

Circuit Lower Bounds from NP-hardness of MCSP under Turing Reductions

Michael Saks

Rutgers University, United States of America
saks@math.rutgers.edu

Rahul Santhanam

University of Oxford, United Kingdom
rahul.santhanam@cs.ox.ac.uk

Abstract

The fundamental Minimum Circuit Size Problem is a well-known example of a problem that is neither known to be in P nor known to be NP-hard. Kabanets and Cai [18] showed that if MCSP is NP-hard under “natural” m-reductions, superpolynomial circuit lower bounds for exponential time would follow. This has triggered a long line of work on understanding the power of reductions to MCSP.

Nothing was known so far about consequences of NP-hardness of MCSP under general Turing reductions. In this work, we consider two structured kinds of Turing reductions: *parametric honest* reductions and *natural* reductions. The latter generalize the natural reductions of Kabanets and Cai to the case of Turing-reductions. We show that NP-hardness of MCSP under these kinds of Turing-reductions imply superpolynomial circuit lower bounds for exponential time.

2012 ACM Subject Classification Theory of computation → Computational complexity and cryptography; Theory of computation → Circuit complexity; Theory of computation → Complexity classes

Keywords and phrases Minimum Circuit Size Problem, Turing reductions, circuit lower bounds

Digital Object Identifier 10.4230/LIPIcs.CCC.2020.26

Funding *Rahul Santhanam*: This work was supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2014)/ERC Grant Agreement No. 615075.

1 Introduction

The Minimum Circuit Size Problem (MCSP) is a fundamental problem in computational complexity that has been studied in various forms since the 1950s [24, 1]. Given the truth table of a Boolean function f and a parameter s , the question is whether f has Boolean circuits of size at most s . It is easy to see that MCSP is in NP: we can simply guess a circuit C of size at most s , and verify that for each input of f that C computes f correctly. Since we are given the truth table of f explicitly, this verification can be done in polynomial time.

MCSP is a rare example of a natural problem in NP for which we neither know that the problem is in P nor that it is NP-complete. There is some evidence that the problem is not in P. The celebrated results on “natural proofs” by Razborov and Rudich [23] can be interpreted as saying that MCSP is not in P if one-way functions exist [18, 2].

Regarding the question of NP-completeness, however, the evidence is more mixed. It is reported that Levin delayed publishing his seminal results on NP-completeness because he was hoping to show that MCSP is NP-complete [8]. More than 50 years later, we still have no idea.

It has been known for a while that the analogue of MCSP for DNFs, where we ask whether the function corresponding to a given truth table has small DNFs, is NP-complete [20, 10, 5]. Intuitively it would appear that checking if a given function has small circuits is at least as



© Michael Saks and Rahul Santhanam;
licensed under Creative Commons License CC-BY
35th Computational Complexity Conference (CCC 2020).

Editor: Shubhangi Saraf; Article No. 26; pp. 26:1–26:13



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 hard a problem as checking if a given function has small DNFs. But it does not seem easy to
 46 make this intuition precise in terms of giving a reduction from the latter to the former¹.

47 Looking at the proof that the analogue of MCSP for DNFs is NP-hard, one notices that
 48 an essential component of the reduction is an explicit function, namely Parity, that is hard
 49 for DNFs. This leads one to wonder whether the difficulty of coming up with an NP-hardness
 50 reduction for MCSP is related to the difficulty of proving Boolean circuit lower bounds for
 51 explicit functions.

52 Kabanets and Cai show that the answer is yes, in an influential paper [18] which led to a
 53 resurgence of interest in MCSP among complexity theorists. They define a notion of "natural
 54 reduction" to MCSP. A natural reduction is a many-one reduction in which the parameter
 55 as well as the size of the output depend only on the size of the input. Kabanets and Cai
 56 observe that standard reductions showing NP-hardness of various well-studied problems are
 57 natural in this sense. They then show that any natural reduction of SAT to MCSP implies
 58 that exponential time requires super-polynomial size Boolean circuits. Note that this does
 59 not give evidence *against* NP-hardness. However, given the well-known difficulty of proving
 60 circuit lower bounds for exponential time, it does indicate that *arguing* for NP-hardness of
 61 MCSP via natural reductions is likely to be difficult.

62 The work of Kabanets and Cai has led to a long line of results on reductions to MCSP
 63 and their implications [2, 3, 4, 7, 21, 17, 16, 6, 22, 15, 14]. Much of this work has focused on
 64 reductions implementable within low-depth complexity classes [7, 21, 6, 22] or on many-one
 65 and truth-table reductions [7, 21, 17, 16].

66 In this work, our focus is on polynomial-time Turing reductions to MCSP, for which very
 67 little is known. On the one hand, we are motivated by the question of NP-hardness of MCSP,
 68 and whether more powerful reductions help in this regard. Allender and Das [3] showed
 69 that every problem in Statistical Zero Knowledge reduces to MCSP via probabilistic Turing
 70 reductions, so at least for some problems that are believed to be hard, Turing reductions
 71 have been found to be useful in reducing to MCSP.

72 On the other hand, from the perspective of complexity theorists who care about circuit
 73 lower bounds, we find the connections between reductions to MCSP and circuit lower bounds
 74 intriguing, and would like to understand the extent to which these connections hold.

75 1.1 Our Results

76 We consider two kinds of structured Turing reductions - *parametric honest* reductions, which
 77 were defined in the context of many-one reductions by Hitchcock and Pavan [17], and *natural*
 78 reductions, which generalize the notion defined by Kabanets and Cai [18]. Parametric honest
 79 reductions are Turing reductions where there is a polynomial lower bound on the parameter
 80 of each query, as a function of input size. Natural reductions are Turing reductions where
 81 the parameter of any query made on an input only depends on the size of the input.

82 Our first main result shows that NP-completeness of MCSP under parametric honest
 83 Turing reductions implies super-polynomial circuit lower bounds for E.

84 ► **Theorem 1.** *If MCSP is NP-hard under parametric honest Turing reductions, then $E \not\subseteq$*
 85 *SIZE(poly).*

86 We remark that every language in P reduces to MCSP under parametric honest Turing
 87 reductions, so we crucially need to exploit the assumption that we reduce from hard languages

¹ Similar issues come up when trying to prove hardness of properly learning Boolean circuits.

88 in NP in order to derive the conclusion of Theorem 1.

89 Our second main result shows that NP-completeness of MCSP under natural Turing
90 reductions implies super-polynomial circuit lower bounds for E.

91 ► **Theorem 2.** *If MCSP is NP-hard under natural Turing reductions, then $E \not\subseteq \text{SIZE}(\text{poly})$.*

92 The proof of Theorem 2 builds upon the ideas of Theorem 1, but is significantly more
93 involved technically, and requires several new ideas.

94 We next describe the main ideas behind our proofs.

95 1.2 Main Ideas

96 Our starting point is the idea of Kabanets and Cai [18] for deriving circuit lower bounds
97 from structured many-one reductions. Suppose we could efficiently generate NO instances of
98 any length for some language L that reduces to MCSP via an m -reduction. We could hope
99 to efficiently generate truth tables of hard functions as follows: we generate a NO instance
100 x_n of L of length n , and then apply the m -reduction f from L to MCSP to generate a truth
101 table y_n and a parameter s_n . Since x_n is a NO instance of L and f is an m -reduction, we
102 are guaranteed that y does not have circuits of size s .

103 If s is large, say at least n^ϵ for some constant ϵ , then we are done, since we have efficiently
104 generated a truth table y of a Boolean function that does not have sub-exponential size
105 Boolean circuits as a function of its input length. But a priori we do not have control over the
106 parameter s associated with y . Imagine for example a reduction that magically knows that
107 each x_n is a NO instance and therefore produces s_n so small that it gives us no non-trivial
108 information about circuit lower bounds.

109 Kabanets and Cai get around this problem by considering natural reductions, where
110 the parameter s_n depends only on the input length n . Now there are two cases: either
111 all but finitely many input lengths yield small parameter $s_n \leq \log(n)^{\log \log(n)}$, or infinitely
112 many input lengths yield parameter $s_n > \log(n)^{\log \log(n)}$. In the first case, Kabanets and Cai
113 use a standard indirect diagonalization argument to argue that circuit lower bounds follow.
114 In the second case, we directly get hard truth tables for infinitely many n by composing
115 the reduction with the efficient generation of NO instances. Thus, in either case, we get
116 $E \not\subseteq \text{SIZE}(\text{poly})$.

117 Note that in the case where infinitely many input lengths yield hard parameters, the
118 Kabanets-Cai argument does not actually exploit the hardness of the language L from which
119 we are reducing. The argument works equally well starting from a language in polynomial
120 time. When we deal with Turing reductions rather than m -reductions, this feature of their
121 argument is an issue, since any language L in P can be decided trivially by Turing reductions
122 that ignore the answers to any queries they ask and simulate the polynomial-time algorithm
123 for L directly to arrive at an answer.

124 Let us consider first the case of parametric honest reductions. Parametric honesty means
125 that we don't need to worry about the parameter being small. Suppose that SAT reduces to
126 MCSP via a parametric honest Turing reduction implemented by an oracle machine M . Our
127 first idea is simple: we define a simulation A of SAT which simply runs M and answers all
128 queries of M by 'yes'. This algorithm A clearly runs in polynomial time. If it is correct on
129 all but finitely many instances, we have that $\text{NP} = \text{P}$ and we can then derive circuit lower
130 bounds using a standard indirect diagonalization technique.

131 If A is wrong on infinitely many instances, however, it is not obvious how to use this.
132 Here we exploit the main idea of Gutfreund, Shaltiel and Ta-Shma [12], who show how any
133 efficient algorithm A that fails to solve SAT can be used to efficiently produce, for infinitely

134 many n , a small set of instances S of sizes polynomially related to n , such that A fails on
 135 some instance in S . This is very useful for us, as an instance on which A fails must ask a
 136 query to M that is answered incorrectly. But since we always answer queries by ‘yes’ in A , a
 137 query is only answered incorrectly if the true answer is “no”, i.e., if the string y that is queried
 138 is the truth table of a hard Boolean function. We don’t know which of the queries this is, but
 139 we do know that there is a *first* wrongly answered query, and that by the parametric honesty
 140 condition, this query has large parameter. Hence, by simply concatenating the queries asked
 141 by A on the instances in S , we obtain hard truth tables of size $\text{poly}(n)$ for infinitely many n ,
 142 and this enables us to conclude that E requires super-polynomial size circuits.

143 This is the argument for the proof of Theorem 1. The assumption of parametric honesty
 144 is an important part of this argument. For the proof of Theorem 2, where we consider natural
 145 reductions, we need to work much harder and use several new ideas.

146 We would like to adopt the following strategy. Suppose there is a natural Turing-reduction
 147 from SAT to MCSP. Choose a “threshold” parameter s and try to solve SAT by running
 148 the reduction and answering all queries with parameter less than the threshold correctly by
 149 doing brute force search, and all queries with parameter greater than the threshold by ‘yes’.

150 Either this simulation succeeds, or it fails. If it succeeds, and the threshold parameter s
 151 is small, NP is easy, and using the standard indirect diagonalization argument, we get that E
 152 requires large circuits. If it fails, we would like to argue that we can efficiently find instances
 153 on which the simulation fails. Such an instance must ask queries with parameters larger than
 154 the threshold that are wrongly answered ‘yes’, so the concatenation of such queries must be
 155 a hard truth table.

156 To efficiently find instances on which the simulation fails, we can again to try to use the
 157 strategy of [12]. Unfortunately, there is a catch. The time taken by the algorithm of [12] to
 158 find hard instances is at least the time taken for the simulation, which is exponential in s .
 159 It might be that the parameter for queries on these instances is just slightly larger than s .
 160 Thus we will take time exponential in s to find truth tables requiring circuits only slightly
 161 larger than s , which won’t give us sufficiently strong circuit lower bounds for E .

162 To remedy this issue, we consider two different simulations, M and M' . M is guaranteed
 163 to be correct and proceeds by just evaluating queries using brute force search. If M is
 164 relatively efficient, we are in good shape, as we can use indirect diagonalization to get the
 165 consequence we want. If M is not always efficient, we get an infinite sequence of input
 166 lengths on which large parameter queries are made. Here “large” means at least s , where s is
 167 super-polylogarithmic in n , but still small enough so that the indirect diagonalization helps
 168 us get circuit lower bounds in case M succeeds.

169 In fact, by using the paddability of SAT in our argument, we get not just an infinite
 170 sequence of input lengths yielding large parameters, but a sequence of long intervals, within
 171 which all input lengths yield large parameters. This will be important later, as the [12]
 172 algorithm does not produce hard instances at a given input length, but rather at one of two
 173 input lengths which are polynomially related.

174 The simulation M' will use a smaller parameter threshold s' , chosen so that s' is
 175 superpolylogarithmic in n , but also so that s is superpolynomial in s' . The simulation M'
 176 implements our initial idea: we answer queries with parameter at most s' by brute force
 177 search, and other queries by ‘yes’.

178 Now either this simulation is correct on all but finitely many input lengths, or it fails on
 179 infinitely many. If it is correct on all but finitely many input lengths, we can use indirect
 180 diagonalization as before, but it is not clear how to use the failure on infinitely many, without
 181 more structural information on where the failures occur.

182 So we argue instead as follows, using a notion of “robust simulation” introduced in [11].
183 If the simulation is correct, robustly often, then we can carry the indirect diagonalization
184 through. Otherwise, the simulation fails on at least one input length in every large enough
185 sequence of input lengths. Using the large stretches of input lengths with large parameters
186 from our first simulation, we get that there are infinitely many long sequences of input
187 lengths where the simulation M' fails and the corresponding parameters of queries are large.

188 Now we use the strategy of [12]. We use simulation M' itself to find a small set of inputs
189 on which M' fails, and on which the parameter of correctly produced queries is at least s .
190 M' will take time only exponential in s' , and since s is super-polynomial in s' , the time
191 taken to find this small set of inputs will be sub-exponential in s . By concatenating the
192 queries made on this small set of inputs, we get a hard truth table, and hence can arrive at
193 our desired conclusion.

194 There are numerous technical details which we have ignored in the above discussion, but
195 hopefully this description helps in understanding the proof of Theorem 2.

196 1.3 Related Work

197 As mentioned earlier in the Introduction, there have been several works in recent years
198 studying reductions to MCSP and their consequences. We briefly survey this work in the
199 context of our results. We focus on work that either shows interesting consequences of
200 hardness for MCSP under certain types of reductions, or unconditionally rules out hardness,
201 as this is closest in spirit to our work here.

202 Kabanets and Cai [18] defined natural m-reductions, and showed that NP-hardness of
203 MCSP under natural m-reductions implies super-polynomial circuit lower bounds for E.
204 Murray and Williams [21] showed unconditionally that MCSP is unconditionally not NP-
205 hard under very efficient local reductions where each output bit only depends on a limited
206 number of input bits. In fact, their technique even rules out P-hardness of MCSP under
207 local reductions. They also show that NP-hardness of MCSP under m-reductions (without
208 the “natural” constraint of [18]) implies $\text{EXP} \neq \text{ZPP}$. Hitchcock and Pavan showed that
209 NP-hardness of MCSP under general truth-table reductions implies $\text{EXP} \neq \text{ZPP}$. Allender,
210 Holden and Kabanets [7] consider the question of hardness for various oracle versions of
211 MCSP and show some unlikely consequences of hardness under efficient reductions when the
212 oracle is powerful.

213 The work that is closest to ours is that of Hirahara and Watanabe [16], who also consider
214 Turing-reductions to MCSP. They consider a kind of reduction they call “oracle-independent”,
215 and show that no language outside of P reduces to MCSP using an oracle-independent
216 reduction. However, there are examples of useful reductions known that are not oracle-
217 independent [6]². It is also shown in [16] that Turing-hardness of approximating the minimum
218 circuit size implies $\text{EXP} \neq \text{ZPP}$. We note that the hardness of approximation factors required
219 for this result are large - the assumption is that it is hard to approximate the logarithm of
220 the circuit size to within any constant factor. Moreover, [16] do not derive a circuit lower
221 bound from this assumption, but the weaker statement that $\text{EXP} \neq \text{ZPP}$.

² On the other hand, we are not aware of any hardness results for MCSP using reductions that are not natural.

222 **2 Preliminaries**223 **2.1 Basic Notions**

224 We refer to the book by Arora and Barak [9] for definitions of standard complexity classes.

225 Given a circuit class \mathfrak{C} , we use $\mathfrak{C}[s(n)]$ to refer to the class of functions computed by
226 \mathfrak{C} -circuits of size $s(n)$.

227 Given a language $L \subseteq \{0, 1\}^*$, $\text{co-}L$ denotes the complement of L . Given language L and
228 $n \in \mathbb{N}$, L_n is used to refer to $L \cap \{0, 1\}^n$.

229 Given a set $S \subseteq \mathbb{N}$ and languages L and L' , we say L agrees with L' on S if for all $n \in S$,
230 $L_n = L'_n$.

231 A time-constructible function $T : \mathbb{N} \rightarrow \mathbb{N}$ is a function such that there is a Turing machine
232 transducer M , which for each n , on input 1^n halts with output $T(n)$ within $O(T(n))$ steps.

233 We need a generalization of the notion of robustly-often simulation from [11]. Given
234 $g : \mathbb{N} \rightarrow \mathbb{N}$, a set $S \subseteq \mathbb{N}$ is called g -robust if for each constant k , there is $m \in \mathbb{N}$ such that for
235 all n such that $m \leq n \leq g(m)^k$, $n \in S$. Note that any g -robust set is also $\text{poly}(g)$ -robust.

236 Given a function g , a language L and a complexity class \mathcal{C} , we say L is g -robustly often
237 in \mathcal{C} if there is $L' \in \mathcal{C}$ for which there is a g -robust set S such that L agrees with L' on S .
238 Given function g and complexity classes \mathcal{C} and \mathcal{C}' , we say that \mathcal{C} is g -robustly often in \mathcal{C}' if
239 for all languages $L \in \mathcal{C}$, L is g -robustly often in \mathcal{C}' . We use the term "robustly often" as a
240 synonym for g -robustly often with g the identity function.

241 The proposition below is a parameterized version of Theorem 12 in [11].

242 **► Proposition 3.** *Let T be a time-constructible function such that $T(n) \geq n$ for all n . If SAT*
243 *is $T(\text{poly}(n))$ -robustly often in $\text{DTIME}(T)$, then $\Sigma_2\text{P}$ is robustly often in $\text{DTIME}(T(\text{poly}(T(\text{poly}(n))))$).*

244 The following is a simple application of standard diagonalization [13].

245 **► Proposition 4.** *Let T be a time-constructible function such that $T = o(2^n)$. Then E is not*
246 *robustly often in $\text{DTIME}(T)$.*

247 **2.2 Resource-Bounded Kolmogorov Complexity**

248 We study various notions of resource-bounded Kolmogorov complexity, and associated decision
249 problems.

250 Fix a universal Turing machine U . The Kt complexity of a string is defined as follows:
251 $\text{Kt}(x) = \min\{|p| + \log(t) \mid U(p) \text{ halts and outputs } x \text{ within } t \text{ steps}\}$.

252 In the MKtP problem, the instance is a string x together with a parameter s , and the
253 question is whether $\text{Kt}(x) \leq s$.

254 It is known that MKtP is complete for E under polynomial-size truth-table reductions [2],
255 but completeness under uniform polynomial-time reductions is unknown.

256 In the MCSP problem, the instance is a string x together with a parameter s , and the
257 question is whether x is the truth table of a Boolean function with circuit complexity at
258 most s .

259 MCSP is easily seen to be in NP , but it is unknown whether MCSP is NP-hard . A
260 brute-force search strategy of trying all circuits C of size at most s and checking if any of
261 them computes the function with truth table x can easily be implemented to run in time
262 $\text{poly}(|x|)s^{O(s)}$.

2.3 Reductions

We will work with constrained notions of polynomial-time Turing reduction.

► **Definition 5.** *A polynomial-time Turing reduction from a language L to MCSP or MKtP is called parametric honest if there is a constant $\epsilon > 0$ such that for every $x \in \{0, 1\}^*$, every query made by the reduction on x has parameter bounded below by $|x|^\epsilon$.*

Note that parametric honesty is not implied by the standard notion of honesty for polynomial-time reductions, which only requires the output length to be bounded below by some polynomial in the input length. Implicit in the definition is that the lower bound on parameters of queries is only required to hold when previous queries are answered correctly; when queries are answered wrongly, "all bets are off".

► **Definition 6.** *A polynomial-time Turing reduction from a language L to MCSP or MKtP is called natural if for any input $x \in \{0, 1\}^*$, the parameter of any query made by the reduction on x depends only on $|x|$.*

This notion of naturalness generalizes the notion used by Kabanets and Cai for m-reductions to Turing reductions. In fact, it is slightly weaker than their notion when restricted to m-reductions, since it does not constrain the length of the output, merely the parameter. As with the definition of parametric honesty, the condition on parameters is only required to hold when previous queries are answered correctly.

Naturalness is incomparable to parametric honesty - parametric honesty allows the parameter to vary but restricts its range of variation, while naturalness allows the parameter to have a large range of variation but insists that it is the same for all queries made on any input of a given length.

3 Parametric Honest Reductions

Our first result unconditionally rules out hardness of MKtP for E under parametric honest Turing reductions. For many results on hardness of MCSP and consequences thereof, MKtP is a 'test case' where analogous results can be shown unconditionally and where the ideas are often simpler [8, 16].

► **Theorem 7.** *MKtP is not hard for E under parametric honest Turing reductions.*

Proof. Let L be a tally set that is in E but not in P. The existence of such a tally set L follows by a standard diagonalization argument.

Now suppose, for the purpose of contradiction, that MKtP is hard for E under parametric honest Turing reductions. This implies that there is an oracle Turing machine M running in time at most cn^k , where c and $k \geq 1$ are constants, such that M decides L correctly with oracle access to MKtP, and moreover there is a constant $\epsilon > 0$ such that each query made by M on an input of length n has parameter at least n^ϵ . We show how to use M to decide $L \in P$, contrary to the assumption on L .

Consider the following polynomial-time machine A that attempts to decide L . Given input 0^n , A runs the oracle machine M , answering each query made by M with 'yes'. When the simulation of M is complete, A accepts iff M accepts. Since M is polynomial-time, so is A . We show that A correctly decides L for large enough n .

Given input 0^n , M makes some sequence of oracle queries $q_1, q_2 \dots q_t$ during the simulation by A , where $0 \leq t \leq cn^k$. We show that for each i , q_i has Kt complexity $O(\log(n))$. Observe that q_i can be generated by the universal machine U implicit in the definition of Kt complexity

306 in time $\text{poly}(n)$ given descriptions of i , n and the oracle machine M : it simply needs to
 307 run M on 0^n answering the first $i - 1$ queries made by M with ‘yes’, and then output the
 308 i ’th query. The description length required is $O(\log(n))$ since $i \leq cn^k$, and moreover the
 309 time taken by U is polynomial in n , hence using the definition of Kt complexity, we have an
 310 $O(\log(n))$ upper bound on the Kt complexity of q_i . Let C be a constant such that the Kt
 311 complexity of each q_i is bounded above by $C \log(n)$.

312 By the assumption that the reduction is parametric honest, the parameter of q_i is at
 313 least n^ϵ , which for large enough n is greater than $C \log(n)$. Hence for large enough n and
 314 for each i , it follows that A ’s assumed answer for the oracle query q_i is correct for each i ,
 315 and hence that A outputs the same answer as M would given the true oracle MKtP. Since
 316 M is an oracle Turing machine correctly implementing a reduction from L to MKtP, A is a
 317 polynomial-time Turing machine correctly deciding L for large enough n , contradicting the
 318 assumption on L .

319

◀

320 We now re-state and prove Theorem 1.

321 ► **Theorem 8.** *If MCSP is NP-hard under parametric honest Turing reductions, then $\text{E} \not\subseteq$*
 322 *SIZE(poly).*

323 **Proof.** Suppose that MCSP is NP-hard under parametric honest Turing reductions. Let M
 324 be a polynomial-time oracle Turing reduction implementing a parametric honest reduction
 325 from SAT to MCSP, and let ϵ be a constant such that the parameter of any query made by
 326 M on any SAT instance of length x is bounded below by $|x|^\epsilon$.

327 Consider the following polynomial-time machine A which attempts to decide SAT. On an
 328 input x of SAT, A runs the oracle Turing machine M . Assume wlog that all queries asked by
 329 M have length a power of two - any queries for which this is not the case may be eliminated
 330 by assuming the answer is ‘no’. For every query asked by M with length a power of two,
 331 A assumes the answer is ‘yes’, and continues the simulation of M . When the simulation
 332 concludes, A accepts iff M accepts.

333 Clearly A runs in polynomial time. If A decides SAT correctly, then we have that $\text{NP} = \text{P}$.
 334 Hence the Polynomial Hierarchy collapses to P. But this implies that $\text{E} \not\subseteq \text{SIZE}(\text{poly})$; if the
 335 contrary were true, then E would collapse to the Polynomial Hierarchy using the Karp-Lipton
 336 theorem for E, and hence to P, contradicting the hierarchy theorem for deterministic time.

337 Suppose now that $\text{NP} \neq \text{P}$. In particular, this means that there are infinitely many
 338 instances on which A does not decide SAT correctly. If not, we could hardcode the finitely
 339 many instances on which A fails to compute SAT into a new polynomial-time machine A'
 340 which runs in polynomial time and solves SAT correctly on all instances, contradicting the
 341 assumption that $\text{NP} \neq \text{P}$.

342 Now, we can use the main result of Gutfreund, Shaltiel and Ta-Shma. They show that
 343 if $\text{NP} \neq \text{P}$, then for any polynomial-time algorithm A that fails to compute SAT, there is
 344 a polynomial-time Turing machine B and a constant d , such on input 1^n , B outputs three
 345 instances x_1, x_2, x_3 each of length n or length n^d such that for infinitely many n , A fails to
 346 decide SAT correctly on at least one of these three instances.

347 We show how to use B to argue again that $\text{E} \not\subseteq \text{SIZE}(\text{poly})$. Consider the three instances
 348 x_1, x_2, x_3 output by B on input 1^n . For each i , $1 \leq i \leq 3$, let Q_i be the concatenation of all
 349 queries made by M when M is simulated by A on input x_i . Let Y be the concatenation of
 350 Q_1, Q_2, Q_3 . Note that $|Y| = \text{poly}(n)$, since M is a polynomial-time oracle Turing machine.
 351 Also note that each query made by M on an input of length n has parameter at least n^ϵ ,
 352 by the parametric honesty of the reduction, as long as all previous queries were answered

353 correctly. Each such query was assumed by A to have a ‘yes’ answer. It could not be the
 354 case that all such answers were correct for all queries made by M on inputs x_1, x_2, x_3 - if
 355 they were, then A would correctly solve SAT on x_1, x_2, x_3 , using the fact that M implements
 356 a correct polynomial-time Turing reduction from SAT.

357 Thus we get that there is a first query made by M on one of x_1, x_2, x_3 that has parameter
 358 at least n^ϵ and is wrongly answered ‘yes’. Fix such a query y . Since M implements a Turing
 359 reduction from SAT to MCSP, y is the truth table of a Boolean function on $O(\log(n))$ bits
 360 that requires circuits of size n^ϵ . We have that y is a substring of Y , hence Y must also
 361 require circuits of size $\Omega(n^\epsilon)$. We are using here the simple fact that if y is a substring (with
 362 length a power of two) of the truth table Y of a Boolean function with circuits of size s , then
 363 y is the truth table of a Boolean function with circuits of size $O(s)$.

364 Hence, for infinitely many n , on input 1^n , we can compute the truth table Y of a Boolean
 365 function on $O(\log(n))$ bits such that the function requires circuits of size n^ϵ , i.e., of exponential
 366 size in the length of the input to the Boolean function. This implies $E \not\subseteq \text{SIZE}(\text{poly})$. ◀

367 4 Natural Reductions

368 We require a technical lemma allowing us to use indirect diagonalization based on robustly-
 369 often simulations.

370 ▶ **Lemma 9.** *Let $T : \mathbb{N} \rightarrow \mathbb{N}$ be a time-constructible function such that $T(\text{poly}(T(\text{poly}(n)))) =$
 371 $o(2^n)$. If SAT is $T(\text{poly}(n))$ -robustly often in time $T(n)$, then $E \not\subseteq \text{SIZE}(\text{poly})$.*

372 **Proof.** Suppose SAT is $T(\text{poly}(n))$ -robustly often in time $T(n)$. By Proposition 3, we have
 373 that $\Sigma_2\text{P}$ is robustly often in time $T(\text{poly}(T(\text{poly}(n))))$. Assume for the sake of contradiction
 374 that $E \subseteq \text{SIZE}(\text{poly})$. The Karp-Lipton theorem for E , credited to Meyer [19], implies that
 375 E collapses to $\Sigma_2\text{P}$. Note that if T is time-constructible, $T' = T(\text{poly}(T(\text{poly}(n))))$ is also
 376 time-constructible. Thus we have that E is robustly often in time T' for a time-constructible
 377 $T' = o(2^n)$, contradicting Proposition 4. ◀

378 We now re-state and prove Theorem 2. In the proof below, we use the fact that SAT
 379 is paddable - there is a polynomial-time computable *padding* function f such that for each
 380 $x \in \{0, 1\}^*$, and for each $m \geq |x|$, $|f(x, 1^m)| = m$ and $f(x, 1^m)$ is in SAT iff x is in SAT.
 381 Of course, this depends on the encoding of SAT, but it suffices for us that there is some
 382 encoding for which paddability is true in the form above, and all other standard properties
 383 of SAT continue to hold.

384 ▶ **Theorem 10.** *If MCSP is NP-hard under natural Turing-reductions, then $E \not\subseteq \text{SIZE}(\text{poly})$*

385 **Proof.** Suppose that MCSP is NP-hard under natural Turing reductions. Let M be an oracle
 386 Turing machine implementing a natural reduction from SAT to MCSP. By the definition of
 387 natural reduction, this means that there is a function $s : \mathbb{N} \rightarrow \mathbb{N}$ such that all queries of M
 388 on input of size n for SAT have parameter $s(n)$. If there are no queries that M makes on an
 389 input of size n , we set $s(n) = 0$.

390 Let $s_1(n) = \log(n)^{\log \log(n)}$, and $s_2(n) = \log(n)^{\log \log \log(n)}$. Moreover let $r(n) = n^{\log(n)}$.
 391 We define two Turing machines M_1 and M_2 that attempt to decide SAT using the oracle
 392 machine M .

393 Machine M_1 operates as follows. On input x of size n , M_1 computes strings $x_1 \dots x_{r(n)-n}$,
 394 where $x_i = f(x, 1^{n+i})$. Here f is the padding function for SAT. Let $x_0 = x$. M_1 iteratively
 395 does the following starting with $i = 0$. It simulates M on x_i . If an oracle query is made with
 396 parameter k , it checks if k is at most $s_1(n + i)$. Note that since M implements a natural

397 reduction, $k = s(n + i)$. If k is at most $s_1(n + i)$, it solves the corresponding MCSP instance
 398 by brute force search in time at most $\text{poly}(n + i)2^{s_1(n+i)^2}$ and continues the simulation.
 399 Note that any succeeding queries must also have parameter at most $s_1(n + i)$, therefore the
 400 simulation continues to completion. M_1 accepts iff M accepts. If no oracle queries are made,
 401 then too the simulation continues to completion, and M_1 returns the same answer as M .

402 In case the parameter k for a query is greater than $s_1(n + i)$, M_1 increments i , and repeats
 403 the process, as long as $i \leq r(n) - n$. If $i > r(n) - n$, M_1 simply runs the simulation of M on
 404 x to completion, now answering all queries by brute-force search regardless of the parameter
 405 size.

406 **Claim:** Either SAT is decidable in time $\text{poly}(r(n))2^{s_1(r(n))^2}$, or there is an infinite
 407 sequence of disjoint intervals $I_j = [n_j, n'_j)$ of input lengths such that $n'_j = r(n_j)$ for each j ,
 408 and moreover for each j , if $n \in I_j$, then $s(n) > s_1(n)$.

409 We establish the Claim. Suppose that SAT is not decidable in time $\text{poly}(r(n))2^{s_1(r(n))^2}$.
 410 We argue that for any $n_0 \in \mathbb{N}$, there exists an input length $n > n_0$ such that $s(n) \geq s_1(n)$,
 411 and moreover $s(m) \geq s_1(m)$ for every $n < m < r(n)$. Indeed, if this were not the case, for
 412 every $n > n_0$ one of the iterations of M_1 would succeed for some $i < r(n) - n$, and this would
 413 mean that the entire simulation would complete in time at most $\text{poly}(r(n))2^{s_1(r(n))^2}$. Note
 414 that when the simulation completes, it does solve SAT correctly, as M is an oracle Turing
 415 reduction correctly reducing SAT to MCSP, and any completing simulation uses correct
 416 answers to oracle queries.

417 Now we pick the intervals I_j inductively as follows. Let n_1 be the least integer such that
 418 $s(m) > s_1(m)$ for every $n_1 \leq m \leq r(n_1)$, and let $n'_1 = r(n_1)$. Inductively, suppose we have
 419 picked $n_1 \dots n_j$. We pick n_{j+1} to be the least integer such that $n_{j+1} > r(n_j)$, and moreover
 420 $s(m) > s_1(m)$ for each $n_{j+1} \leq m \leq r(n_{j+1})$. Such an integer exists by the argument of the
 421 previous paragraph. Let $n'_{j+1} = r(n_j)$, and set $I_j = [n_j, n'_j)$. This sequence of intervals is
 422 clearly disjoint, and satisfies the property in the Claim.

423 The first disjunct of the Claim implies that $\text{E} \not\subseteq \text{SIZE}(\text{poly})$, using Lemma 9 and the fact
 424 that $T(n) = \text{poly}(r(n))n^{s_1(r(n))}$ is time-constructible and satisfies $T(\text{poly}(T(\text{poly}(n)))) =$
 425 $o(2^n)$, for the given choice of r and s_1 . Hence, in the rest of our proof, we assume that the
 426 second disjunct holds, i.e., that there is an infinite sequence of long intervals within which
 427 each input length generates queries with a large parameter.

428 Now we define machine M_2 . On input x of length n , M_2 uses M to try to solve the search
 429 version of SAT, i.e., to find a satisfying assignment for x if one exists. M_2 uses the standard
 430 search-to-decision reduction for SAT based on self-reducibility, simulating M to answer any
 431 decision questions as described below. In order to find a satisfying assignment for x , this
 432 search-to-decision procedure might call the decision procedure for SAT on inputs x' such that
 433 $|x'| < |x|$; in such cases, M_2 pads up x' to length $|x|$ by using the padding function $f(x', 1^{|x|})$
 434 and runs the simulation of M on $f(x', 1^{|x|})$. This guarantees that all inputs on which M
 435 is run during the simulation have the same length. At the end of the search-to-decision
 436 procedure, if a satisfying assignment is successfully found, M_2 accepts, else it rejects.

437 Next we describe how M_2 uses M to decide if an input is in SAT. M_2 simulates the oracle
 438 machine M on the input. If M asks a query with parameter greater than $s_2(n)$, M_2 assumes
 439 the answer is ‘yes’ and continues the simulation. If M asks a query with parameter at most
 440 $s_2(n)$, M_2 solves the query by brute force search and continues the simulation. At the end of
 441 the simulation, M_2 accepts iff M accepts.

442 M_2 always halts within time $\text{poly}(n)2^{s_2(n)^2}$. Unlike machine M_1 , machine M_2 is not
 443 guaranteed to solve SAT correctly. However, since M_2 has been modified to solve the search
 444 version of SAT, the only way M_2 can make a mistake on x is to answer ‘no’ when the true

445 answer is ‘yes’. We will argue that whether or not M_2 solves SAT correctly, circuit lower
446 bounds follow.

447 We consider two cases. The first is that there is a $2^{s_2(n)^3}$ -robust set S on which M_2 solves
448 SAT correctly. In this case, we can apply Lemma 9 with $T(n) = \text{poly}(n)2^{s_2(n)^2}$, noting that
449 T is time-constructible and that $T(\text{poly}(n)) < 2^{s_2(n)^3}$ for large enough n . Thus we get that
450 $E \not\subseteq \text{SIZE}(\text{poly})$ in this case.

451 The remaining case is that there is no $2^{s_2(n)^3}$ -robust set S on which M_2 solves SAT
452 correctly. In this case, we apply the main idea of [12], using in addition our assumption that
453 there is an infinite sequence $\{I_j\}$ of disjoint intervals $[n_j, r(n_j))$ of input lengths such that
454 for each j and $n \in [n_j, r(n_j))$, we have that $s(n) > s_1(n)$.

455 If there is no $2^{s_2(n)^3}$ -robust set S on which M_2 solves SAT correctly, there must be a
456 constant k such that for each n , there is $m < 2^{ks_2(n)^3}$ such that M_2 fails to solve SAT
457 correctly on some input of length m . Now we use the existence of the sequence $\{I_j\}$ of
458 intervals from the Claim. Since $r(n) = n^{\log(n)}$ is significantly larger than $2^{ks_2(n)^3}$ for large
459 enough n , it follows that there is an infinite sequence of input lengths $\{m_i\}$ such that the
460 following hold: (i) M_2 fails to solve SAT correctly on some input x of length m_i such that
461 M on input x asks a query with parameter at least $s_1(m_i)$ (ii) For all input lengths m such
462 that $m_i \leq m \leq 2^{s_2(m_i)^4}$, $s(m) \geq s_1(m)$.

463 Now we apply the main idea of [12], by attempting to use the simulation M_2 itself to find
464 a small set of inputs for which M_2 fails to decide SAT correctly on at least one of these inputs.
465 We define a Turing machine A as follows. On input 1^m , A formulates the question of whether
466 there is an instance of size m on which M_2 fails to solve SAT correctly and asks a query with
467 parameter at least $s_2(m)$ as an instance y of SAT. Since M_2 runs in time $\text{poly}(n)2^{s_2(n)^2}$,
468 such an instance y of size at most $2^{s_2(m)^3}$ for large enough m can be computed in time at
469 most $2^{s_2(m)^3}$. A then runs M_2 on y to find a satisfying assignment for y , as in the proof
470 of the main result of [12]. For all $m = m_i$ for some i , either A succeeds, in which case it
471 finds an instance x of length m on which M_2 fails, or it finds a set of at most three instances
472 y_1, y_2, y_3 such that $|y_j| = |y|$ for each $1 \leq j \leq 3$, and M_2 fails on at least one of these three
473 instances. Note that by item (ii) in the para above, in the first case x must ask some query
474 with parameter at least $s_1(m)$ that is answered wrongly, and in the second case, one of
475 y_1, y_2, y_3 must ask some query with parameter at least $s_1(|y|)$ that is answered wrongly. In
476 either case, a wrongly answered query must be the truth table of a hard function. In the
477 first case, we have that the concatenation Z_i of all queries asked by x is the truth table of a
478 function on $O(\log(m))$ input bits that does not have circuits of size $s_1(m)$, and in the second
479 case we have that the concatenation Z_i of all queries asked by y_1, y_2, y_3 must be the truth
480 table of a function on $O(\log(|y|))$ input bits that does not have circuits of size $s_1(|y|)$. A
481 outputs Z_i .

482 In either case, A runs in time $\text{poly}(|Z_i|)2^{s_2(|Z_i|)^3}$ and outputs a string Z_i that does not
483 have circuits of size $s_1(|Z_i|)$. Since $s_2(|Z|) = s_1(|Z|)^{o(1)}$, this implies that $E \not\subseteq \text{SIZE}(\text{poly})$. ◀

484 5 Open Problems

485 The main open problem is to derive circuit lower bounds from unrestricted Turing reductions
486 from SAT to MCSP. One obstacle here is that we don’t we know that MKtP is not hard for
487 E under Turing reductions - this is a potentially simpler ‘test case’ that could be indicative.
488 We note that [16] have shown that a gap version of MKtP is not hard for E, when the gap
489 between the upper and lower thresholds for Kt-complexity is $\omega(\log(n))$.

490 An easier problem is to derive circuit lower bounds from unrestricted many-one reductions

491 from SAT to MCSP. Even this is unknown, though we do know that such reductions imply
 492 $\text{EXP} \neq \text{ZPP}$ [21, 17].

493 — **References** —

- 494 **1** Eric Allender. The complexity of complexity. In *Computability and Complexity - Essays*
 495 *Dedicated to Rodney G. Downey on the Occasion of His 60th Birthday*, pages 79–94, 2017.
- 496 **2** Eric Allender, Harry Buhrman, Michal Koucký, Dieter van Melkebeek, and Detlef Ronneburger.
 497 Power from random strings. *SIAM J. Comput.*, 35(6):1467–1493, 2006.
- 498 **3** Eric Allender and Bireswar Das. Zero knowledge and circuit minimization. In *Symposium on*
 499 *Mathematical Foundations of Computer Science (MFCS)*, pages 25–32, 2014.
- 500 **4** Eric Allender, Joshua A. Grochow, and Cristopher Moore. Graph isomorphism and circuit
 501 size. *CoRR*, abs/1511.08189, 2015.
- 502 **5** Eric Allender, Lisa Hellerstein, Paul McCabe, Toniann Pitassi, and Michael E. Saks. Minimizing
 503 DNF formulas and AC0 circuits given a truth table. *Electronic Colloquium on Computational*
 504 *Complexity (ECCC)*, (126), 2005.
- 505 **6** Eric Allender and Shuichi Hirahara. New insights on the (non-)hardness of circuit minimization
 506 and related problems. In *International Symposium on Mathematical Foundations of Computer*
 507 *Science (MFCS)*, pages 54:1–54:14, 2017.
- 508 **7** Eric Allender, Dhiraj Holden, and Valentine Kabanets. The minimum oracle circuit size
 509 problem. In *International Symposium on Theoretical Aspects of Computer Science (STACS)*,
 510 pages 21–33, 2015.
- 511 **8** Eric Allender, Michal Koucký, Detlef Ronneburger, and Sambuddha Roy. The pervasive reach
 512 of resource-bounded Kolmogorov complexity in computational complexity theory. *J. Comput.*
 513 *Syst. Sci.*, 77(1):14–40, 2011.
- 514 **9** Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge
 515 University Press, 1st edition, 2009.
- 516 **10** Sebastian Czort. The complexity of minimizing disjunctive normal form formulas. Master’s
 517 Thesis, University of Aarhus, 1999.
- 518 **11** Lance Fortnow and Rahul Santhanam. Robust simulations and significant separations.
 519 *Information and Computation*, 256:149–159, 2017.
- 520 **12** Dan Gutfreund, Ronen Shaltiel, and Amnon Ta-Shma. If NP languages are hard on the worst-
 521 case, then it is easy to find their hard instances. *Computational Complexity*, 16(4):412–441,
 522 2007.
- 523 **13** Juris Hartmanis and Richard Stearns. On the computational complexity of algorithms.
 524 *Transactions of the American Mathematical Society*, pages 285–306, 1965.
- 525 **14** Shuichi Hirahara, Igor Carboni Oliveira, and Rahul Santhanam. Np-hardness of minimum
 526 circuit size problem for OR-AND-MOD circuits. In *33rd Computational Complexity Conference,*
 527 *CCC 2018*, pages 5:1–5:31, 2018.
- 528 **15** Shuichi Hirahara and Rahul Santhanam. On the average-case complexity of MCSP and its
 529 variants. In *Computational Complexity Conference (CCC)*, pages 7:1–7:20, 2017.
- 530 **16** Shuichi Hirahara and Osamu Watanabe. Limits of minimum circuit size problem as oracle. In
 531 *Conference on Computational Complexity (CCC)*, pages 18:1–18:20, 2016.
- 532 **17** John M. Hitchcock and Aduri Pavan. On the NP-completeness of the minimum circuit size
 533 problem. In *Conference on Foundation of Software Technology and Theoretical Computer*
 534 *Science (FSTTCS)*, pages 236–245, 2015.
- 535 **18** Valentine Kabanets and Jin-yi Cai. Circuit minimization problem. In *Symposium on Theory*
 536 *of Computing (STOC)*, pages 73–79, 2000.
- 537 **19** Richard Karp and Richard Lipton. Some connections between nonuniform and uniform
 538 complexity classes. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing,*
 539 *April 28-30, 1980, Los Angeles, California, USA*, pages 302–309, 1980.
- 540 **20** William J. Masek. Some NP-complete set covering problems. Unpublished Manuscript, 1979.

- 541 21 Cody Murray and Ryan Williams. On the (non) NP-hardness of computing circuit complexity.
542 In *Conference on Computational Complexity (CCC)*, pages 365–380, 2015.
- 543 22 Igor Carboni Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit
544 lower bounds, and pseudorandomness. In *Computational Complexity Conference (CCC)*, pages
545 18:1–18:49, 2017.
- 546 23 Alexander A. Razborov and Steven Rudich. Natural proofs. *J. Comput. Syst. Sci.*, 55(1):24–35,
547 1997.
- 548 24 Boris A. Trakhtenbrot. A survey of Russian approaches to perebor (brute-force searches)
549 algorithms. *IEEE Annals of the History of Computing*, 6(4):384–400, 1984.