'}))$, where $t' \leq t$, and $x = x(a)$ be the corresponding input of Q . There is a circuit $D_{n,t'}$ of size $\leq 2^{k_2\beta n}$ such that if $Q_{\#}(x) \leq n^3$, then $D_{n,t'}(a)$ outputs a string describing all such circuits.*

Proof. It follows from $\text{NP} \subseteq \text{Circuit}[\text{poly}]$ using a standard argument that $\text{PH} \subseteq \text{Circuit}[\text{poly}]$. In addition, it is not hard to define a relation in PH (using a padded input containing the string 1^t) that checks if a given input a satisfies $Q_{\#}(x(a)) \leq n^3$. Consequently, checking if a string λ describes a list of such circuits for a can be done by a circuit of size at most $\text{poly}(t)$. Using again that $\text{NP} \subseteq \text{Circuit}[\text{poly}]$ and a self-reduction, we obtain circuits $D_{n,t'}$ as in the statement of the lemma. \square

Lemma 22 (Completing the list of strings). *There is a constant $k_3 \in \mathbb{N}$ for which the following holds. For every $n \geq 1$ there is a circuit E_n of size $\leq 2^{n+k_3\beta n}$ that given access to a truth-table $f \in \{0, 1\}^{2^n}$ and a string $w \in \{0, 1\}^{2^{\beta n}}$ describing a circuit C_w of size $s \leq 2^{\beta n}/10n$ that does not compute f , $E_n(f, w)$ outputs a string y such that $C(y) \neq f(y)$.*

Proof. First, E_n evaluates C_w on every string $z \in \{0, 1\}^n$. This can be easily done by a circuit of size $2^n \cdot \text{poly}(|w|)$ under a reasonable encoding of the circuit C_w . Then E_n inspects one-by-one each string z and stores the first string where C_w and f differ. Note that a circuit of size $\leq 2^n \cdot \text{poly}(n)$ can print this string from the truth-table of f and C_w . It follows that the overall complexity of E_n is $2^{n+k_3\beta n}$ for some constant k_3 . \square

The previously established results will allow us to find in each round a string y_i that significantly reduces the number of remaining circuits (while at least one such string exists), and then to complete the list so that no circuit of bounded size is consistent with all strings in the final list. We show next that if f is hard and a reasonable number of circuits of bounded size are consistent with the current list of strings, then a good string y_i exists.

For convenience, we introduce a function to capture the fraction of strings encoding circuits that are consistent with a set of inputs and their corresponding labels. Given $a = ((z_1, b_1), \dots, (z_i, b_i))$, let $x = x(a)$ be the corresponding input to Q under our choice of parameters. Furthermore, let $m = |x|$, and recall that $Q \subseteq \bigcup_{m \geq 1} \{0, 1\}^m \times \{0, 1\}^{m^{1/5}}$. In order to maintain the same underlying domain size when considering the fraction of consistent circuits, we assume without loss of generality using appropriate padding that the encoding of x has a fixed length $m = m(n)$ for each choice of n (i.e., the choice of $1 \leq i \leq t$ does not affect $m^{1/5}$). In addition, we can take $m(n) \leq 2^{11\beta n}$, which will be useful when upper bounding the number of necessary rounds. We let $\phi(a) \in [0, 1]$ denote the ratio $Q_{\#}(x(a))/2^{m^{1/5}}$. (Thus in our formal argument we count circuits using their descriptions as binary strings.)

Lemma 23 (Existence of a good string y_i). *For every integer $i \geq 1$ and for every $z_1, \dots, z_{i-1} \in \{0, 1\}^n$, let $a = ((z_1, f(z_1)), \dots, (z_{i-1}, f(z_{i-1})))$. If*

$$f \notin \text{Circuit}[2^{\beta n}] \quad \text{and} \quad Q_{\#}(x(a)) \geq 4n^2,$$

then there is some string $y_i \in \{0, 1\}^n$ such that if a' denotes the sequence a augmented with $(y_i, f(y_i))$, then

$$\phi(a') \leq \phi(a) \cdot (1 - 1/2n).$$

Proof. The argument is inspired by a combinatorial principle discussed in [Kra95]. Consider the tuple a and the string $x = x(a)$ as in the statement of the lemma. Moreover, let $Q(x) = \{w \in \{0, 1\}^{m^{1/5}} : (x, w) \in Q\}$. For convenience, let $r = |Q(x)| = Q_{\#}(x) \geq 4n^2$, using our assumption. Define an auxiliary undirected bipartite graph $G = (L, R, E)$ as follows. Set $L = \{0, 1\}^n$, $R = \binom{Q(x)}{n}$, and $(y, \{w^1, \dots, w^n\}) \in E(G)$ if and only if for $\leq n/2$ of the circuits C_{w^i} we have $f(y) = C_{w^i}(y)$.

Note that for any right vertex $v = (w^1, \dots, w^n) \in R$ there is a left vertex $y \in L$ such that $(y, v) \in E$. If not, then $D = \text{Majority}_n(C_{w^1}(x), \dots, C_{w^n}(x))$ is a circuit that computes f on every input string y . The size of D is at most $n \cdot (2^{\beta n}/10n) + 5n \leq 2^{\beta n}$, using the definition of Q and that the majority function can be computed (with room to spare) by a circuit of size at most $5n$ [Weg87]. This contradicts the hardness of f .

By an averaging argument, there is a left vertex y^* that is connected to at least $|R|/|L| = \binom{r}{n}/2^n$ vertices in R . We show below (Claim 24) that for at least $r/2n$ strings $w \in Q(x)$, the corresponding circuit C_w satisfies $C_w(y^*) \neq f(w^*)$. This implies that by taking y^* as the string y_i described in the statement of the lemma, we get $Q_{\#}(x(a')) \leq r - r/2n = r(1 - 1/2n)$, and consequently

$$\phi(a') = \frac{Q_{\#}(x(a'))}{2^{m^{1/5}}} \leq \frac{r(1 - 1/2n)}{2^{m^{1/5}}} = \frac{Q_{\#}(x(a)) \cdot (1 - 1/2n)}{2^{m^{1/5}}} = \phi(a) \cdot (1 - 1/2n).$$

Claim 24. *Let $y^* \in L$ be a left-vertex connected to at least $\binom{r}{n} \cdot 2^{-n}$ right-vertices in R , where $r \geq 4n^2$ and n is sufficiently large. Then, for at least $r/2n$ distinct strings $w \in Q(x)$, we have $C_w(y^*) \neq f(y^*)$.*

Proof. The claim follows using a standard counting argument. If the conclusion were false, the vertex y^* would be connected to *strictly* less than (assuming for simplicity that n is even and $r/2n$ is an integer)

$$\sum_{j=0}^{n/2} \binom{r/2n}{\frac{n}{2}+j} \cdot \binom{r}{\frac{n}{2}-j} \leq \binom{r}{n} \cdot 2^{-n} \quad (\text{as upper bounded below})$$

vertices in R , which is contradictory. It remains to verify this inequality, which can be done using some careful estimates. First, note that

$$\begin{aligned} \sum_{j=0}^{n/2} \binom{r/2n}{\frac{n}{2}+j} \binom{r}{\frac{n}{2}-j} &\leq \sum_{j=0, \dots, \frac{n}{2}-1} \frac{r^n / (2n)^{\frac{n}{2}+j}}{(\frac{n}{2}+j)! (\frac{n}{2}-j)!} + \frac{r^n}{n! (2n)^n} && (\text{using } \binom{n}{k} \leq \frac{n^k}{k!}) \\ &\leq \sum_{j=0, \dots, \frac{n}{2}-1} \frac{e^n r^n / (2n)^{\frac{n}{2}+j}}{e^2 (\frac{n}{2}+j)^{\frac{n}{2}+j} (\frac{n}{2}-j)^{\frac{n}{2}-j}} + \frac{e^n r^n}{e n^n (2n)^n} && (\text{since } e \left(\frac{n}{e}\right)^n \leq n!) \\ &\leq \sum_{j=0, \dots, \frac{n}{2}-1} \frac{e^n r^n / (2n)^{\frac{n}{2}}}{e^2 (\frac{n}{2})^j (\frac{n}{2}+j)^{\frac{n}{2}} (\frac{n}{2}-j)^{\frac{n}{2}}} + \frac{e^n r^n}{e n^n (2n)^n} && (*) \end{aligned}$$

By considering the cases $j < \frac{n}{4}$ and $\frac{n}{2} > j \geq \frac{n}{4}$, we get $(\frac{n}{2})^j ((\frac{n}{2})^2 - j^2)^{\frac{n}{2}} \geq (n/8)^{3n/4}$, so

$$\begin{aligned} (*) &\leq \sum_{j=0, \dots, \frac{n}{2}-1} \frac{e^n r^n}{e^2 (n/8)^{3n/4} (2n)^{n/2}} + \frac{e^n r^n}{e n^n (2n)^n} \\ &\leq \frac{n e^n r^n}{e^2 (n/8)^{3n/4} (2n)^{n/2}} \leq \frac{\sqrt{2\pi} r^n}{e^2 n^{1/2} (2n)^n} \leq \frac{\sqrt{2\pi} r^r r^{1/2}}{e^2 (r-n)^{r-n+1/2} n^{n+1/2}} \cdot \frac{1}{2^n} \\ &\leq \binom{r}{n} / 2^n, \end{aligned}$$

where n is assumed to be sufficiently large, $r > n$, and the last inequality makes use of Stirling's approximation $\sqrt{2\pi} (\frac{n}{e})^n n^{1/2} \leq n! \leq e (\frac{n}{e})^n n^{1/2}$. This completes the proof of Claim 24. \square

This completes the proof of Lemma 23. \square

We are ready to combine these results and define a circuit C of size $\leq 2^{n+k\beta n}$ with the property stated in Lemma 17. This circuit on an input $f \in \{0, 1\}^N$ where $N = 2^n$ computes as follows.

1. C sequentially computes the string $a^{(i)} = (y_1, f(y_1)), \dots, (y_i, f(y_i))$ for $1 \leq i \leq t'$ and $t' = 2^{10\beta n} - n^3$.

During stage i , C inspects all strings $y \in \{0, 1\}^n$, using the circuit $C_{n,i}$ (Proposition 20) to fix y_i as the string that minimizes $C_{n,i}(a^{(i)})$.

2. C uses the circuit $D_{n,t'}$ (Lemma 21) to print the descriptions of n^3 circuits of size at most $s = 2^{\beta n} / 10n$.

3. Finally, C invokes n^3 copies of the circuit E_n (Lemma 22) to complete the list y_1, \dots, y_t of strings, where $t = t' + n^3 = 2^{10\beta n}$.

Correctness of the construction follows from the properties of the circuits $C_{n,i}$, $D_{n,t'}$, and E_n in combination with Lemma 23. More precisely, if $f \notin \text{Circuit}[2^{\beta n}]$, then for every $1 \leq i \leq t'$, either $\phi(a^{(i)}) \leq (1 - 1/4n)^i$ or $Q_{\#}(x(a^{(i-1)})) < 4n^2$. To see this, note that if the latter condition does not hold, then for some string y^* as in Lemma 23 we get with respect to the corresponding extension $a^{(i)}$ that $\phi(a^{(i)}) \leq \phi(a^{(i-1)}) \cdot (1 - 1/2n)$. Since C tries all strings during its computation in step 1 when in stage i , and the relative approximation given by circuit $C_{n,i}$ is sufficiently precise, we are guaranteed in this case (using an inductive argument) to fix a string y_i such that $\phi(a^{(i)}) \leq \phi(a^{(i-1)}) \cdot (1 - 1/4n) \leq (1 - 1/4n)^i$. On the other hand, if the condition $Q_{\#}(x(a^{(i-1)})) < 4n^2$ holds for some $i \leq t'$, then by monotonicity it is maintained until we reach $i = t'$. Consequently, using that initially $\phi(\epsilon) = 1$, $t' = 2^{10\beta n} - n^3$, $m(n) \leq 2^{11\beta n}$, and recalling that the second input of the relation Q has length $m^{1/5}$ and that this parameter is related to the definition of ϕ , when C reaches $i = t'$ at the end of step 1 we have

$$\begin{aligned} Q_{\#}(x(a^{(t')})) &\leq \max\{4n^2, (1 - 1/4n)^{t'} \cdot 2^{m^{1/5}}\} \\ &\leq n^3. \end{aligned}$$

This implies using Lemmas 21 and 22 and the description of C that if $f \notin \text{Circuit}[2^{\beta n}]$ then every circuit of size at most $s = 2^{\beta n}/10n$ disagrees with f on some input string among y_1, \dots, y_t .

Finally, we upper bound the circuit size of C . For every $i \leq t'$ in step 1 and each string $y \in \{0, 1\}^n$, C feeds $C_{n,i}$ with the appropriate bit in the input string f and the previously computed string $a^{(i-1)}$. This produces an estimate $v_y \in \mathbb{N}$ represented as a string of length $2^{O(\beta n)}$ that is stored as a pair (y, v_y) . Using Proposition 20, all pairs (y, v_y) can be simultaneously computed by a circuit of size at most $2^n \cdot 2^{O(\beta n)}$. By inspecting each such pair in sequence, C can pick the string $y_i \in \{0, 1\}^n$ minimizing v_{y_i} using a sub-circuit of size $2^n \cdot \text{poly}(2^{O(\beta n)})$. Also note that the bit $f(y_i)$ can be easily computed from y_i and f by a circuit of size $O(N \log N)$. Therefore, each stage i can be done by a circuit of size at most $2^{n+O(\beta n)}$, and since there are $t' \leq 2^{10\beta n}$ stages, the computation in step 1. can be done by a circuit of size $2^{n+O(\beta n)}$. Lastly, steps 2 and 3 can be each implemented by a circuit of size at most $2^{O(\beta n)}$ using the upper bounds on circuit size provided by Lemmas 21 and 22, respectively, and the description of C . It follows that the overall circuit size of C is at most $2^{n+k\beta n}$, where k is a constant that only depends on the circuits provided by the initial assumption that $\text{NP} \subseteq \text{Circuit}[\text{poly}]$.

A remark on formulas vs. circuits. An obstacle to producing the anti-checker using formulas of size $N^{1+o(1)}$ under the assumption that $\text{NP} \subseteq \text{Formula}[\text{poly}]$ comes from the sequential aspect of the construction. A string y_j produced after the j -th round is inspected during each subsequent round of the construction. In the case of formulas, the corresponding bits need to be recomputed each time, and the overall complexity becomes prohibitive. (There are other intermediate computations that one may not be able to simulate so easily with sub-quadratic formulas, such as selecting the best string y_i during each round.)

4.3 The Anti-Checker Hypothesis

The existence of anti-checkers of bounded size witnessing the hardness of Boolean functions is far from obvious. In this section, we explore consequences of a hypothetical phenomenon manifesting on a higher level: the existence of a small collection of anti-checker sets witnessing hardness of all hard functions. We show that a certain formulation of this Anti-Checker Hypothesis (AH) implies

unconditional lower bounds. Complementing this result, we prove unconditionally that (AH) holds for functions that are hard in the average case.

For simplicity, we adopt a concrete setting of parameters for the hypothesis and in the results presented in this section. Understanding the validity of (AH) with respect to other non-trivial setting of parameters would also be interesting.

The Anti-Checker Hypothesis (AH). *For every $\lambda \in (0, 1)$, there are $\varepsilon > 0$ and a collection $\mathcal{Y} = \{Y_1, \dots, Y_\ell\}$ of sets $Y_i \subseteq \{0, 1\}^n$, where $\ell = 2^{(2-\varepsilon)n}$ and each $|Y_i| = 2^{n^{1-\varepsilon}}$, for which the following holds.*

If $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and $f \notin \text{Circuit}[2^{n^\lambda}]$, then some set $Y \in \mathcal{Y}$ forms an anti-checker for f : For each circuit C of size $2^{n^\lambda}/10n$, there is an input $y \in Y$ such that $C(y) \neq f(y)$.

The Anti-Checker Hypothesis can be shown to imply the hardness of a specific meta-computational problem in NP (which is not necessarily NP-complete).

Definition 25 (Succinct MCSP). *Let $s, t: \mathbb{N} \rightarrow \mathbb{N}$ be functions. The Succinct Minimum Circuit Size Problem with parameters s and t , abbreviated Succinct-MCSP(s, t), is the problem of deciding given a list of $t(n)$ pairs (y_i, b_i) , where $y_i \in \{0, 1\}^n$ and $b_i \in \{0, 1\}$, if there exists a circuit C of size $s(n)$ computing the partial function defined by these pairs, i.e., $C(y_i) = b_i$ for every $i \in [t]$.*

Note that Succinct-MCSP(s, t) \in NP whenever s and t are constructive functions.

Theorem 26. *Assume (AH) holds, and let $\varepsilon = \varepsilon(\lambda) > 0$ be the corresponding constant for $\lambda = 1/2$. Then Succinct-MCSP($2^{n^{1/2}}/10n, 2^{n^{1-\varepsilon}}$) \notin Formula[poly]. In particular, NP $\not\subseteq$ Formula[poly].*

Proof. The proof is by contradiction. Take $\lambda = 1/2$ in the Anti-Checker Hypothesis, and let $\varepsilon = \varepsilon(\lambda) > 0$ be the given constant. In addition, let $F_m: \{0, 1\}^N \rightarrow \{0, 1\}$ be a formula of size m^k for Succinct-MCSP($2^{n^{1/2}}/10n, 2^{n^{1-\varepsilon}}$), where $m \leq \text{poly}(n) \cdot 2^{n^{1-\varepsilon}}$ is the total input length for this problem. We argue below that from these assumptions it follows that Gap-MCSP[$2^{n^{1/3}}, 2^{n^{2/3}}$] \in Formula[$N^{2-\delta}$] for some $\delta > 0$. This contradicts Theorem 5 if α is taken to be a sufficiently small constant, which completes the proof.

We define a formula $E: \{0, 1\}^N \rightarrow \{0, 1\}$ that solves Gap-MCSP[$2^{n^{1/3}}, 2^{n^{2/3}}$]. It projects the appropriate bits of the input f to produce $T = 2^{(2-\varepsilon)n}$ instances of Succinct-MCSP($2^{n^{1/2}}/10n, 2^{n^{1-\varepsilon}}$) obtained from f and from the collection \mathcal{Y} in the natural way. The formula E is defined as the conjunction of T independent copies of the formula F_m from above. Note that E has at most $T \cdot m^k \leq N^{2-\delta}$ leaves, where $\delta = \delta(\varepsilon) > 0$. Finally, it is easy to see that it correctly solves Gap-MCSP using our choice of parameters and (AH). \square

We say that a Boolean function f with n inputs is hard on average for circuits of size s if every circuit of size s fails to compute f on at least $1/s$ fraction of all inputs.

Proposition 27 (Average-Case AH). *For every $\lambda \in (0, 1)$ there is $\varepsilon > 0$ such that for every large enough $n \in \mathbb{N}$ there is a collection $\mathcal{Y} = \{Y_1, \dots, Y_\ell\}$ of $\ell = 2^n$ sets $Y_i \subseteq \{0, 1\}^n$ of size $2^{n^{1-\varepsilon}}$ for which the following holds. If $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is hard on average for circuits of size 2^{n^λ} , then some set $Y \in \mathcal{Y}$ constitutes an anti-checker for f : For each circuit C of size 2^{n^λ} there is a string $y \in Y$ such that $C(y) \neq f(y)$.*

Proof. Let \mathcal{H} be the set of all Boolean functions f over n inputs that are hard on average for circuits of size $s = 2^{n^\lambda}$. Then we can generate anti-checkers for $f \in \mathcal{H}$ by choosing n -bit strings uniformly at random: for each $i \in [2^n]$, we let \mathbf{Y}_i be the set obtained by sampling with repetition $2^{n^{1-\varepsilon}}$ random strings in $\{0, 1\}^n$, where $1 - \varepsilon > \lambda$. Then, for every large enough n , for each circuit C of size at most 2^{n^λ} and for each $f \in \mathcal{H}$,

$$\Pr[C|_{\mathbf{Y}_i} \equiv f|_{\mathbf{Y}_i}] \leq (1 - 1/2^{n^\lambda})^{2^{n^{1-\varepsilon}}} \leq \exp(-2^{n^{1-\varepsilon}}/2^{n^\lambda}).$$

Now by a union bound over all such circuits, for a fixed $f \in \mathcal{H}$ we get

$$\Pr[\mathbf{Y}_i \text{ is not an anti-checker set for } f] \leq \exp(O(n \cdot 2^{n^\lambda})) \cdot \exp(-2^{n^{1-\varepsilon}}/2^{n^\lambda}) < 1/4,$$

where the last inequality used our choice of ε . Finally,

$$\Pr[\exists f \in \mathcal{H} \text{ s.t. none of } \mathbf{Y}_1, \dots, \mathbf{Y}_{2^n} \text{ is an anti-checker set for } f] \leq 2^{2^n} \cdot (1/4)^{2^n} < 1.$$

There is therefore a collection \mathcal{Y} with the desired properties. □

Theorem 26 and Proposition 27 show a connection between establishing super-polynomial formula size lower bounds for NP and understanding the difference between worst-case and average-case collections of anti-checkers.

Acknowledgements

We thank Avishay Tal for bringing [Tal16] to our attention in connection to the problem of proving lower bounds against $U_2\text{-Formula}\oplus$. We are also grateful to Jan Krajíček for discussions related to hardness magnification. This work was supported in part by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2014)/ERC Grant Agreement no. 615075.

References

- [Aar17] Scott Aaronson. $P \stackrel{?}{=} NP$. *Electronic Colloquium on Computational Complexity* (ECCC), 24:4, 2017.
- [ABK⁺06] Eric Allender, Harry Buhrman, Michal Koucký, Dieter van Melkebeek, and Detlef Ronneburger. Power from random strings. *SIAM J. Comput.*, 35(6):1467–1493, 2006.
- [ABSRW04] Michael Alekhovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Pseudorandom generators in propositional proof complexity. *SIAM Journal of Computing*, 34(1):67–88, 2004.
- [AK10] Eric Allender and Michal Koucký. Amplifying lower bounds by means of self-reducibility. *J. ACM*, 57(3):14:1–14:36, 2010.
- [All01] Eric Allender. When worlds collide: Derandomization, lower bounds, and Kolmogorov complexity. In *Foundations of Software Technology and Theoretical Computer Science* FSTTCS, pages 1–15, 2001.

- [All17] Eric Allender. The complexity of complexity. In *Computability and Complexity - Essays Dedicated to Rodney G. Downey on the Occasion of His 60th Birthday*, pages 79–94, 2017.
- [BFT98] Harry Buhrman, Lance Fortnow, and Thomas Thierauf. Nonrelativizing separations. In *Conference on Computational Complexity (CCC)*, pages 8–12, 1998.
- [Bog18] Andrej Bogdanov. Small-bias require large formulas. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, 2018.
- [BS90] Ravi B. Boppana and Michael Sipser. The complexity of finite functions. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, pages 757–804. 1990.
- [CILM18] Marco Carmosino, Russell Impagliazzo, Shachar Lovett, and Ivan Mihajlin. Hardness amplification for non-commutative arithmetic circuits. In *Computational Complexity Conference (CCC)*, 2018.
- [DM16] Irit Dinur and Or Meir. Toward the KRW composition conjecture: Cubic formula lower bounds via communication complexity. In *Conference on Computational Complexity (CCC)*, pages 3:1–3:51, 2016.
- [FGHK16] Magnus Gausdal Find, Alexander Golovnev, Edward A. Hirsch, and Alexander S. Kulikov. A better-than- $3n$ lower bound for the circuit complexity of an explicit function. In *Symposium on Foundations of Computer Science (FOCS)*, pages 89–98, 2016.
- [Gol08] Oded Goldreich. *Computational Complexity - A Conceptual Perspective*. Cambridge University Press, 2008.
- [Hås98] Johan Håstad. The shrinkage exponent of de Morgan formulas is 2. *SIAM J. Comput.*, 27(1):48–64, 1998.
- [HS17] Shuichi Hirahara and Rahul Santhanam. On the average-case complexity of MCSP and its variants. In *Computational Complexity Conference (CCC)*, pages 7:1–7:20, 2017.
- [IKV18] Russell Impagliazzo, Valentine Kabanets, and Ilya Volkovich. The power of natural properties as oracles. In *Computational Complexity Conference (CCC)*, 2018.
- [IKW02] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. In search of an easy witness: exponential time vs. probabilistic polynomial time. *J. Comput. Syst. Sci.*, 65(4):672–694, 2002.
- [IM02] Kazuo Iwama and Hiroki Morizumi. An explicit lower bound of $5n - o(n)$ for boolean circuits. In *Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 353–364, 2002.
- [IMZ12] Russell Impagliazzo, Raghu Meka, and David Zuckerman. Pseudorandomness from shrinkage. In *Symposium on Foundations of Computer Science (FOCS)*, pages 111–119, 2012.

- [IPS97] Russell Impagliazzo, Ramamohan Paturi, and Michael E. Saks. Size-depth tradeoffs for threshold circuits. *SIAM J. Comput.*, 26(3):693–707, 1997.
- [Juk12] Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*. Springer, 2012.
- [Jus72] Jørn Justesen. Class of constructive asymptotically good algebraic codes. *IEEE Trans. Information Theory*, 18(5):652–656, 1972.
- [KC00] Valentine Kabanets and Jin-yi Cai. Circuit minimization problem. In *Symposium on Theory of Computing (STOC)*, pages 73–79, 2000.
- [Kra95] Jan Krajíček. Extensions of models of PV. In *ASL/Springer Series – Lecture Notes in Logic – Proceedings of the Logic Colloquium*, 1995.
- [Kra01] Jan Krajíček. On the weak pigeonhole principle. *Fundamenta Mathematicae*, 170(1-3):123–140, 2001.
- [Kra04] Jan Krajíček. Dual weak pigeonhole principle, pseudo-surjective functions, and provability of circuit lower bounds. *Journal of Symbolic Logic*, 69(1):265–286, 2004.
- [KRT17] Ilan Komargodski, Ran Raz, and Avishay Tal. Improved average-case lower bounds for de Morgan formula size: Matching worst-case lower bound. *SIAM J. Comput.*, 46(1):37–57, 2017.
- [KW16] Daniel M. Kane and Ryan Williams. Super-linear gate and super-quadratic wire lower bounds for depth-two and depth-three threshold circuits. In *Symposium on Theory of Computing (STOC)*, pages 633–643, 2016.
- [Lev84] Leonid Levin. Randomness conservation inequalities; information and independence in mathematical theories. *Information and Control*, 61:15–37, 1984.
- [LW13] Richard J. Lipton and Ryan Williams. Amplifying circuit lower bounds against polynomial time, with applications. *Computational Complexity*, 22(2):311–343, 2013.
- [LY94] Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *Symposium on Theory of Computing (STOC)*, pages 734–740, 1994.
- [MP17] Moritz Müller and Ján Pich. Feasibly constructive proofs of succinct weak circuit lower bounds. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:144, 2017.
- [MV15] Eric Miles and Emanuele Viola. Substitution-permutation networks, pseudorandom functions, and natural proofs. *J. ACM*, 62(6):46:1–46:29, 2015.
- [MW18] Cody Murray and R. Ryan Williams. Circuit lower bounds for nondeterministic quasipolytime: an easy witness lemma for NP and NQP. In *Symposium on Theory of Computing (STOC)*, pages 890–901, 2018.
- [Neč66] È. Nečiporuk. On a boolean function. *Doklady of the Academy of the USSR*, 169(4):765–766, 1966.

- [OS17] Igor Carboni Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit lower bounds, and pseudorandomness. In *Computational Complexity Conference (CCC)*, pages 18:1–18:49, 2017.
- [OS18] Igor C. Oliveira and Rahul Santhanam. Hardness magnification for natural problems. *Manuscript*, 2018.
- [Pic14] Ján Pich. *Complexity Theory in Feasible Mathematics*. PhD thesis, Charles University in Prague, 2014.
- [Pic15] Ján Pich. Circuit lower bounds in bounded arithmetics. *Annals of Pure and Applied Logic*, 166(1), 2015.
- [Raz91] Alexander A. Razborov. Lower bounds for deterministic and nondeterministic branching programs. In *Symposium on Fundamentals of Computation Theory (FCT)*, pages 47–60, 1991.
- [Raz94] Alexander A. Razborov. On provably disjoint NP-pairs. *Basic Research in Computer Science Center*, 1994.
- [Raz95] Alexander A. Razborov. Unprovability of lower bounds on circuit size in certain fragments of bounded arithmetic. *Izvestiya of Russian Academy of Science*, 59:201–224, 1995.
- [Raz15] Alexander A. Razborov. Pseudorandom generators hard for k -DNF resolution and polynomial calculus. *Annals of Mathematics*, 182(2):415–472, 2015.
- [RR97] Alexander A. Razborov and Steven Rudich. Natural proofs. *J. Comput. Syst. Sci.*, 55(1):24–35, 1997.
- [San09] Rahul Santhanam. Circuit lower bounds for merlin–arthur classes. *SIAM J. Comput.*, 39(3):1038–1061, 2009.
- [Spi96] Daniel A. Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Trans. Information Theory*, 42(6):1723–1731, 1996.
- [Sri03] Aravind Srinivasan. On the approximability of clique and related maximization problems. *J. Comput. Syst. Sci.*, 67(3):633–651, 2003.
- [SS96] Michael Sipser and Daniel A. Spielman. Expander codes. *IEEE Trans. Information Theory*, 42(6):1710–1722, 1996.
- [Sto83] Larry J. Stockmeyer. The complexity of approximate counting (Preliminary Version). In *Symposium on Theory of Computing (STOC)*, pages 118–126, 1983.
- [Tal14] Avishay Tal. Shrinkage of de morgan formulae by spectral techniques. In *Symposium on Foundations of Computer Science (FOCS)*, pages 551–560, 2014.
- [Tal16] Avishay Tal. The bipartite formula complexity of inner-product is quadratic. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:181, 2016.

- [Tel18] Roei Tell. Quantified derandomization of linear threshold circuits. In *Symposium on Theory of Computing (STOC)*, 2018.
- [Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.
- [Weg87] Ingo Wegener. *The complexity of Boolean functions*. Wiley, 1987.
- [Wil16] Ryan Williams. Natural proofs versus derandomization. *SIAM J. Comput.*, 45(2):497–529, 2016.

A Unconditional Lower Bounds for Gap-MKtP and Gap-MCSP

A.1 MKtP – A near-quadratic lower bound against U_2 -formulas

In this section, we provide the proof of Theorem 3.

Proof Idea. We employ the technique of random restrictions to show that Gap-MKtP requires near-quadratic size formulas. The idea is that, with high probability, a formula F of sub-quadratic size simplifies under a random restriction $\rho: [N] \rightarrow \{0, 1, *\}$. This will allow us to complete a fixed restriction ρ either to a string w^y of Kt complexity $\leq s_1$, or to a string w^n of Kt complexity $\geq s_2$. Because the simplified formula $F \upharpoonright_\rho$ depends on few input variables in $\rho^{-1}(*)$, if we define w^y and w^n appropriately $F \upharpoonright_\rho$ won't be able to distinguish the two instances. Consequently, F does not compute Gap-MKtP $_{[s_1, s_2]}$.

In order for this idea to work, we cannot use a truly random restriction. This is because our restrictions will set most of the variables indexed in $[N]$ to simplify a near-quadratic size formula, and a typical random restriction cannot be completed to a string of low Kt complexity. We use instead pseudorandom restrictions, which can be computed from a much smaller number of random bits. Previous work established that such restrictions also simplify sub-quadratic size formulas. As a consequence, we are able to extend any restriction in the support of a pseudorandom distribution of restrictions to either an “easy” or a “hard” string, as explained in the paragraph above. (We remark that in order to improve our parameter s_1 in Gap-MKtP $_{[s_1, s_2]}$, it is useful to compose a sequence of pseudodeterministic restrictions.)

We proceed with the technical details. Let $\rho: [N] \rightarrow \{0, 1, *\}$ be a *restriction*, and $\boldsymbol{\rho}$ be a *random restriction*, i.e., a distribution of restrictions. We say that $\boldsymbol{\rho}$ is p -*regular* if $\Pr[\boldsymbol{\rho}(i) = *] = p$ and $\Pr[\boldsymbol{\rho}(i) = 0] = \Pr[\boldsymbol{\rho}(i) = 1] = (1 - p)/2$ for every $i \in [N]$. In addition, $\boldsymbol{\rho}$ is k -*wise independent* if any k coordinates of $\boldsymbol{\rho}$ are independent.

Lemma 28 (cf. [IMZ12, Vad12]). *There exist q -regular k -wise independent random restrictions $\boldsymbol{\rho}$ distributed over $\rho: [N] \rightarrow \{0, 1, *\}$ samplable with $O(k \log(N) \log(1/q))$ bits. Furthermore, each output coordinate of the random restriction can be computed in time polynomial in the number of random bits.*

As a consequence, we get p -regular k -wise independent random restrictions where each restriction in the support has bounded Kt complexity. In order to define the Kt complexity of a restriction $\rho: [N] \rightarrow \{0, 1, *\}$, we view it as a $2N$ -bit string $\text{encoding}(\rho)$ where each symbol in $\{0, 1, *\}$ is encoded by an element in $\{0, 1\}^2$. We abuse notation and write $\text{Kt}(\rho)$ to denote $\text{Kt}(\text{encoding}(\rho))$.

Proposition 29. *There is a distribution $\mathcal{D}_{q,k}$ of q -regular k -wise independent restrictions such that each restriction $\rho: [N] \rightarrow \{0, 1, *\}$ in the support of $\mathcal{D}_{q,k}$ satisfies $\text{Kt}(\rho) = O(k \log(N) \log(1/q))$. Furthermore, this is witnessed by a pair (M, w_ρ) where the machine M does not depend on ρ .*

Proof. By Lemma 28, each output coordinate of ρ can be computed in time $\text{poly}(\ell)$ from a seed w_ρ of length $\ell = O(k \log(N) \log(1/q))$. Therefore, the binary string describing ρ can be computed in time $O(N \cdot \text{poly}(\ell))$ from a string w_ρ with $\text{Kt}(w_\rho) = O(k \log(N) \log(1/q))$. It follows from Proposition 8 that $\text{Kt}(\rho) = O(k \log(N) \log(1/q))$. The furthermore part follows from the fact that the machine M is obtained from the generator provided by Lemma 28, i.e., in order to produce different restrictions one only needs to modify the input seeds, which are encoded in w_ρ . \square

Let $N = 2^n$. Given a function $F: \{0, 1\}^N \rightarrow \{0, 1\}$ and a restriction $\rho: [N] \rightarrow \{0, 1, *\}$, we let $F \upharpoonright_\rho$ be the function in $\{0, 1\}^{\rho^{-1}(*)} \rightarrow \{0, 1\}$ obtained in the natural way from F and ρ . In this section, we use $L(F)$ to denote the size (number of leaves) of the smallest U_2 -formula that computes a function F .

The next result allows us to shrink the size of a formula using a pseudorandom restriction. This restriction can be obtained by a composition of restrictions. This reduces the amount of randomness and the corresponding complexity of the restriction.

Lemma 30 (Shrinkage from pseudorandom restrictions ([HS17, Theorem 28]; cf. [IMZ12, KRT17])). *Let $F: \{0, 1\}^N \rightarrow \{0, 1\}$, $q = p^{1/r}$ for an integer $r \geq 1$, and $L(F) \cdot p^2 \geq 1$. Moreover, let $\mathcal{R}_{p,k}^r$ be a distribution obtained by the composition of r independent q -regular k -wise independent random restrictions supported over $[N] \rightarrow \{0, 1, *\}$, where $k = q^{-2}$. Finally, assume that $q \leq 10^{-3}$. Then,¹⁰*

$$\mathbb{E}_{\rho \in \mathcal{R}_{p,k}^r} [L(F \upharpoonright_\rho)] \leq c^r p^2 L(F),$$

where $c \geq 1$ is an absolute constant.

Proposition 31. *There is a (p -regular k -wise independent) distribution $\mathcal{R}_{p,k}^r$ obtained by the composition of r independent q -regular k -wise independent random restrictions supported over $[N] \rightarrow \{0, 1, *\}$, where $k = q^{-2}$ and $q = p^{1/r}$, such that each restriction $\rho: [N] \rightarrow \{0, 1, *\}$ in the support of $\mathcal{R}_{p,k}^r$ satisfies $\text{Kt}(\rho) = O(rk \log(N) \log(1/q))$.*

Proof. We use the distribution $\mathcal{D}_{q,k}$ of restrictions provided by Proposition 29. A restriction ρ in the support of $\mathcal{R}_{p,k}^r$ is therefore obtained through the composition of r restrictions ρ_1, \dots, ρ_r in the support of $\mathcal{D}_{q,k}$. For each $i \in [r]$, $\text{Kt}(\rho_i) = O(k \log(N) \log(1/q))$. Moreover, each Kt upper bound is witnessed by a pair (M, w_i) , where M can be taken to be the same machine for all $i \in [r]$. It is not hard to see that for the string $w = 1^{|w_1|} 0 w_1 1^{|w_2|} 0 w_2 \dots 1^{|w_r|} 0 w_r$ there is a machine M' satisfying $|\langle M' \rangle| \leq |\langle M \rangle| + O(1)$ and running in time $t_{M'}(w) \leq r \cdot \max_i t_M(w_i) + \text{poly}(rN)$ such that the pair (M', w) witnesses that $\text{Kt}(\rho) = O(rk \log(N) \log(1/q))$. \square

We will also need the following simple proposition, which holds even with respect to Kolmogorov complexity instead of Kt complexity.

¹⁰The assumption that $q \leq 10^{-3}$ does not appear in [HS17, Theorem 28]. The proof sketch appearing there does not seem to address the cases where $p^\Gamma L(\psi) < 1$ in their analyses of formula shrinkage in Lemma 27 and Theorem 28. This can be easily fixed using appropriate expressions of the form $1 + p^2 L(\psi)$. Lemma 27 is only affected by a constant factor. Then, proceeding by induction as in the proof of their Theorem 28 but also addressing this possibility, one gets instead an upper bound of the form $1 + cq^\Gamma (1 + cq^\Gamma (\dots))$, which translates to $1 + (cq^\Gamma) + (cq^\Gamma)^2 + \dots + (cq^\Gamma)^{r-1} + (cq^\Gamma)^r L(f)$. This can still be upper bounded by $c^r p^2 L(F)$ (for a different universal constant c as in the statement of Lemma 30) using that q is sufficiently small and therefore $cq^\Gamma \leq 1/2$ (note that $\Gamma = 2$ and $c \leq 500$ in [HS17]).

Proposition 32. *Let $S \subseteq [N]$ be a set of size at least two. There exists a function $h: S \rightarrow \{0, 1\}$ such that for every string $w \in \{0, 1\}^N$, if w agrees with h over S then $\text{Kt}(w) \geq |S| - 5 \log |S|$.*

Proof. It is easy to encode a pair (M, a) (as in Definition 6) satisfying $|\langle M \rangle| + |a| < |S| - 5 \log |S|$ by a binary string of length at most $2 \log |S| + 2 + |\langle M \rangle| + |a| < |S|$. Since each pair (M, a) outputs at most one binary string of length N , it follows by a counting argument that for some choice of $h: S \rightarrow \{0, 1\}$, no string w of length N that agrees with h over S has $\text{Kt}(w) < |S| - 5 \log |S|$. \square

The next lemma describes the high-level strategy of the lower bound proof.

Lemma 33 (Adaptation of Lemma 27 from [HS17]). *There exists a constant $a \geq 1$ such that the following holds. Let $\rho: [N] \rightarrow \{0, 1, *\}$ be a restriction, $V = \rho^{-1}(*)$, and let $F: \{0, 1\}^N \rightarrow \{0, 1\}$ be a function such that $L(F \upharpoonright_\rho) \leq M$. If*

$$\text{Kt}(\rho) + a \cdot n \leq s_1(n) \quad \text{and} \quad (|V| - M) - 5 \log(|V| - M) \geq s_2(n) \quad \text{and} \quad |V| \geq M + a,$$

then F does not compute $\text{Gap-MKtP}[s_1(n), s_2(n)]$, where $n = \log N$.

Proof. Under these assumptions, we define a positive instance $w^y \in \mathcal{YES}_N$ and a negative instance $w^n \in \mathcal{NO}_N$ such that $F(w^y) = F(w^n)$.

– $w^y \in \{0, 1\}^N$ is obtained from ρ by additionally setting each $*$ -coordinate of this restriction to 0. Note that, given the $2N$ -bit binary string encoding ρ , w^y can be computed in time polynomial in N . It follows from Proposition 8 that $\text{Kt}(w^y) \leq \text{Kt}(\rho) + a \cdot n$, for some universal constant $a \geq 1$. Since this bound is at most $s_1(n)$, we get that $w^y \in \mathcal{YES}_N$.

– $w^n \in \{0, 1\}^N$ is defined as follows. Since $L(F \upharpoonright_\rho) \leq M$, $F \upharpoonright_\rho$ depends on at most M input coordinates (indexed by elements in V). Let $W \subseteq V \subseteq [N]$ be this set of coordinates. Moreover, let $S = V \setminus W$. The string $w^y \in \{0, 1\}^N$ is obtained from ρ by additionally setting each $*$ -coordinate of this restriction in W to 0, and then setting each remaining $*$ -coordinate in S to agree with the function $h: S \rightarrow \{0, 1\}$ provided by Proposition 32. Since $|S| \geq |V| - M$ and the real-valued function $\phi(x) = x - 5 \log x$ is non-decreasing if $x \geq a$ for a large enough constant a , our assumptions and Proposition 32 imply that $\text{Kt}(w^n) \geq s_2(n)$. Consequently, $w^n \in \mathcal{NO}_N$.

Using that F restricted to ρ depends only on variables from $W \subseteq \rho^{-1}(*)$, and that the strings w^y and w^n agree over coordinates in $\rho^{-1}(\{0, 1\}) \cup W$, it follows that $F(w^y) = F(w^n)$. Since w^y is a positive instance while w^n is a negative instance, F does not compute $\text{Gap-MKtP}[s_1(n), s_2(n)]$. \square

We are now ready to set parameters in order to complete the proof of Theorem 3. For a sufficiently large constant $C' \geq 1$, let

$$n \stackrel{\text{def}}{=} \log N, \quad p \stackrel{\text{def}}{=} N^{-1+\alpha/2}, \quad r \stackrel{\text{def}}{=} n/C', \quad q \stackrel{\text{def}}{=} p^{1/r}, \quad k \stackrel{\text{def}}{=} q^{-2},$$

and assume that N is sufficiently large. Note that, under this choice of parameters, $q = 2^{C'(-1+\alpha/2)} = \Omega(1)$ and $q \leq 10^{-3}$.

Proposition 34 (Concentration Bound for $|\rho^{-1}(*)|$). *For $\rho \sim \mathcal{R}_{p,k}^r$ with parameters as above, we have $\Pr[|\rho^{-1}(*)| \geq pN/2] \geq 1/2$.*

Proof. Note that ρ is p -regular and pairwise independent (i.e., $k \geq 2$ for our choice of parameters). The result then follows from Chebyshev's inequality using mean $\mu = pN$, variance $\sigma^2 = Np(1-p)$, and the value of p . \square

Using Proposition 31, we can sample a random restriction $\rho \in_R \mathcal{R}_{p,k}^r$ as described in the statement of Lemma 30 such that each $\rho: [N] \rightarrow \{0, 1, *\}$ in the support of $\mathcal{R}_{p,k}^r$ satisfies

$$\text{Kt}(\rho) = O(rk \log(N) \log(1/q)) = O((n/C')q^{-2}n \log(1/q)) \leq (C/2)n^2,$$

if C is a sufficiently large constant.

Toward a contradiction, let $F: \{0, 1\}^N \rightarrow \{0, 1\}$ be a formula of size $L(F)$ that supposedly computes $\text{Gap-MKtP}[Cn^2, 2^{(\alpha/2)n-2}]$, where $p^2 L(F) = 1$ (note that $L(F) = N^{2-\alpha}$), and let

$$M \stackrel{\text{def}}{=} 10 \cdot c^r p^2 L(F) = 10 \cdot c^r \leq 2^{(\alpha/4)n},$$

for a constant $c \geq 1$ as in Lemma 30, and using that $C' = C'(\alpha)$ is large enough in the definition of r .

Invoking Lemma 30 and Markov's inequality, Proposition 34, and a union bound, there is a fixed restriction $\rho: [N] \rightarrow \{0, 1, *\}$ for which the following holds:

- For $V \stackrel{\text{def}}{=} \rho^{-1}(*)$, we have $|V| \geq pN/2 = 2^{(\alpha/2)n}/2$;
- $\text{Kt}(\rho) \leq (C/2)n^2$.
- $L(F \upharpoonright_\rho) \leq M \leq 2^{(\alpha/4)n}$.

Using these parameters in the statement of Lemma 33, it is easy to check that its hypotheses are satisfied given our choices of $s_1(n) = Cn^2$ and $s_2(n) = 2^{(\alpha/2)n-2}$. This is a contradiction to our assumption that F computes Gap-MKtP for these parameters, which completes the proof.

A.2 MKtP – Stronger lower bounds for large parameters

The goal of this section is to prove Theorem 2. First, we need a definition. We say that a generator $G: \{0, 1\}^r \rightarrow \{0, 1\}^N$ δ -fools a function $f: \{0, 1\}^N \rightarrow \{0, 1\}$ if

$$\left| \Pr_{x \in_R \{0,1\}^N} [f(x) = 1] - \Pr_{y \in_R \{0,1\}^r} [f(G(y)) = 1] \right| \leq \delta.$$

Similarly, G δ -fools a class of functions \mathcal{F} if G δ -fools every function $f \in \mathcal{F}$. The parameter r is called the *seed-length* of G . We say that G is *explicit* if it can be uniformly computed in time $\text{poly}(N, 1/\delta)$.

Theorem 35 ([IMZ12]). *Let $c > 0$ be an arbitrary constant. The following hold:*

1. *There is an explicit generator $G^{U_2}: \{0, 1\}^r \rightarrow \{0, 1\}^N$ using a seed of length $r = s^{1/3+o(1)}$ that s^{-c} -fools the class $U_2\text{-Formula}[s(N)]$ of formulas on N input variables.*
2. *There is an explicit generator $G^{B_2}: \{0, 1\}^r \rightarrow \{0, 1\}^N$ using a seed of length $r = s^{1/2+o(1)}$ that s^{-c} -fools the class $B_2\text{-Formula}[s(N)]$ of formulas on N input variables.*
2. *There is an explicit generator $G^{BP}: \{0, 1\}^r \rightarrow \{0, 1\}^N$ using a seed of length $r = s^{1/2+o(1)}$ that s^{-c} -fools the class $\text{BP}[s(N)]$ of branching programs on N input variables.*

We now prove Theorem 2 (Part 1). The other cases are similar. We instantiate G^{U_2} with $s(N) = N^{3-\varepsilon}$ and $c = 1$. Then $G^{U_2}: \{0, 1\}^{N^{1-\delta'}} \rightarrow \{0, 1\}^N$ for some $\delta' = \delta'(\varepsilon) > 0$.

Proposition 36. For every string $w \in \{0,1\}^{N^{1-\delta'}}$, let $G^{U_2}(w) \in \{0,1\}^N$ be the N -bit output of G^{U_2} on w . Then

$$\text{Kt}(G^{U_2}(w)) \leq 2^{(1-\delta'/2)n}$$

for every large enough $n = \log N$.

Proof. This follows from Proposition 8 using that G^{U_2} is explicit and therefore runs in time $\text{poly}(N)$ under our choice of parameters. \square

As a consequence of Proposition 36, every output of G^{U_2} is always an N -bit string of Kt complexity at most $2^{(1-\delta)n}$, for a fixed $\delta > 0$. On the other hand, it is well-known that a random N -bit string (where $N = 2^n$) has Kolmogorov complexity (and thus Kt complexity) at least 2^{n-1} with high probability. It follows that $\text{Gap-MKtP}[2^{(1-\delta)n}, 2^{n-1}] \notin U_2\text{-Formula}[N^{3-\varepsilon}]$, since otherwise this would violate the security of the generator G^{U_2} against formulas of this type and size.

A.3 MCSP – A similar near-quadratic lower bound against U_2 -formulas

In this Section, we sketch the proof of Theorem 5, which is the analogue of Theorem 3 in the context of MCSP. More precisely, we explain why the argument carries over when we measure the complexity of a string by circuit size instead of via Kt complexity, modulo small changes to the involved parameters.

As explained in Section A.1, the crucial idea in the proof of Theorem 3 is that a pseudorandom restriction simplifies a U_2 -formula of bounded size. For technical reasons, we employ a composition of restrictions of small complexity, so that the overall complexity of the combined restriction is bounded. This allows us to trivialize any small formula F using a fixed restriction ρ of bounded complexity, where $|\rho^{-1}(*)|$ is sufficiently large compared to other relevant parameters of the argument. Then, Lemma 33 employs a counting argument (via Proposition 32) to extend this restriction to a positive instance w^y and to a negative instance w^n such that $F(w^y) = F(w^n)$. This can be used to show that no small formula correctly computes Gap-MKtP for our choice of parameters.

In order to establish Theorem 5, we make two observations. Firstly, Lemma 28 already gives individual restrictions of low *circuit complexity* instead of low Kt complexity. Secondly, the *counting argument* used to extend ρ to a negative instance w^n works for most complexity measures including circuit size, Kolmogorov complexity, etc.

Using these two observations, the proof goes through under minor adjustments of the relevant parameters. We remark that one obtains a lower bound for $\text{Gap-MCSP}[n^d, 2^{(\alpha/2-o(1))n}]$ instead of $\text{Gap-MCSP}[Cn^2, 2^{(\alpha/2)n-2}]$ because of a polynomial circuit complexity overhead in the argument, which is not present in the case of Kt complexity since there one takes the logarithmic of the running time when measuring complexity, and because the circuit complexity (measured by number of gates) of a random string can be slightly smaller than its Kt complexity.

B Hardness Magnification and Proof Complexity

The initial instance of hardness magnification from [OS18] says that if an average-case version of MCSP (with inputs being truth tables of Boolean functions) is worst-case hard for formulas of super-linear size, then its succinct version (with inputs being lists of input-output tuples representing partial Boolean functions) is hard for NC^1 (cf. [OS18, Theorem 1]).

Hardness magnification for MCSP thus attacks strong circuit lower bounds by 1. employing the natural proofs barrier which states a conditional hardness of MCSP, and 2. exploiting the difference between feasible (succinct) and infeasible (uncompressed) formulations of a meta-computational problem like MCSP.

This strategy has a history in proof complexity. The work of Razborov [Raz94, Raz95] and Krajíček [Kra04] formulated the natural proofs barrier as a conditional proof complexity lower bound expressing hardness of tautologies encoding circuit lower bounds. This idea was further developed in the theory of proof complexity generators [Kra01, ABSRW04]. It has led, in particular, to Razborov’s conjecture [Raz15] about hardness of Nisan-Wigderson generators for strong proof systems. Razborov’s conjecture is designed to imply hardness of circuit lower bounds formalized in a way so that the whole truth table of the hard function is hardwired into the formula.

The realization that a feasible formulation of circuit lower bounds should be much harder than the infeasible truth table formulas inspired the result about unprovability of circuit lower bounds in theories of bounded arithmetic such as VNC^1 , cf. [Pic15], and the proposal [Pic14, 0.1 Circuit lower bounds and Complexity-Theoretic tautologies] to study exponentially harder lb formulas. Once the definitions are given, it is for example clear that polynomial-size proofs of the lb formulas transform into almost linear-size proofs of the truth table formulas. Another instance of this phenomena says that:

If the truth table formulas encoding a polynomial circuit lower bound require superlinear-size proofs in AC^0 -Frege systems, then lb formulas encoding the same polynomial circuit lower bound require (NC^1) -Frege proofs of super-polynomial size (implicit in the proof of [MP17, Proposition 4.14]).

Since AC^0 -Frege lower bounds are known, this suggests a way for attacking Frege lower bounds. ([OS18] established analogous results in circuit complexity, where it might be easier to prove lower bounds. However, their version of the MCSP problem refers to the average-case complexity of truth-tables, which seems harder to analyse. We refer to [OS18] for further discussion.)

The lb formulas result from the feasible witnessing of circuit lower bounds. In [MP17], the witnessing was provided by a theorem of Lipton and Young [LY94] establishing the existence of anti-checkers, described in Section 1.2. This allows to express the hardness of f without using its whole truth table. The present paper extends the idea of anti-checkers into the context of hardness magnification in circuit complexity for the standard worst-case formulation of MCSP.