

1 $AC^0[p]$ Lower Bounds and NP-Hardness for 2 Variants of MCSP

3 **Rahul Ilango**

4 Rutgers University, Piscataway, USA

5 rahul.ilango@rutgers.edu

6 — Abstract —

7 The Minimum Circuit Size Problem (MCSP) asks whether a (given) Boolean function has a circuit
8 of at most a (given) size. Despite over a half-century of study, we know relatively little about the
9 computational complexity of MCSP. We do know that questions about the complexity of MCSP
10 have significant ramifications on longstanding open problems. In a recent development, Golovnev *et*
11 *al.* [11] improve the status of unconditional lower bounds for MCSP, showing that $MCSP \notin AC^0[p]$
12 for any prime p . While their results generalize to most “typical” circuit classes, it fails to generalize
13 to the circuit minimization problem for depth- d formulas, denoted (AC_d^0) -MCSP. In particular, their
14 result relies on a Lipchitz hypothesis that is unknown (and possibly false) in the case of (AC_d^0) -MCSP.
15 Despite this, we show that (AC_d^0) -MCSP $\notin AC^0[p]$ by proving even the failure of the Lipchitzness for
16 AC_d^0 formulas implies that $MAJORITY \leq_{tt}^{AC^0} (AC_d^0)$ -MCSP. Somewhat remarkably, our proof (in the
17 case of non-Lipchitzness) uses completely different techniques than [11]. To our knowledge, this is
18 the first MCSP reduction that uses modular properties of a function’s circuit complexity.

19 We also define MOCSP, an oracle version of MCSP that takes as input a Boolean function f , a
20 size threshold s , and oracle Boolean functions f_1, \dots, f_t , and determines whether there is an oracle
21 circuit of size at most s that computes f when given access to f_1, \dots, f_t . We prove that MOCSP
22 is NP-complete under non-uniform AC^0 many-one reductions as well as (uniform) ZPP truth table
23 reductions. We also observe that improving this ZPP reduction to a deterministic polynomial-time
24 reduction requires showing $EXP \neq ZPP$ (using theorems of Hitchcock and Pavan [17] and Murray and
25 Williams [22]). Optimistically, these MOCSP results could be a first step towards NP-hardness results
26 for MCSP. At the very least, we believe MOCSP clarifies the barriers towards proving hardness for
27 MCSP and provides a useful “testing ground” for questions about MCSP.

28 **2012 ACM Subject Classification** Theory of computation \rightarrow Circuit complexity; Theory of compu-
29 tation \rightarrow Problems, reductions and completeness

30 **Keywords and phrases** Minimum Circuit Size Problem, reductions, NP-completeness, circuit lower
31 bounds

32 **Acknowledgements** I would like to give a special thanks to Eric Allender for innumerable suggestions
33 and perspectives during all stages of this work. To name just a single example, one of his suggestions
34 led me to improve a PARITY-reduction to the presented MAJORITY-reduction for (AC_d^0) -MCSP. I
35 would also like to thank Abhishek Bhrushundi and Aditi Dudeja for their help in results about
36 constant depth formulas that lead to this paper. I am grateful to Ryan Williams for asking interesting
37 questions and helping to improve the exposition of the paper. Finally, I thank Harry Buhrman,
38 Lance Fortnow, Aditya Potukuchi, and Michael Saks for answering my questions and engaging in
39 many useful discussions.

40 **1 Introduction**

41 The Minimum Circuit Size Problem (MCSP) takes as input a Boolean function f (represented
42 by its truth table) and a size parameter s and asks if there is a circuit of size at most s
43 computing f . Study of this problem began in the 1950s by complexity theorists in the
44 Soviet Union [30], where MCSP was of such great interest that Levin is said to have delayed

45 publishing his initial NP-completeness results in hope of showing that MCSP is NP-complete.¹
 46 Interest in MCSP was revitalized when Kabanets and Cai [19] connected the problem with
 47 the natural proofs framework of Razborov and Rudich [27]. Since then, MCSP has been the
 48 subject of intense research. We begin by reviewing some of this work.

49 1.1 Known lower bounds, hardness, and non-hardness for MCSP

50 It is easy to see that MCSP is in NP (the circuit of size at most s can be used as a witness),
 51 but, despite work by numerous researchers, the exact complexity of MCSP remains unknown.

52 **Lower bounds and hardness results.** We believe MCSP is not easy to compute. Kabanets
 53 and Cai [19] show that MCSP \notin P conditioned on a widely-believed cryptographic hypothesis,
 54 and Allender and Das [2] show that MCSP is hard for SZK under BPP-Turing reductions.

55 Unconditionally, we know lower bounds against MCSP for restricted classes of circuits.
 56 Hirahara and Santhanam [15] show that MCSP requires nearly quadratic sized DeMorgan
 57 formulas, and Allender *et al.* [1] prove that MCSP \notin AC⁰. In a recent paper, Golovnev *et al.*
 58 [11] improve the latter result, showing that MCSP requires exponential-sized AC⁰[p] circuits
 59 by proving MAJORITY \in (AC⁰)^{MCSP}. The MAJORITY hardness result of [11] generalizes to
 60 the circuit minimization problem for many circuit classes, however, the techniques fail in the
 61 case of constant depth formulas.

62 Under weak reductions, we know MCSP is hard for some subclasses of P. Oliveira and
 63 Santhanam [25] prove that MCSP is hard for DET under TC⁰ truth table reductions, and
 64 Golovnev *et al.* [11] use the results of [25] to show that NC¹ \subseteq (AC⁰)^{MCSP}. Surprisingly, we
 65 know stronger results for the “program” variant of MCSP, MKTP. Allender and Hirahara
 66 [3] show that MKTP is hard for DET under NC⁰ many-one reductions, and Hirahara and
 67 Santhanam [15] show average-case lower bounds for MKTP against AC⁰[p].

68 The most natural question is whether MCSP is NP-complete. As of yet, we have not
 69 managed to uncover even strong supporting evidence for, or against, MCSP being NP-
 70 complete. We do know that the circuit minimization problem is NP-complete for some
 71 restricted classes of circuits: DNF circuits by Masek [20] and OR \circ AND \circ MOD _{m} circuits
 72 by Hirahara, Oliveira, and Santhanam [14]. Impagliazzo, Kabanets, and Volkovich [18]
 73 show that if there exist Indistinguishability Obfuscators against randomized polynomial-time
 74 algorithms, then MCSP \in ZPP \iff NP = ZPP.

75 **Known non-hardness results.** The unconditional non-hardness results for MCSP rule out
 76 NP-hardness under certain types of reductions. For example, Hirahara and Watanabe [16]
 77 show that “oracle-independent reductions” cannot show that MCSP is hard for either a class
 78 larger than P under polynomial-time Turing reductions or a class larger than AM \cap coAM
 79 under BPP reductions with one query to MCSP. Moreover, while most NP-complete problems
 80 are complete under rather weak reductions such as TIME[$n^{o(1)}$] or AC⁰ many-one reductions,
 81 Murray and Williams [22] prove that MCSP is not NP-hard under TIME[$n^{.49}$] reductions,
 82 and Allender, Ilango, and Vafa [5] show that a super-linear approximations of MCSP cannot
 83 be NP-hard under even non-uniform AC⁰ many-one reductions.

84 Conditioned on a widely-believed cryptographic hypothesis, Allender and Hirahara [3]
 85 show that a very weak approximation of MCSP is NP-intermediate.

¹ [6] cites a personal communication from Levin regarding this story.

1.2 Implications of lower bounds and hardness for MCSP

While we have not managed to establish the complexity of MCSP, a series of works, beginning with Kabanets and Cai [19], connect the computational complexity of MCSP and its variants to longstanding open questions in the field.

Separations of complexity classes. Several works ([17], [22], [2], [19]) show that MCSP being NP-hard, under various notions of reducibility, implies unknown class separations. For example, Hitchcock and Pavan [17] and Murray and Williams [22] show that if MCSP is NP-hard under polynomial-time truth-table reductions, then $ZPP \neq EXP$, a major open problem.²

Worst-case versus average-case complexity for NP. Using tools developed by Nisan and Wigderson [23] and Carmosino, Impagliazzo, Kabanets, and Kolokova [9], Hirahara [13] gives a “worst-case to average-case” reduction for NP conditioned on a certain approximation to MCSP being NP-hard. Thus, if one could show this approximation to MCSP is NP-hard, the worst-case and average-case complexity of NP would be equivalent.

Circuit Lower Bounds. Recent work by Oliveira, Pich, and Santhanam ([26] and [24])³ explores a phenomenon they term “hardness magnification,” whereby even weak circuit lower bounds on certain computational problems imply strong lower bounds on other problems. For example, [26] shows that if MCSP cannot be solved on average with no error by linear-size formulas, then NP does not have polynomial-size formulas. [24] shows that if a certain approximation to MCSP cannot be computed by circuits of size $n^{1+\epsilon}$, then NP does not have polynomial-sized circuits.

1.3 Our Contributions

In this work, we focus on hardness results for variants of MCSP, in particular establishing an $AC^0[p]$ lower bound and an NP-hardness result.

MAJORITY-hardness for (AC_d^0) -MCSP

As mentioned previously, Golovnev *et al.* [11] proves that $MAJORITY \in (AC^0)^{MCSP}$. Using similar techniques, they also show that, for restricted classes of circuits \mathcal{C} such as formulas and constant depth circuits, the \mathcal{C} -circuit minimization problem, denoted (\mathcal{C}) -MCSP, is hard for MAJORITY under AC^0 reductions. For these MAJORITY reductions to work, [11] requires that the size of the minimum \mathcal{C} -circuit on truth tables of length n is roughly $(n^{.49})$ -Lipchitz.

This Lipchitzness hypothesis is unknown (and perhaps even false) in the class of depth- d formulas, which we denote AC_d^0 .⁴ Despite this, we prove MAJORITY-hardness for (AC_d^0) -MCSP by giving a MAJORITY reduction that works in the case that Lipchitzness fails. Applying the lower bounds of Razborov [27] and Smolensky [29] then gives an $AC^0[p]$ lower bound for (AC_d^0) -MCSP.

² [22] only shows the result under many-one reductions, but their techniques easily generalize to the truth table case. [17] explicitly proves the truth table result using a different approach than [22].

³ Pich is an author on [24] but not [26].

⁴ We will always use the notation AC_d^0 to refer to depth- d formulas and never depth- d circuits.

121 ► **Theorem 1.1.** *Let $d \geq 2$. Then MAJORITY $\leq_{tt}^{AC^0}$ (AC_d⁰)-MCSP. Consequently, (AC_d⁰)-MCSP \notin*
 122 *AC⁰[p] for any prime p .*

123 Remarkably, the techniques used for this MAJORITY reduction (in the case of non-
 124 Lipchitzness) are entirely different than the ones used by [11] for general MCSP. Indeed, the
 125 non-Lipchitz case reduction we present is of a very different flavor than, to our knowledge,
 126 all known MCSP hardness results. As far as the author knows, it is the only MCSP hardness
 127 result that does not easily generalize to an approximation of MCSP. This is because the key
 128 step in the reduction is determining, exactly, a Boolean function’s circuit complexity modulo
 129 a certain prime.

130 ► **Open Question 1.2.** *Can one extend Theorem 1.1 to an approximation of MCSP?*

131 We also remark that our notion of size for AC_d⁰ formulas is critical for Theorem 1.1. We
 132 define the size of an AC_d⁰ formula to be the number of input leaves. While this is the standard
 133 definition of formula size, we make heavy use of elementary direct product theorems known,
 134 specifically, for this notion of formula size. It is not clear to us how to generalize Theorem
 135 1.1 to the case when the size of an AC_d⁰ formula is, say, the number of gates or the number of
 136 wires.

137 NP-Hardness of oracle MCSP (MOCSP)

138 Some work has been done trying to approach the NP-hardness of MCSP “from below,” that
 139 is, proving that the circuit minimization problem is NP-hard for restricted classes of circuits.
 140 As mentioned previously, we know that (DNF)-MCSP [20] and (OR ◦ AND ◦ MOD_m)-MCSP
 141 [14] are NP-hard.

142 Instead, we attempt to approach MCSP from “above.” We formulate the Minimum Oracle
 143 Circuit Size Problem, denoted MOCSP, that takes as input a truth table T , a size parameter
 144 $s \in \mathbb{N}$, and auxiliary truth tables T_1, \dots, T_t and asks whether there is an oracle circuit
 145 of size at most s that computes T when given access to T_1, \dots, T_t . It is easy to see that
 146 MOCSP \in NP (the oracle circuit of size s acts as a witness).

147 We note that this is not the first time someone has considered an “oracle version” of
 148 MCSP. Allender *et al.* [1] and Allender, Holden, and Kabanets [4] consider the problem of
 149 minimizing oracle circuits for a fixed oracle A . We will denote this problem MCSP^A. An
 150 important result for this problem that [1] proves is that MCSP^{QBF} is complete for PSPACE
 151 under ZPP reductions. MOCSP differs from MCSP^A in that the oracle circuit gets access to
 152 a finite number of Boolean functions, not a language, and the functions the oracle circuit has
 153 access to are *inputs* to the problem.

154 In our view, MOCSP has two advantages over MCSP^A. First, MOCSP \in NP while the
 155 complexity of MCSP^A depends on the oracle A . Second, there is an easy reduction from
 156 MCSP to MOCSP, simply provide no oracle truth tables. Therefore, we can use MOCSP as a
 157 testing ground for hardness results we conjecture for MCSP. Thus, the most natural question
 158 is whether we can prove that MOCSP is NP-hard. We prove that MOCSP is indeed NP-hard
 159 under non-uniform AC⁰ reductions and under uniform randomized reductions.

160 ► **Theorem 1.3.** ■ NP $\leq_m^{AC^0}$ MOCSP

161 ■ NP \leq_m^{RP} MOCSP

162 ■ NP \leq_{tt}^{ZPP} MOCSP

163 These NP-hardness results are all proved by giving a reduction from approximating r -
 164 bounded set cover to MOCSP. It is worth noting that the NP-hardness results of (DNF)-MCSP
 165 [20] and (OR ◦ AND ◦ MOD_m)-MCSP [14] are also proved via set cover problems.

166 Given that we can show MOCSP is NP-hard under randomized reductions, one might even
 167 begin to hope that we can prove hardness under, say, polynomial-time truth table reductions.
 168 Unfortunately, this seems difficult. Essentially the same proofs Murray and Williams [22]
 169 or Hitchcock and Pavan [17] use to show that MCSP being NP-hard under polynomial-time
 170 truth table reductions implies $\text{EXP} \neq \text{ZPP}$ also works for MOCSP.

171 ► **Theorem 1.4** (Essentially proven in [17] and [22]). *If $\text{NP} \leq_{tt}^P \text{MOCSP}$, then $\text{EXP} \neq \text{ZPP}$.*

172 Thus, improving our ZPP reduction to a P reduction requires separating EXP from ZPP, a
 173 longstanding open problem. For completeness, we give the MOCSP version of Murray and
 174 Williams’ proof in Appendix B.

175 Even so, we expect that the ground truth is that MOCSP is NP-hard under, at least,
 176 polynomial-time Turing reductions.

177 ► **Conjecture 1.5.** $\text{NP} \leq_T^P \text{MOCSP}$.

178 We give some details on why we believe Conjecture 1.5 near the end of Section 4. Even so,
 179 we believe proving such a hardness result is beyond current techniques. Perhaps one could
 180 even prove there are some barriers.

181 **A perspective on MOCSP and some questions.**

182 In light of the fact that hardness for MCSP beyond SZK under even non-uniform reductions
 183 is unknown, we found these MOCSP hardness results to be quite surprising. To an optimist,
 184 NP-hardness results for MOCSP could even be a first step towards proving hardness for
 185 MCSP. Indeed, a PSPACE-hardness result was first proved by Buhrman and Torenvliet [8]
 186 for an “oracle” version of space-bounded Kolmogorov complexity before Allender *et al.* [1]
 187 showed PSPACE-hardness for the non-oracle version about four years later.⁵

188 Even if stronger hardness results for MCSP remain out of reach, MOCSP could still yield
 189 valuable insights about MCSP. For instance, it would be interesting to see which of the
 190 barriers and non-hardness results known for MCSP carry over to MOCSP.

191 ► **Open Question 1.6.** *Can one show that other barriers or non-hardness results that hold
 192 for MCSP also hold for MOCSP?*

193 As an example of the insight given by answers to this question, consider Murray and Williams’
 194 [22] result that proving MCSP is NP-hard under polynomial-time many-one reductions implies
 195 $\text{EXP} \neq \text{ZPP}$. A natural question one might ask is whether we can improve this theorem
 196 to show that MCSP being NP-hard under randomized reductions implies unknown class
 197 separations. As we note in Theorem 1.4, however, Murray and Williams’ proof carries over
 198 to MOCSP, and Theorem 1.3 shows MOCSP is indeed NP-hard under randomized reductions.
 199 Thus, any improvement of Murray and Williams’ result to randomized reductions likely
 200 requires a fact about MCSP that we do not know for MOCSP.

201 In another direction, our results seem to imply that proving hardness for MOCSP is easier
 202 than proving hardness for MCSP. Indeed, since MCSP easily reduces to MOCSP, any hardness
 203 result that is true for MCSP must also be true for MOCSP. Therefore, we can use MOCSP
 204 as a testing ground for hardness results we conjecture about MCSP. For example, Hirahara’s
 205 [13] worst-case to average-case reduction for NP can be based on a certain approximation
 206 of MCSP being NP-hard, which would imply that a certain approximation of MOCSP is

⁵ The conference versions of [8] and [1] are four years apart.

207 also NP-hard. Given that we can prove the NP-hardness of MCSP under uniform ZPP
 208 reductions and non-uniform AC⁰ reductions, we ask if one can prove something similar for
 209 the approximation version of MCSP.

210 ► **Open Question 1.7.** *Can one prove that, for some $\epsilon > 0$, approximating MCSP on*
 211 *n -inputs to a factor of n^ϵ is NP-hard under, say, P/poly reductions? Conversely, can one*
 212 *prove that there is any barrier to showing such a hardness result?*

213 We note that the techniques we use to prove NP-hardness results for MCSP seem to break
 214 down completely in the case of even super-constant approximation, so answering this question
 215 will likely require new ideas.

216 1.4 Proof Overviews

217 In this section, we give fairly detailed overviews of our proofs. In doing this, we will often
 218 state results without filling in low-level details. To make clear to the reader when we are
 219 doing this, we mark such sentences with an italicized *we observe*.

220 Majority Hardness for (AC_d⁰)-MCSP

221 Recall, AC_d⁰ is the class of depth- d formulas. We also define AND \circ AC_{d-1}⁰ and OR \circ AC_{d-1}⁰
 222 be the classes of AC_d⁰ formulas with a top AND and top OR gate respectively. For $\mathcal{F} \in$
 223 $\{\text{AC}_d^0, \text{AND} \circ \text{AC}_{d-1}^0, \text{OR} \circ \text{AC}_{d-1}^0\}$ and a truth table T , we let $\text{CC}_{\mathcal{F}}(T)$ denote the size of the
 224 minimum \mathcal{F} -formula computing T where the size of a formula is the number of input leaves.

225 Our analysis proceeds by considering each $n \in \mathbb{N}$ and splitting into cases depending on
 226 whether $\text{CC}_{\text{AC}_d^0}$ is Lipchitz on truth tables of length around n . In more detail, fix some
 227 sufficiently large n . Let $q = \Theta(n^2)$ be a power of two. We divide into cases depending on
 228 whether there exists an $m \in \{q^{10}, q^{50}\}$ such that $\text{CC}_{\text{AC}_d^0}$ is (m^{25}) -Lipchitz on truth tables of
 229 length m .

230 Case 1: Lipchitzness holds for some m .

231 If there does exist an $m \in \{q^{10}, q^{50}\}$ such that $\text{CC}_{\text{AC}_d^0}$ is (m^{25}) -Lipchitz on truth tables of
 232 length m , then the techniques of [11] yield an AC⁰ truth table reduction from MAJORITY
 233 on n -bits to (AC_d⁰)-MCSP on m -bits. For completeness, we include a self-contained proof of
 234 this case in Appendix A.

235 Case 2: Lipchitzness fails.

236 Assume that for all $m \in \{q^{10}, q^{50}\}$ $\text{CC}_{\text{AC}_d^0}$ is not (m^{25}) -Lipchitz on truth tables of length m .
 237 Let $u = q^{10}$ and $v = q^{50}$.

238 **Lipchitzness failing \implies functions easier to compute with a top AND gate.** *We observe,*
 239 *as a straight forward consequence of Lipchitzness failing, that there exists a truth table of*
 240 *length u that has an optimal formula with large top fan-in and and a truth table of length v*
 241 *that is easier to compute with a top AND gate:*

- 242 1. There exists a Boolean function f^u that takes $\log u$ inputs and an AC_d⁰ formula ϕ^u such
 243 that ϕ^u is an optimal AC_d⁰ formula for f^u and $\phi^u = \phi_1^u \wedge \dots \wedge \phi_t^u$ for some $t \geq n$ and
 244 some $\phi_1^u, \dots, \phi_t^u \in \text{AC}_{d-1}^0$.

245 2. There exists a Boolean function f^v that takes $\log v$ inputs such that $\text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(f^v) >$
 246 $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^u) + u \log u$.
 247 We will make use of f^u and ϕ^u to reduce MAJORITY to $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}$ and we will use f^v to
 248 reduce $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}$ to $\text{CC}_{\text{AC}_d^0}$.

249 **Using $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}$ and optimal subformulas of ϕ^u to compute a dot product.** The heart
 250 of our MAJORITY reduction is a fairly elementary observation about optimal $(\text{AND} \circ \text{AC}_{d-1}^0)$
 251 formulas. Recall, $\phi^u = \phi_1^u \wedge \cdots \wedge \phi_t^u$ is an optimal (AC_d^0) formula for f^u and, hence, also an
 252 optimal $(\text{AND} \circ \text{AC}_{d-1}^0)$ formula for f^u . We observe that for any $A \subseteq [t]$, the $(\text{AND} \circ \text{AC}_{d-1}^0)$ -
 253 optimality of ϕ^u implies that $\bigwedge_{i \in A} \phi_i^u$ is also an optimal $(\text{AND} \circ \text{AC}_{d-1}^0)$ formula for the
 254 function it computes.

255 Introducing some notation, for a string $x \in \{0, 1\}^n$, we let f_x^u be the function given
 256 by $\bigwedge_{i \in O_x} \phi_i^u$ where $O_x \subseteq [n]$ are the bits in x that are one. Using the above observation
 257 about subformulas being optimal, we have that $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u) = \sum_{i \in O_x} |\phi_i^u|$.⁶ Thus,
 258 one can think of $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)$ as computing the dot product between x and the vector
 259 $\langle |\phi_1^u|, \dots, |\phi_n^u| \rangle$.

260 Note that that the definition of f_x^u depends on the labeling of $\phi_1^u, \dots, \phi_t^u$, in particular the
 261 choice of which ϕ_i^u have $i \leq n$. We will later choose an labeling of the ϕ_i^u that is convenient.

262 **Computing MAJORITY (non-uniformly) using $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}$.** Our goal is to compute
 263 MAJORITY on a string $x \in \{0, 1\}^n$ using the above “dot product” observation. Before
 264 we show how to do this, we give some intuition on how we came up with the idea.

265 Instead of trying to compute MAJORITY, suppose we relaxed the problem to computing
 266 PARITY given access to the integer produced by the dot product $x \cdot \langle |\phi_1^u|, \dots, |\phi_n^u| \rangle$. Well,
 267 if it so happened that all the entries in the vector $\langle |\phi_1^u|, \dots, |\phi_n^u| \rangle$ were odd, then it is clear
 268 that the integer produced by $x \cdot \langle |\phi_1^u|, \dots, |\phi_n^u| \rangle$ is odd if and only if x has an odd number of
 269 ones. Our approach for MAJORITY is a generalization of this.

270 Let $p = O(n)$ be prime greater than n . We observe, via an averaging argument, that
 271 there exists integers $k \geq 0$ and $1 \leq r \leq p - 1$ such that (after relabeling the ϕ_i^u)

$$272 \quad |\phi_1^u|/p^k \equiv \cdots \equiv |\phi_n^u|/p^k \equiv r \pmod{p}.$$

273 Thus, we can determine the weight w of x (and hence compute MAJORITY of x) by computing
 274 the value of

$$275 \quad \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)/p^k = \sum_{i \in O_x} |\phi_i^u|/p^k \equiv rw \pmod{p}$$

276 and multiplying by the inverse of r modulo p .⁷

277 **Reducing computing $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}$ to computing $\text{CC}_{\text{AC}_d^0}$.** Ultimately, we want to compute
 278 MAJORITY using $\text{CC}_{\text{AC}_d^0}$ not $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}$. By the above procedure, it suffices to show how
 279 to compute $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)$ using $\text{CC}_{\text{AC}_d^0}$.

⁶ Recall, our notion of formula size is the number of input leaves.

⁷ In case the reader is unsure of whether the last parts of this procedure are implementable in AC^0 , realize that the output of $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)$ is a binary string of length $O(\log n)$ and that any function on a string of length $O(\log n)$ can be computed by a polynomial-sized DNF. See the proof in Section 3 for more details.

280 We can make such a computation as follows. Recall that f^v is a function satisfying

$$281 \quad \text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(f^v) > \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v) + u \log u$$

282 whose existence is guaranteed by the failure of Lipchitzness. Take the direct product of f_x^u
 283 with f^v to obtain a function $g_x(y, z) = f_x^u(y) \wedge f^v(z)$. Since the difference between computing
 284 f^v with a top AND gate and a top OR gate is larger than $u \log u$ (which is the maximum
 285 complexity of f_x^u), we observe⁸ any optimal AC_d⁰ formula for g_x must have a top AND gate,
 286 so

$$287 \quad \text{CC}_{\text{AC}_d^0}(g_x) = \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(g_x).$$

288 Then, we observe that

$$289 \quad \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(g_x) = \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u) + \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v).$$

290 Hence, if we are non-uniformly given the value of $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v)$, we can subtract
 291 $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v)$ from $\text{CC}_{\text{AC}_d^0}(g_x)$ to find $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)$.

292 NP-hardness of Oracle MCSP (MOCSP)

293 We define the size of an oracle circuit to be the total number of AND, OR, and oracle gates.
 294 The Minimum Oracle Circuit Size Problem, MOCSP, takes as input a truth table T , a
 295 threshold $s \in \mathbb{N}$, and auxiliary truth tables T_1, \dots, T_t and outputs whether there is an oracle
 296 circuit of size at most s that computes T when given oracle access to T_1, \dots, T_t . We denote
 297 the output of MOCSP on such an input as $\text{MOCSP}(T, s; T_1, \dots, T_t)$. We denote the minimum
 298 size of any oracle circuit computing T when given access to T_1, \dots, T_t as $\text{CC}^{T_1, \dots, T_t}(T)$.

299 We prove that MOCSP is NP-hard under various reductions by giving a reduction from
 300 4-approximating r -bounded set cover, denoted 4-SetCover_r, to MOCSP. As a reminder,
 301 4-SetCover_r is the promise problem takes as input sets $S_1, \dots, S_t \subseteq [n]$ of cardinality at most
 302 r whose union is $[n]$ as well as an integer $c \in [n]$ and requires outputting YES when $c \geq \ell$
 303 and NO when $c < \ell/4$ where ℓ is the optimal cover size, i.e.

$$304 \quad \ell = \min\{|I| : I \subseteq [t] \text{ and } \bigcup_{i \in I} S_i = [n]\}.$$

305 For sufficiently large r , 4-SetCover_r is known to be NP-hard (see Theorem 2.6).

306 **Informal idea.** We begin by giving a high-level overview of the reduction to orient the
 307 reader. (It will be very informal, but we are building to a more detailed description.) Say
 308 we are given sets $S_1, \dots, S_t \subseteq [n]$ of cardinality at most r whose union is $[n]$. One can
 309 think of each of these sets S_i as “seeing” a small portion of the ground set $[n]$. For some
 310 carefully chosen truth table T of length $m \geq n$, we let each set S_1, \dots, S_t induce truth tables
 311 T_{S_1}, \dots, T_{S_t} respectively where each truth table T_{S_i} “sees” roughly the same part of T as S_i
 312 “sees” of $[n]$. Finally, we ask how hard it is for a circuit to compute T given oracle access to
 313 T_1, \dots, T_t , and we show that, if T has a certain property, then the answer to this question is
 314 the answer to the set cover problem up to a constant factor.

315 We now illustrate the algorithm in more detail. Fix sets $S_1, \dots, S_t \subseteq [n]$ of cardinality at
 316 most r whose union is $[n]$ and fix a truth table T of length m . Assume the optimal cover
 317 size of $[n]$ by S_1, \dots, S_t is ℓ .

⁸ both the “we observe” statements in this paragraph are consequences of standard direct product theorems for formulas.

318 **The truth tables induced by S_1, \dots, S_t and T .** We rigorously define the truth tables
 319 T_{S_1}, \dots, T_{S_t} of length m induced by S_1, \dots, S_t and T . First, we fix a partition of $[m]$ into n
 320 sets. It does not really matter what partition we choose as long as the sets are roughly the
 321 same size and the partition is easily computable, but, for concreteness, let $P_1, \dots, P_n \subseteq [m]$
 322 be the partition of $[m]$ given by $P_i = \{j \in [m] : j \equiv i \pmod n\}$.

323 We can then use this partition to “lift” any subset of $[n]$ into a subset of $[m]$ as follows.
 324 For a subset S of $[n]$, let S^m denote the subset of $[m]$ given by $S^m = \bigcup_{i \in S} P_i$.

325 Next, for a subset $P \subseteq [m]$, we let T_P be the truth table of length m that “sees” T on
 326 the elements of P and zeroes everywhere else, that is the i th bit of T_P is

$$327 \begin{cases} \text{the } i\text{th bit of } T & , \text{ if } i \in P \\ 0 & , \text{ otherwise.} \end{cases}$$

328 Finally, we define the truth table T_{S_i} induced by S_i to be the truth table $T_{S_i^m}$ given by
 329 the above notation (we are dropping the m superscript for concision).

330 **$\text{CC}^{T_{S_1}, \dots, T_{S_t}}(T)$ is at most 2ℓ .** Suppose, without loss of generality, that $S_1 \cup \dots \cup S_\ell$ is an
 331 optimal cover of $[n]$. Then, by construction, the function computed by $T_{S_1} \vee \dots \vee T_{S_\ell}$ is T .
 332 This is an oracle circuit of size $2\ell - 1$, so $\text{CC}^{T_{S_1}, \dots, T_{S_t}}(T) \leq 2\ell$.

333 **If T is (rn) -irritable, then $\text{CC}^{T_{S_1}, \dots, T_{S_t}}(T) > \ell/2$.** Recall the notation defined previously
 334 that, for a set $P \subseteq [m]$, T_P denotes the string of length m that equals T on the bits in P
 335 and is zero everywhere else. Also recall that we fixed a partition P_1, \dots, P_n of $[m]$. We say
 336 that T is (s) -irritable if for all $i \in [n]$ we have that

$$337 \text{CC}^{T_{P_1}, \dots, T_{P_{i-1}}, T_{P_{i+1}}, \dots, T_{P_n}}(T) > s.$$

338 Informally, T being (rn) -irritable means that if you take away access to any particular
 339 T_{P_i} oracle, then computing T requires an oracle circuit of size greater than rn , which is a
 340 $(r/2)$ -factor jump over the trivial $2n$ -sized oracle circuit given by $T_{P_1} \vee \dots \vee T_{P_n}$ if one had
 341 full oracle access.

342 Now assume T is (rn) -irritable. We need to show that $\text{CC}^{T_{S_1}, \dots, T_{S_t}}(T) > \ell/2$. For
 343 contradiction, suppose that C is an oracle circuit computing T with at most $\ell/2$ gates. Then
 344 C uses at most $q \leq \ell/2$ unique oracle gates. Without loss of generality, assume C uses
 345 as oracles only T_{S_1}, \dots, T_{S_q} . Next, note that for any $i \in [q]$, T_{S_i} can, by construction, be
 346 computed by the oracle circuit $\bigvee_{j \in S_i} T_{P_j}$. Moreover, this is an oracle circuit for T_{S_i} of size
 347 at most $2r$ since $|S_i| \leq r$. Thus, replacing each T_{S_i} oracle gate in C with the oracle circuit
 348 $\bigvee_{j \in S_i} T_{P_j}$, we can transform C into an oracle circuit D of size at most $r \cdot |C| \leq r\ell$ such that
 349 D computes T when given access to the oracles in the set $O = \{T_{P_j} : j \in S_1 \cup \dots \cup S_q\}$.
 350 However, since $q \leq \ell/2$ is less than the optimal cover size, $|S_1 \cup \dots \cup S_q| < n$ and so $|O| < n$,
 351 so O is missing $T_{P_{i^*}}$ for some $i^* \in [n]$. But then D is an oracle circuit of size at most $r\ell \leq rn$
 352 that computes T when given access to T_{P_1}, \dots, T_{P_n} without using $T_{P_{i^*}}$ as an oracle gate,
 353 which contradicts that T is (rn) -irritable.

354 **RP, ZPP and AC^0 reductions.** At this point, we have shown that one can compute whether
 355 S_1, \dots, S_t admits a c -cover (up to a 4-approximation) by outputting $\text{MOCSP}(T, 2c; T_{S_1}, \dots, T_{S_t})$
 356 for some T that is (rn) -irritable.

357 We observe by a counting argument that a truth table T of length $m \geq n^3$ picked uniformly
 358 at random is (rn) -irritable with high probability. Thus, picking a random truth table T of

length $\Theta(n^3)$ and outputting $\text{MOCSP}(T, 2c; T_{S_1}, \dots, T_{S_t})$ gives an RP many-one reduction from 4-SetCover_r to MOCSP (note that we get one-sided error because irritability was only required for the $\ell/2$ lower bound and not required for the 2ℓ upper bound). Additionally, since we can check if a random T is (rn) -irritable using an oracle to MOCSP , we observe that $4\text{-SetCover}_r \leq_{tt}^{\text{ZPP}} \text{MOCSP}$. Finally, we observe that there is an AC⁰ circuit C such that $C(T, c, S_1, \dots, S_t) = (T, 2c; T_{S_1}, \dots, T_{S_t})$. Therefore, by non-uniformly hardcoding an (rn) -irritable truth table T into C , we get that 4-SetCover_r reduces to MOCSP under (non-uniform) AC⁰ many-one reductions.

1.5 Paper Organization

In Section 2 we fix notation and review precise definitions. In Section 3 we prove that MAJORITY reduces to $(\text{AC}_d^0)\text{-MCSP}$, and in Section 4 we prove our MOCSP results.

2 Preliminaries

For an integer n , we let $[n]$ denote the set $\{1, \dots, n\}$. For a binary string $x \in \{0, 1\}^*$, the *weight* of x , denoted $\text{wt}(x)$, is the number ones in x . We identify a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ with its truth table $T \in \{0, 1\}^{2^n}$ and often use them interchangeably.

We let \log denote the base-2 logarithm and \ln represent the base- e logarithm. For functions f and g , we say $f = \tilde{O}(g)$ if there exists a c such that $f(x) \leq \log^c(g(x))g(x)$ for all sufficiently large x . We say that $f = \tilde{\Omega}(g)$ if $g = \tilde{O}(f)$.

We say a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ is c -Lipchitz if for all $x, y \in \{0, 1\}^n$ that differ in at most one bit, $|f(x) - f(y)| \leq c$.

Complexity classes and reductions

We assume the reader is familiar with the standard complexity classes such as AC⁰, P, ZPP, RP, NP, E and the notion of Turing machines. For background on these, we refer to Arora and Barak's excellent textbook [7]. For us, AC⁰ always refers to *non-uniform* AC⁰.

We review the types of reductions we use in case the reader is not familiar with randomized reductions, truth table reductions, or our notation.

Many-one reductions. We will make use of the follow notions of many-one reduction.

- $L \leq_m^{\text{AC}^0} L'$ if there is a non-uniform (polynomial-sized) AC⁰ circuit C such that $x \in L \iff C(x) \in L'$.
- $L \leq_m^{\text{P}} L'$ if there is a polynomial-time Turing machine M such that $x \in L \iff M(x) \in L'$.
- $L \leq_m^{\text{RP}} L'$ if there is a polynomial-time probabilistic Turing machine M taking in a “random” auxiliary input r such that

$$x \in L \implies \forall r M(x, r) \in L', \text{ and}$$

$$x \notin L \implies \Pr_r[M(x, r) \in L'] \geq 2/3,$$

and $|r|$ is polynomial in the length of x .

396 **Truth table reductions.** We will also make use of the following notions of truth table
397 reductions.

- 398 ■ We say an oracle circuit C is a *truth table oracle circuit* if there is no directed path
399 between oracle gates in C .
- 400 ■ $L \leq_{tt}^{AC^0} L'$ if there is a non-uniform (polynomial-sized) AC^0 truth table oracle circuit C
401 such that C computes L when given oracle access to L' .
- 402 ■ $L \leq_{tt}^{ZPP} L'$ if L can be computed with zero-error by a polynomial-time probabilistic oracle
403 Turing machine M with oracle access to L' with the caveat that all of M 's oracle queries
404 must be answered simultaneously (i.e. so no oracle query can depend on another oracle
405 query). On any single input, M is allowed to output “don't know” with probability at
406 most $1/2$.

407 AC_d^0 formulas, (AC_d^0) -MCSP, and $CC_{AC_d^0}$

408 For an integer $d \geq 2$, we let AC_d^0 denote the class of depth- d formulas that use AND and OR
409 gates with unbounded fan-in and fan-out 1 and that takes as “input leaves” the bits of a
410 binary string and the negation of each of those bits.

411 For an AC_d^0 formula ϕ , we define the size of ϕ , denoted $|\phi|$, to be the total number of
412 input leaves ϕ uses. For a Boolean function f , we let $CC_{AC_d^0}(f)$ be the size of the smallest
413 AC_d^0 formula computing f .

414 ► **Definition 2.1** (Minimum Circuit Size Problem for constant depth formulas). (AC_d^0) -MCSP,
415 is the language given by

$$416 \{(T, s) \in \{0, 1\}^* \times \mathbb{N} : T \text{ is the truth table of a Boolean function, and } CC_{AC_d^0}(T) \leq s\}.$$

417 We will also make use of the classes of formulas $OR \circ AC_{d-1}^0$ and $AND \circ AC_{d-1}^0$, defined
418 as the subclasses of AC_d^0 formulas with a top OR gate and a top AND gate respectively. For
419 $\mathcal{C} \in \{OR \circ AC_{d-1}^0, AND \circ AC_{d-1}^0\}$, we define We define $CC_{\mathcal{C}}$ and (\mathcal{C}) -MCSP analogous to $CC_{AC_d^0}$
420 and (AC_d^0) -MCSP.

421 We also require the following elementary lemmas regarding AC_d^0 formulas.

422 ► **Lemma 2.2.** *Let f be a Boolean function. Then $CC_{AC_d^0}(f) = CC_{AC_d^0}(\neg f)$.*

423 **Proof.** One can use DeMorgan's laws to turn any AC_d^0 formula for f of size s into an AC_d^0
424 formula for $\neg f$ of size s . ◀

425 We note that our specific notion of AC_d^0 formula size is crucial for the next lemma.

426 ► **Lemma 2.3** (Direct product theorem for formulas). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $g :$
427 $\{0, 1\}^m \rightarrow \{0, 1\}$ be Boolean functions that are both not the constant zero function. Define
428 $h : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}$ by $h(x, y) = f(x) \wedge g(y)$. Then*

$$429 CC_{AND \circ AC_{d-1}^0}(h) = CC_{AND \circ AC_{d-1}^0}(f) + CC_{AND \circ AC_{d-1}^0}(g), \text{ and}$$

$$430 CC_{OR \circ AC_{d-1}^0}(h) \geq CC_{OR \circ AC_{d-1}^0}(f) + CC_{OR \circ AC_{d-1}^0}(g).$$

432 **Proof.** It is easy to see that

$$433 CC_{AND \circ AC_{d-1}^0}(h) \leq CC_{AND \circ AC_{d-1}^0}(f) + CC_{AND \circ AC_{d-1}^0}(g).$$

434 On the other hand, since f is not the constant 0 function, it has a 1-valued input x_0 .
435 Then, $h(x_0, y)$ computes $g(y)$. Thus, if ϕ is an $OR \circ AC_{d-1}^0$ formula for h , then ϕ has at least

436 $CC_{AND \circ AC_{d-1}^0}(g)$ y -input leaves. A similar argument shows that ϕ has at least $CC_{AND \circ AC_{d-1}^0}(f)$
 437 x -input leaves. Hence

$$438 \quad CC_{AND \circ AC_{d-1}^0}(h) \geq CC_{AND \circ AC_{d-1}^0}(f) + CC_{AND \circ AC_{d-1}^0}(g).$$

439 A similar argument shows that

$$440 \quad CC_{OR \circ AC_{d-1}^0}(h) \geq CC_{OR \circ AC_{d-1}^0}(f) + CC_{OR \circ AC_{d-1}^0}(g).$$

441 ◀

442 Oracle Circuits and Oracle MCSP: MOCSP

443 An oracle circuit C is made up of NOT gates, fan-in two AND and OR gates, and oracle
 444 gates g_1, \dots, g_t with fan-in i_1, \dots, i_t respectively for some integers $t, i_1, \dots, i_t \geq 1$. When
 445 given functions $f_1 : \{0, 1\}^{i_1} \rightarrow \{0, 1\}, \dots, f_t : \{0, 1\}^{i_t} \rightarrow \{0, 1\}$, we let $C^{f_1, \dots, f_t}(x)$ be the
 446 value obtained when evaluating C on input x by using the f_1, \dots, f_t functions as the outputs
 447 of the g_1, \dots, g_t gates respectively.

448 We define the size of an oracle circuit C , denoted $|C|$, to be the sum of the number of
 449 OR gates, the number of AND gates, and the number of oracle gates in C .

450 For Boolean functions f, f_1, \dots, f_t , we let $CC^{f_1, \dots, f_t}(f)$ be the size of the smallest oracle
 451 circuit that computes f when given access to f_1, \dots, f_t . Analogous to MCSP, we define the
 452 following.

453 **► Definition 2.4** (The Minimum Oracle Circuit Size Problem). *The Minimum Oracle Circuit*
 454 *Size Problem, denoted MOCSP, takes as input a truth table T , a threshold $s \in \mathbb{N}$, and*
 455 *oracle truth tables T_1, \dots, T_t and outputs whether $CC^{T_1, \dots, T_t}(T) \leq s$. The output of MOCSP*
 456 *on such an input is written as $MOCSP(T, s; T_1, \dots, T_t)$.*

457 r -Bounded Set Cover

458 We will make use of the following well known NP-complete problem.

459 **► Definition 2.5** (r -Bounded Set Cover). *r -Bounded Set Cover, denoted SetCover_r , is the*
 460 *problem that takes as input a tuple (n, c, S_1, \dots, S_t) , where $n \in \mathbb{N}$ is a universe size, $c \in \mathbb{N}$ is*
 461 *a proposed cover size $1 \leq c \leq n$, and $S_1, \dots, S_t \subseteq [n]$ are sets of cardinality at most r whose*
 462 *union is $[n]$, and outputs whether $c \geq \ell$ where ℓ is the optimal cover size given by*

$$463 \quad \ell = \min\{|I| : I \subseteq [t] \text{ and } \cup_{i \in I} S_i = [n]\}.$$

464 We will also make use of the following restricted version of set cover. Let $\text{SetCover}_{r,n,t}$
 465 denote r -bounded set cover on t subsets of $[n]$. We encode inputs to $\text{SetCover}_{r,n,t}$ as the
 466 tuple (c, S_1, \dots, S_t) (with n implicit) where c is represented in binary and the set S_i , for each
 467 $i \in [t]$, is represented by a bit string of length $r \lceil \log(n+1) \rceil$ that is a concatenated list of the
 468 elements of S_i in binary, padded with zeroes if $|S_i| < r$ (note that zero is not an element of
 469 $[n]$, so padding with zeroes is not ambiguous).

470 We will use that SetCover_r is NP-hard even to approximate to a roughly $\ln r$ factor.

471 **► Theorem 2.6** (Feige [10] and Trevisan [31]). *Let r be a sufficiently large constant, and let*
 472 *L be a language. If for every instance $x = (n, c, S_1, \dots, S_t)$ of SetCover_r , we have that both*
 473 *that*

474 **■** $c \geq \ell$ implies $x \in L$, and

475 **■** $c \leq \ell / (\ln r - O(\ln \ln r))$ implies $x \notin L$,

476 *where ℓ is the optimal cover size, then L is NP-hard under polynomial-time many-one*
 477 *reductions.*

478 **3** MAJORITY $\leq_{tt}^{\text{AC}^0}$ (AC_d⁰)-MCSP

479 Let $d \geq 2$. Our goal in this section is to prove the following theorem.

480 ▶ **Theorem 3.1.** MAJORITY $\leq_{tt}^{\text{AC}^0}$ (AC_d⁰)-MCSP.

481 We will do this by showing that for all sufficiently large $n \in \mathbb{N}$, there exists an AC⁰ truth
482 table oracle circuit that computes MAJORITY on n -bits when given access to (AC_d⁰)-MCSP.

483 Fix some n , and let q be the least power of two such that $n \leq \sqrt{q}/2$. We will split our
484 analysis into cases depending on whether there exists an $m \in \{q^{10}, q^{50}\}$ such that CC_{AC_d⁰} is
485 ($m^{.25}$)-Lipchitz on inputs of length m .

486 **3.1 Case 1: Lipchitzness Holds**

487 If Lipchitzness holds, then the desired (AC_d⁰)-MCSP oracle circuit C exists for computing
488 MAJORITY on n -inputs by the work of Golovnev *et al.* [11]. At a high-level, C works
489 by using the input string to sample a random variable whose circuit complexity spikes (in
490 expectation) depending on the weight of the input and using Lipchitzness to show that this
491 spike happens with such high probability that it can be derandomized using non-uniformity.

492 For completeness, we give a self-contained proof of this case in Appendix A.

493 **3.2 Case 2: Lipchitzness fails**

494 Assume that for all $m \in \{q^{10}, q^{50}\}$, CC_{AC_d⁰} is not ($m^{.25}$)-Lipchitz on inputs of length m . Thus,
495 for all $m \in \{q^{10}, q^{50}\}$ there exist functions $f^m, h^m : \{0, 1\}^{\log m} \rightarrow \{0, 1\}$ that differ only on a
496 single input $z^m \in \{0, 1\}^{\log m}$ such that $\text{CC}_{\text{AC}_d^0}(h^m) - \text{CC}_{\text{AC}_d^0}(f^m) > m^{.25}$.

497 We assume, without loss of generality, that for all $m \in \{q^{10}, q^{50}\}$, $f^m(z^m) = 0$ and
498 $h^m(z^m) = 1$. (If this is not the case, then replace f^m and h^m by $\neg f^m$ and $\neg h^m$ respectively
499 and apply Lemma 2.2.)

500 First, we show that the failure of Lipchitzness implies the existence of functions that
501 are much easier to compute by formulas with an AND gate on top. For $m \in \{q^{10}, q^{50}\}$ let
502 $\mathbb{1}_{z^m} : \{0, 1\}^{\log m} \rightarrow \{0, 1\}$ denote the indicator function that accepts just the string z^m .

503 ▶ **Proposition 3.2.** Let $m \in \{q^{10}, q^{50}\}$. For sufficiently large n , $\text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(f^m) \geq$
504 $\text{CC}_{\text{AC}_d^0}(f^m) + m^{.24}$, and so any optimal AC_d⁰ formula for f^m has an AND gate on top.

505 **Proof.** Suppose ϕ is an $\text{OR} \circ \text{AC}_{d-1}^0$ formula computing f^m , that is, f^m is computed by
506 $\phi = \phi_1 \vee \dots \vee \phi_t$ for some AC_{d-1}⁰ formulas ϕ_1, \dots, ϕ_t . Then

507
$$\mathbb{1}_{z^m} \vee \phi_1 \vee \dots \vee \phi_t$$

508 computes h^m . Since $\mathbb{1}_{z^m}$ can be computed by a single AND gate of formula size $\log m$,
509 this shows that $\text{CC}_{\text{AC}_d^0}(h^m) \leq |\phi| + \log m$. Combining this with the fact that $\text{CC}_{\text{AC}_d^0}(h^m) -$
510 $\text{CC}_{\text{AC}_d^0}(f^m) \geq m^{.25}$ gives the desired result. ◀

511 At this point, we will need to refer to both q^{10} and q^{50} individually, so for convenience
512 let $u = q^{10}$ and $v = q^{50}$.

513 Let ϕ^u be an optimal AC_d⁰ formula for f^u . By Proposition 3.2, for sufficiently large n ,
514 we know that $\phi^u = \phi_1^u \wedge \dots \wedge \phi_t^u$ for some AC_{d-1}⁰ formulas $\phi_1^u, \dots, \phi_t^u$. Moreover, we can
515 assume, without loss of generality, that the top gate of ϕ_i^u is OR for all $i \in [t]$. (If some ϕ_i^u
516 has an AND gate on top, then this AND can be carried out by the AND gate on top of ϕ^u
517 without increasing the size of the formula.)

518 Our next Proposition shows that ϕ^u has high top fan-in.

519 ► **Proposition 3.3.** *For sufficiently large n ,*

520
$$t \geq u^{24}.$$

521 **Proof.** We divide into cases depending on d .

522 **Case 1:** $d \geq 3$. Realize that

523
$$(\phi_1^u \vee \mathbb{1}_{z^u}) \wedge \cdots \wedge (\phi_t^u \vee \mathbb{1}_{z^u})$$

524 computes h^u . Since $\mathbb{1}_{z^u}$ can be computed by a single AND gate of formula size $\log u$ and
 525 the top gate of each ϕ_i^u is an OR gate and $d \geq 3$, this yields a depth- d formula for h^u of size
 526 $CC_{AC_d^0}(f^u) + t \log u$. Since $CC_{AC_d^0}(h^u) - CC_{AC_d^0}(f^u) \geq u^{25}$, the desired bound on t follows.

527 **Case 2:** $d = 2$. Let $\mathbb{1}_{z^u, j} : \{0, 1\}^{\log u} \rightarrow \{0, 1\}$ be the function that accepts a string x if
 528 and only if the j th bit of x equals the j th bit of z^u . Observe that, since $\bigwedge_{i \in [t]} \phi_i^u$ computes
 529 f^u , we have that

530
$$\bigwedge_{i \in [t]} \bigwedge_{j \in [\log u]} (\phi_i^u \vee \mathbb{1}_{z^u, j})$$

531 computes h^u . Since $\mathbb{1}_{z^u, j}$ is computed by a single input leaf and ϕ_i^u has an OR gate on top,
 532 this yields a depth-2 formula for h^u of size $(|\phi^u| + 1) \log u$. Since ϕ^u is an optimal CNF,
 533 each clause ϕ_i^u of ϕ^u is the OR of at most $\log u$ input leaves. In other words, $|\phi_i^u| \leq \log u$.
 534 Therefore, we have that

535
$$CC_{AC_d^0}(h^u) \leq (|\phi^u| + 1) \log u \leq \left(\sum_{i \in [t]} |\phi_i^u| + 1 \right) \log u \leq (t \log u + 1) \log u = t \log^2 u + \log u.$$

536 On the other hand, we know by assumption that

537
$$CC_{AC_d^0}(h^u) > CC_{AC_d^0}(f^u) + u^{25} \geq u^{25}.$$

538 Combining these two inequalities gives us the desired bound on t .

539 ◀

540 Let p be smallest prime greater than n . (Note that $p \leq 2n$ by Bertrand's postulate,
 541 also known as Chebyshev's theorem. See [12] for a proof.) We say that an integer j is
 542 (k, r) -good for integers $k \geq 0$ and $1 \leq r \leq p - 1$ if p^k divides j and $j/p^k \equiv r \pmod{p}$. In other
 543 words, an integer j is (k, r) -good for $k \geq 0$ and $r \in [p - 1]$ if the k th largest entry of the
 544 base- p representation of the integer j equals r and all previous entries equal zero. From this
 545 "base- p " perspective, it is clear that all positive integers j are (k, r) -good for some $k \geq 0$ and
 546 $r \in [p - 1]$.

547 We show that, for some k and r , a large subset of the $|\phi_i^u|$ are (k, r) -good.

548 ► **Proposition 3.4.** *For all sufficiently large n , there exist integers $k \geq 0$ and $1 \leq r \leq p - 1$
 549 and a set $S \subseteq [t]$ of cardinality n such that, for all $i \in S$, the integer $|\phi_i^u|$ is (k, r) -good.*

550 **Proof.** We do this by an averaging argument. First, we show that each $|\phi_i^u|$ is (k, r) -good
 551 for a k not too large.

552 ► **Claim 3.5.** *For all $i \in [t]$, $|\phi_i^u|$ is (k, r) -good for some $0 \leq k \leq \log(u \log u) + 1$ and some
 553 $r \in [p - 1]$.*

554 Proof. Fix some $i \in [t]$. $|\phi_i^u|$ is a positive integer, so $|\phi_i^u|$ is (k, r) -good for some $k \geq 0$ and
 555 some $r \in [p-1]$. We still need to upper bound this k . Note that the size of $|\phi_i^u|$ is at most
 556 by $u \log u$ since ϕ^u is optimal for f^u and f^u can be computed by a DNF of size $u \log u$. Thus,
 557 for p^k to divide $|\phi_i^u|$, we must have that $k \leq \log_p(u \log u) + 1 \leq \log(u \log u) + 1$. \triangleleft

558 Since for all $i \in [t]$ we have shown that $|\phi_i^u|$ is (k, r) good for some $0 \leq k \leq \log(u \log u) + 1$
 559 and some $r \in [p-1]$, a standard averaging argument implies that there exists a set $S \subseteq [t]$
 560 of cardinality at least

$$561 \quad \frac{t}{(\log(u \log u) + 1)(p-1)}$$

562 such that for all $i \in S$, $|\phi_i^u|$ is (k, r) -good for some fixed $k \geq 0$ and $1 \leq r \leq p-1$. For
 563 sufficiently large n , we have that

$$564 \quad \frac{t}{(\log(u \log u) + 1)(p-1)} \geq \frac{u^{24}}{4n \log u} \geq n$$

565 using that $u = q^{10} \geq n^{10}$. We then can truncate S so that it has only n elements as
 566 desired. \blacktriangleleft

567 Assume that n is large enough that all the sufficiently large hypotheses in Propositions
 568 3.2, 3.3, and 3.4 apply. For convenience, relabel ϕ_1, \dots, ϕ_t so that the set S guaranteed by
 569 Proposition 3.4 is just $S = [n]$. Fix $k \geq 0$ and $r \in [p-1]$ to be the values such that for all
 570 $i \in S = [n]$, $|\phi_i^u|$ is (k, r) -good.

571 Introducing notation, for a set $A \subseteq [n]$, let f_A^u be the function computed by $\bigwedge_{i \in A} \phi_i^u$.

572 **► Lemma 3.6.** *Let $A \subseteq [n]$. Then $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_A^u) = \sum_{i \in A} |\phi_i^u|$.*

573 **Proof.** By construction, we have that $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_A^u) \leq \sum_{i \in A} |\phi_i^u|$. Suppose for contradic-
 574 tion that $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_A^u) < \sum_{i \in A} |\phi_i^u|$.

575 Let $\theta_1 \wedge \dots \wedge \theta_\ell$ be a minimum-sized $(\text{AND} \circ \text{AC}_{d-1}^0)$ formula for f_A^u . By assumption, we
 576 have that $\sum_{j=1}^{\ell} |\theta_j| < \sum_{i \in A} |\phi_i^u|$. We can thus replace the $\bigwedge_{i \in A} \phi_i^u$ in the optimal formula
 577 for f^u with $\theta_1 \wedge \dots \wedge \theta_\ell$ and get a smaller formula. In more detail, we have that

$$578 \quad f^u = f_A^u \wedge \left(\bigwedge_{i \in [t] \setminus A} \phi_i^u \right) = (\theta_1 \wedge \dots \wedge \theta_\ell) \wedge \left(\bigwedge_{i \in [t] \setminus A} \phi_i^u \right)$$

579 which is a formula of size

$$580 \quad \sum_{j=1}^{\ell} |\theta_j| + \sum_{i \in [t] \setminus A} |\phi_i^u| < \sum_{i \in A} |\phi_i^u| + \sum_{i \in [t] \setminus A} |\phi_i^u| = \sum_{i=1}^t |\phi_i^u| = |\phi^u|$$

581 which contradicts the optimality of ϕ^u for f^u . \blacktriangleleft

582 For a string $x \in \{0, 1\}^n$, let f_x^u be shorthand for $f_{A_x}^u$ where $A_x \subseteq [n]$ is the set of indices
 583 where x is one.

584 **► Proposition 3.7.** *Let $x \in \{0, 1\}^n$. Then x has weight w if and only if the integer
 585 $\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)$ is (k, rw) -good.*

586 **Proof.** By Lemma 3.6 and the fact that $|\phi_i^u|$ is (k, r) -good for all $i \in [n]$, we have that

$$587 \quad \frac{\text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)}{p^k} = \frac{\sum_{i \in A_x} |\phi_i^u|}{p^k} \equiv w \cdot r \pmod{p}$$

588 where $A_x \subseteq [n]$ are bits of x that are ones. The “only if” part of the statement is guaranteed
 589 by the fact that $1 \leq r \leq p-1$ has a multiplicative inverse modulo p since p is prime. \blacktriangleleft

590 ► **Theorem 3.8.** *Assume n is sufficiently large. Then there is a depth-8 AC⁰ truth table*
 591 *oracle circuit C with $O(n^{250})$ wires such that $C^{(\text{AC}_d^0)\text{-MCSP}}$ computes MAJORITY on n -bits.*

592 **Proof.** It suffices to show that for every $w \in [n]$, there exists a depth-7 AC⁰ oracle circuit C_w
 593 with $O(n^{249})$ wires such that $C_w^{(\text{AC}_d^0)\text{-MCSP}}(x) = 1 \iff \text{wt}(x) = w$. Then $\text{MAJORITY}(x) =$
 594 $\bigvee_{w \geq n/2} C_w(x)$.

595 Fix some $w \in [n]$. The circuit C_w works as follows. On input $x \in \{0, 1\}^n$, first check
 596 if x is the all zeroes string. If so, then reject. Otherwise, compute the truth table of the
 597 direct product function $g_x : \{0, 1\}^{\log u} \times \{0, 1\}^{\log v} \rightarrow \{0, 1\}$ given by $g_x(y, z) = f_x^u(y) \wedge f^v(z)$.
 598 Compute $s = \text{CC}_{\text{AC}_d^0}(g_x)$ in binary using oracle access to $(\text{AC}_d^0)\text{-MCSP}$. Finally accept if the
 599 integer s has the property that $s - \text{CC}_{\text{AC}_d^0}(f^v)$ is (k, rw) -good. Reject otherwise.

600 We now verify this yields a (non-uniform) AC⁰ truth table oracle circuit. We can check if
 601 x is the all zeroes string with a single OR gate. This requires one level of depth and $O(n)$
 602 wires. Next, realize the j th bit in the truth table of g_x is either zero for all x or equal to

$$603 \quad f_x^u(j) = \bigvee_{i \in [n]: \phi_i^u(j)=1} x_i$$

604 where x_i denotes the i th bit of x . Thus, using non-uniformity, we can compute the truth table
 605 of g_x with $O(nuv) = O(nq^{60}) = O(n^{121})$ wires and depth-one. Next, we can compute $s =$
 606 $\text{CC}_{\text{AC}_d^0}(g_x)$ in binary with $O(uv \log(uv))$ calls to $(\text{AC}_d^0)\text{-MCSP}$ using the fact that $\text{CC}_{\text{AC}_d^0}(g_x) \leq$
 607 $uv \log(uv)$ by the DNF bound and the fact that

$$608 \quad \text{CC}_{\text{AC}_d^0}(g_x) = s \iff (\text{AC}_d^0)\text{-MCSP}(g_x, s) = 1 \text{ and } (\text{AC}_d^0)\text{-MCSP}(g_x, s - 1) = 0.$$

609 This takes at most $\tilde{O}((uv)^2) = O(n^{241})$ wires, an additional three layers of depth, and
 610 $2uv \log(uv)$ oracle calls that all do not depend on each other. Finally, the DNF upper bound
 611 guarantees that $\text{CC}_{\text{AC}_d^0}(g_x) \leq uv \log(uv) \leq n^{61}$, so the length of the integer $s = \text{CC}_{\text{AC}_d^0}(g_x)$ in
 612 binary is at most $61 \log n$. Therefore we can check if s has the property that $s - \text{CC}_{\text{AC}_d^0}(f^v)$
 613 is (k, rw) -good using a DNF with at most n^{62} wires and at most an additional two layers of
 614 depth. Combining all this yields a AC⁰ circuit of depth-7 with at most $O(n^{241})$ wires and no
 615 directed path between oracle gates.

616 Next, we argue for correctness. Clearly, C_w rejects the all zero string, so assume $x \neq 0^n$.
 617 By Proposition 3.7, it suffices to show that, for $s = \text{CC}_{\text{AC}_d^0}(g_x)$,

$$618 \quad s - \text{CC}_{\text{AC}_d^0}(f^v) = \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u).$$

619 We confirm that neither f_x^u nor f^v is the constant zero function, so that we can use the
 620 direct product theorems in Lemma 2.3.

621 ► **Claim 3.9.** Neither f_x^u nor f^v is the constant zero function.

622 **Proof.** If f^v were the constant zero function, then $\text{CC}_{\text{AC}_d^0}(h^v) \leq \log v$ by DNF computation
 623 which contradicts that

$$624 \quad \text{CC}_{\text{AC}_d^0}(h^v) - \text{CC}_{\text{AC}_d^0}(f^v) \geq v^{.25}.$$

625 Next, let $i \in [n]$ be a bit of x that is not zero. (Recall, we assumed that $x \neq 0^n$.) Then
 626 f_x^u has accepts every input that that ϕ_i^u accepts. For contradiction, suppose that ϕ_i^u had no
 627 ones. Then we can remove ϕ_i^u from the optimal formula $\phi^u = \phi_1^u \wedge \dots \wedge \phi_i^u$ for f^u and get a
 628 smaller formula for f^u which contradicts the optimality of ϕ^u . ◀

Next we show that the optimal AC_d^0 formula for g_x has an AND gate on top.

▷ **Claim 3.10.** $\text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(g_x) > \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(g_x)$. Consequently, $\text{CC}_{\text{AC}_d^0}(g_x) = \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(g_x)$.

Proof. Let $\Delta = \text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(g_x) - \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(g_x)$. We need to show $\Delta > 0$.

$$\begin{aligned}
\Delta &\geq \text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(f^v) - \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v) + \text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(f_x^u) && \text{(by Lemma 2.3)} \\
&\quad - \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u) \\
&\geq (v)^{24} + \text{CC}_{\text{OR} \circ \text{AC}_{d-1}^0}(f_x^u) - \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u) && \text{(by Proposition 3.2)} \\
&\geq (v)^{24} - u \log u && \text{(by DNF bound on } f_x^u) \\
&\geq n^{50 \cdot 24} - n^{10} \log(n^{10}) && \text{(by definition of } u \text{ and } v) \\
&> 0 && \text{(for sufficiently large } n)
\end{aligned}$$

◁

Using the claim we have that

$$\begin{aligned}
s - \text{CC}_{\text{AC}_d^0}(f^v) &= \text{CC}_{\text{AC}_d^0}(g_x) - \text{CC}_{\text{AC}_d^0}(f^v) && \text{(by definition)} \\
&= \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(g_x) - \text{CC}_{\text{AC}_d^0}(f^v) && \text{(by Claim 3.10)} \\
&= \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u) + \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v) - \text{CC}_{\text{AC}_d^0}(f^v) && \text{(by Lemma 2.3)} \\
&= \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u) + \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v) - \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f^v) && \text{(by Proposition 3.2)} \\
&= \text{CC}_{\text{AND} \circ \text{AC}_{d-1}^0}(f_x^u)
\end{aligned}$$

as desired. ◀

► **Remark 3.11.** We remark that the only time we use the failure of Lipschitzness in Case 2 is to show the existence of functions like f^u with high top fan-in and functions like f^v with a large difference between top AND gate and top OR gate complexity. Using known PARITY lower bounds and depth-hierarchy theorems for AC^0 circuits, we can unconditionally prove the existence of f^u and f^v respectively but with slightly worse parameters that would yield a quasi-polynomial reduction (at least in the $d \geq 3$ case) rather than the polynomial reduction we present.

4 On the NP-hardness of MOCSP

First, we introduce some useful notation and definitions. For a truth table T of length m and a set $P \subseteq [m]$, let T_P be the truth table of length m where the j th bit of T_P equals

$$\begin{cases} \text{the } j\text{th bit of } T & , \text{ if } j \in P \\ 0 & , \text{ otherwise.} \end{cases}$$

Next, we say a truth table T of length m is (s)-irritable on a partition $\mathcal{P} = (P_1, \dots, P_n)$ of $[m]$ if for all $i \in [n]$

$$\text{CC}^{T_{P_1}, \dots, T_{P_{i-1}}, T_{P_{i+1}}, \dots, T_{P_n}}(T) > s.$$

Finally, for a partition $\mathcal{P} = (P_1, \dots, P_n)$ of $[m]$ and any set $S \subseteq [n]$, we define the \mathcal{P} -lift of S , denoted $S^{\mathcal{P}}$, to be the subset of $[m]$ given by

$$S^{\mathcal{P}} = \bigcup_{i \in S} P_i.$$

Our first theorem shows that one can use an irritable truth table and MOCSP to approximate r -bounded set cover.

► **Theorem 4.1.** *Let $S_1, \dots, S_t \subseteq [n]$ be sets of cardinality at most r that cover $[n]$. Let T be a truth table of length m , and let $\mathcal{P} = (P_1, \dots, P_n)$ be a partition of $[m]$. Then*

■ $\text{CC}^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}}(T) \leq 2\ell$, and

■ $\text{CC}^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}}(T) > \ell/2$ if T is (rn) -irritable on \mathcal{P}

where ℓ is size of the optimal cover of $[n]$ by S_1, \dots, S_t .

Proof. We split this proof into two claims to make clear that our two “without loss of generality” assumptions do not conflict with each other.

▷ **Claim 4.2.** $\text{CC}^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}}(T) \leq 2\ell$

Proof. Without loss of generality, assume that the optimal cover size ℓ is witnessed by $S_1 \cup \dots \cup S_\ell = [n]$. Then, by construction, the function computed by oracle circuit $T_{S_1^{\mathcal{P}}} \vee \dots \vee T_{S_\ell^{\mathcal{P}}}$ of size $2\ell - 1$ is T . In more detail,

$$\bigvee_{i \in [\ell]} T_{S_i^{\mathcal{P}}} = \bigvee_{i \in [\ell]} \bigvee_{j \in S_i} T_{P_j} = \bigvee_{j \in S_1 \cup \dots \cup S_\ell} T_{P_j} = \bigvee_{j \in [n]} T_{P_j} = T.$$

Therefore $\text{CC}^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}}(T) \leq 2\ell$. ◁

▷ **Claim 4.3.** If T is (rn) -irritable on \mathcal{P} , then $\text{CC}^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}} > \ell/2$.

Proof. For contradiction, suppose there is an oracle circuit D with at most $\ell/2$ gates such that $D^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}}$ computes T . Since D has at most $q \leq \ell/2$ unique oracle gates, assume, without loss of generality, that $D^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_q^{\mathcal{P}}}}$ computes T .

Recall, by the definition of $T_{S_i^{\mathcal{P}}}$, we have that $T_{S_i^{\mathcal{P}}} = \bigvee_{j \in S_i} T_{P_j}$. Note that $\bigvee_{j \in S_i} T_{P_j}$ is an oracle circuit of size at most $2r$ since $|S_i| \leq r$. Thus, by replacing each $T_{S_i^{\mathcal{P}}}$ oracle gate in D with the oracle circuit $\bigvee_{j \in S_i} T_{P_j}$, we can transform D into an oracle circuit E of size at most $2r|C| \leq r\ell$ that computes T when given access to the oracles in the set $O = \{T_{P_j} : j \in S_1 \cup \dots \cup S_q\}$. However, since the optimal cover of n is of size ℓ and $q \leq \ell/2 < \ell$, it follows that $|S_1 \cup \dots \cup S_q| < n$ and hence $|O| < n$. Thus, there is an element $i^* \in [n]$ such that $T_{P_{i^*}} \notin O$. Therefore, the circuit E witnesses that

$$\text{CC}^{T_{P_1}, \dots, T_{P_{i^*-1}}, T_{P_{i^*+1}}, \dots, T_{P_n}}(T) \leq r\ell \leq rn$$

which contradicts that T is (rn) -irritable on \mathcal{P} . ◁

Moreover, given a truth table T and the partition \mathcal{P} on which T is sufficiently irritable, we show it is easy to build a constant-depth circuit that approximates r -bounded set cover using MOCSP.

► **Theorem 4.4.** *Let $r \in \mathbb{N}$. There exists a polynomial-time algorithm A such that if $\mathcal{P} = (P_1, \dots, P_n)$ is a partition of $[m]$, and T is a truth table of length m , then*

1. $A(0^n, 0^t, T, \mathcal{P})$ outputs a depth-2 AC⁰ circuit $C_{n,t,T,\mathcal{P}}$ with $O(mnrt)$ wires,

2. $\text{MOCSP} \circ C_{n,t,T,\mathcal{P}}$ accepts all YES instances of $\text{SetCover}_{r,n,t}$, and

3. $\text{MOCSP} \circ C_{n,t,T,\mathcal{P}}$ computes a 4-approximation to $\text{SetCover}_{r,n,t}$ if T is (rn) -irritable on \mathcal{P} .

703 **Proof.** First, we show that there is a polynomial-time algorithm A that outputs a small
 704 depth-2 circuit computing a specific function. Then we show that this specific function is
 705 helpful in computing $\text{SetCover}_{r,n,t}$.

706 \triangleright **Claim 4.5.** There is a polynomial-time algorithm A such that $A(0^n, 0^t, T, \mathcal{P})$ outputs a
 707 depth-2 AC^0 circuit $C_{n,t,T,\mathcal{P}}$ with $O(mnrt)$ wires satisfying

$$708 \quad C_{n,t,T,\mathcal{P}}(c, S_1, \dots, S_t) = (T, 2c; T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}})$$

709 for any instance (c, S_1, \dots, S_t) of $\text{SetCover}_{r,n,t}$.

710 **Proof.** On input $(0^n, 0^t, T, \mathcal{P})$, A builds the circuit $C_{n,t,T,\mathcal{P}}$ as follows. First, A will hardwire
 711 $C_{n,t,T,\mathcal{P}}$ to output T . This requires $O(m)$ wires and depth one. Next, A adds circuitry to
 712 $C_{n,t,T,\mathcal{P}}$ that outputs $2c$ by adding an extra zero to the binary expansion of c . This uses
 713 $O(\log n)$ wires and depth one.

714 Finally, A adds circuitry to $C_{n,t,T,\mathcal{P}}$ that outputs $T_{S_i^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}$ as follows. Observe that
 715 for any $i \in [t]$ and $j \in [m]$, the j th bit of $T_{S_i^{\mathcal{P}}}$ is one if and only if S_i contains the unique
 716 element $k_j \in [n]$ such that $j \in P_{k_j}$. Thus, since A has access to \mathcal{P} , A can calculate k_j for all
 717 $j \in [m]$ and then add circuitry to $C_{n,t,T,\mathcal{P}}$ that calculates the j th bit of $T_{S_i^{\mathcal{P}}}$ by ORing over
 718 all the elements of S_i and using an AND to check if any one of those elements is k_j . This
 719 requires $O(mnrt)$ wires and depth-two.

720 Therefore, $C_{n,t,T,\mathcal{P}}$ is a depth-2 AC^0 circuit with $O(mnrt)$ wires as desired. Moreover, it
 721 is clear from this description that A runs in polynomial-time. \triangleleft

722 It remains to show that the algorithm A given in Claim 4.5 satisfies (2) and (3). Let
 723 (c, S_1, \dots, S_t) be an instance of $\text{SetCover}_{r,n,t}$. Let ℓ be the minimum size of any cover of $[n]$
 724 by S_1, \dots, S_t .

725 First, we show (2) holds. Suppose $c \geq \ell$. Then Theorem 4.1 implies that

$$726 \quad CC^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}}(T) \leq 2\ell \leq 2c$$

727 SO

$$728 \quad \text{MOCSP}(C_{n,t,T,\mathcal{P}}(c, S_1, \dots, S_t)) = \text{MOCSP}(T, 2c; T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}) = \text{YES}$$

729 as desired.

730 Finally, we show (3) holds. Suppose $c < \ell/4$ and T is (rn) -irritable. Then Theorem 4.1
 731 implies that

$$732 \quad CC^{T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}}(T) > \ell/2 \geq 2c$$

733 SO

$$734 \quad \text{MOCSP}(C_{n,t,T,\mathcal{P}}(c, S_1, \dots, S_t)) = \text{MOCSP}(T, 2c; T_{S_1^{\mathcal{P}}}, \dots, T_{S_t^{\mathcal{P}}}) = \text{NO}.$$

735 Hence (3) holds. \blacktriangleleft

736 Of course to make use of Theorem 4.4, we need to actually find truth tables T and
 737 partitions \mathcal{P} on which T is sufficiently irritable. Fortunately, such T and \mathcal{P} are abundant.
 738 We show that, with high probability, any choice of \mathcal{P} and a random choice of a truth table T
 739 suffices.

740 ► **Lemma 4.6.** *Let $n, r \in \mathbb{N}$. Let m be the least power of two greater than n^3 . Let $\mathcal{P} =$
 741 (P_1, \dots, P_n) be any partition of $[m]$ such that $|P_i| \geq m/n - 1$ for all $i \in [n]$. Pick a truth
 742 table $T \in \{0, 1\}^m$ uniformly at random. Then T is (rn) -irritable on \mathcal{P} except with probability
 743 $2^{-\Omega(n^2)}$.*

744 **Proof.** We prove this by bounding the probability that, for some fixed i^* and some fixed
 745 oracle circuit C ,

$$746 \quad C^{T_{P_1}, \dots, T_{P_{i^*-1}}, T_{P_{i^*+1}}, \dots, T_{P_n}} \text{ computes } T$$

747 and then union bounding over all i^* and all oracle circuits of size at most rn .

748 Realize that the function computed by $C^{T_{P_1}, \dots, T_{P_{i^*-1}}, T_{P_{i^*+1}}, \dots, T_{P_n}}$ does not depend on
 749 any of the bits in T that lie in P_{i^*} . Therefore, since $|P_{i^*}| \geq m/n - 1$, this means that the
 750 probability that $C^{T_{P_1}, \dots, T_{P_{i^*-1}}, T_{P_{i^*+1}}, \dots, T_{P_n}}$ computes T is at most $2^{-m/n+1} \leq 2^{-n^2+1}$.

751 Now, we union bound. Clearly, there are at most n choices of i^* . Next, we need to count
 752 the number of C of size at most rn . For concision, let $s = rn$. We bound the number of
 753 oracle circuits of size s , allowing for identity gates to catch circuits of smaller size. For each
 754 of the s gates, there are $4 + n$ gate types to chose from (AND, OR, NOT, identity, and the n
 755 oracle gates). Then, for each of the s gates, we have to choose the at most $\log(m)$ wires that
 756 feed into that gate and there are at most $(s + \log(m))$ choices for where each of these wires
 757 comes from. Hence, we get a bound of

$$758 \quad (4 + n)^s (s + \log(m))^{s \log(m)}$$

759 whose logarithm is

$$760 \quad s \log(4 + n) + s \log(m) \log(s + \log(m)) = \tilde{O}(rn).$$

761 Thus, probability that T fails to be (rn) -irritable on \mathcal{P} is at most

$$762 \quad n 2^{-n^2 + \tilde{O}(rn)} \leq 2^{-\Omega(n^2)}.$$

763 ◀

764 Thus, using a random choice of T gives us an NP-hardness result under RP-reductions.

765 ► **Corollary 4.7.** $\text{NP} \leq_m^{\text{RP}} \text{MOCSP}$.

766 **Proof.** Let r be sufficiently large that computing a 4-approximation to r -bounded set cover
 767 is NP-hard (such an r exists by Theorem 2.6). We will reduce giving a 4-approximation of
 768 r -bounded set cover to MOCSP.

769 The reduction R works as follows. On an instance (c, S_1, \dots, S_t) of $\text{SetCover}_{r,n,t}$, R first
 770 computes the integer m that is the least power of two greater than n^3 . Next, R computes
 771 the partition $\mathcal{P} = (P_1, \dots, P_n)$ of $[m]$ where for all $i \in [n]$

$$772 \quad P_i = \{j \in [m] : j \equiv i \pmod{n}\}.$$

773 Then, R picks a truth table T of length m uniformly at random. After that, R runs the
 774 algorithm A from Theorem 4.4 on the input $(0^n, 0^t, T, \mathcal{P})$ to obtain the circuit $C_{n,t,T,\mathcal{P}}$.
 775 Finally, R outputs $\text{MOCSP}(C_{n,t,T,\mathcal{P}}(c, S_1, \dots, S_t))$.

776 Now, we argue for correctness. Theorem 4.4 guarantees that $C_{n,t,T,\mathcal{P}}$ correctly answers all
 777 YES instances of $\text{SetCover}_{r,n,t}$, so R also correctly answers all YES instances of $\text{SetCover}_{r,n,t}$.

778 On the other hand, observe that our construction of \mathcal{P} guarantees that $|P_i| \geq m/n - 1$
 779 for all $i \in [n]$, so Lemma 4.6 implies that T is (rn) -irritable on \mathcal{P} with high probability.
 780 Therefore, Theorem 4.1 further implies that with high probability $C_{n,t,T,\mathcal{P}}$ (and hence R)
 781 computes a 4-approximation to r -bounded set cover. ◀

782 Using more queries to MOCSP, we can improve the RP reduction to a ZPP reduction by
 783 checking if the randomly chosen T is indeed (rn) -irritable on \mathcal{P} .

784 ► **Corollary 4.8.** $\text{NP} \leq_{tt}^{\text{ZPP}} \text{MOCSP}$.

785 **Proof.** Run the same reduction as in the proof of Corollary 4.7 except check whether T
 786 is (rn) -irritable on \mathcal{P} using the MOCSP oracle. This can be done at the same time the
 787 as the MOCSP oracle answers $\text{MOCSP}(C_{n,t,T,\mathcal{P}}(c, S_1, \dots, S_t))$. If T is indeed (rn) -irritable
 788 on \mathcal{P} , then we know the output given by $\text{MOCSP}(C_{n,t,T,\mathcal{P}}(c, S_1, \dots, S_t))$ is correct using
 789 Theorem 4.1. Otherwise, output “don’t know.” ◀

790 We can also use non-uniform bits to provide the reduction with a truth table T and a
 791 partition \mathcal{P} such that T is sufficiently irritable on \mathcal{P} . This yields an AC^0 many-one reduction.

792 ► **Corollary 4.9.** *MOCSP is NP-hard under (non-uniform) AC^0 many-one reductions.*

793 **Proof.** Let r be large enough that computing a 4-approximation to r -bounded set cover is
 794 NP-hard. It suffices to show that for all sufficiently large n and t , there is an AC^0 circuit C
 795 such that $\text{MOCSP} \circ C$ computes a 4-approximation to $\text{SetCover}_{r,n,t}$.

796 By Lemma 4.6 for sufficiently large n , there exists a truth table T of length $O(n^3)$ and
 797 a partition $\mathcal{P} = (P_1, \dots, P_n)$ of $[m]$ such that T is (rn) -irritable on \mathcal{P} . Thus, letting A be
 798 the algorithm from Theorem 4.4, $A(0^n, 0^t, T, \mathcal{P})$ outputs a depth-2 AC^0 circuit C such that
 799 $\text{MOCSP} \circ C$ computes a 4-approximation to $\text{SetCover}_{r,n,t}$, as desired. ◀

800 At this point, one might begin to speculate whether we can prove that MOCSP is NP-
 801 hard under deterministic polynomial-time reductions. Unfortunately, this seems difficult.
 802 This is because Murray and Williams’ [22] and Hitchcock and Pavan’s [17] result that
 803 $\text{NP} \leq_{tt}^{\text{P}} \text{MCSP} \implies \text{EXP} \neq \text{ZPP}$ also holds for MOCSP with essentially the same proof. For
 804 completeness, we give the MOCSP version of Murray and Williams’ proof in Appendix B.

805 ► **Theorem 4.10** (Essentially proved in [22] and [17]). *If $\text{NP} \leq_{tt}^{\text{P}} \text{MOCSP}$, then $\text{EXP} \neq \text{ZPP}$.*

806 Still, it seems plausible to us that MOCSP is hard for NP under Turing reductions. Indeed,
 807 Theorem 4.4 implies that, to prove a P-Turing reduction, it suffices to show that there is
 808 a polynomial time algorithm B , with oracle access to MOCSP, such that for all large n , $B(0^n)$
 809 outputs a truth table T and a partition $\mathcal{P} = (P_1, \dots, P_n)$ such that T is (rn) -irritable on \mathcal{P} .
 810 We stress that B has access to MOCSP, so B can actually check whether T is (rn) -irritable
 811 on \mathcal{P} and make adjustments accordingly.

812 ► **Conjecture 4.11.** $\text{NP} \leq_{\text{T}}^{\text{P}} \text{MOCSP}$.

813 Maybe it is even possible to prove that such a B exists if $\text{E} \not\subseteq \text{i.o-SIZE}[2^{O(n)}]$.

814 ► **Open Question 4.12.** *Can one show that $\text{E} \not\subseteq \text{i.o-SIZE}[2^{O(n)}]$ implies $\text{NP} \leq_{\text{T}}^{\text{P}} \text{MOCSP}$?*

815 ——— References ———

- 816 1 Eric Allender, Harry Buhrman, Michal Koucký, Dieter van Melkebeek, and Detlef Ronneburger.
 817 Power from random strings. *SIAM J. Comput.*, 35(6):1467–1493, 2006.
- 818 2 Eric Allender and Bireswar Das. Zero knowledge and circuit minimization. *Information and*
 819 *Computation*, 256:2–8, 2017.
- 820 3 Eric Allender and Shuichi Hirahara. New insights on the (non-)hardness of circuit minimization
 821 and related problems. In *Symposium on Mathematical Foundations of Computer Science*
 822 *(MFCS)*, pages 54:1–54:14, 2017.

- 823 4 Eric Allender, Dhiraj Holden, and Valentine Kabanets. The minimum oracle circuit size
824 problem. *Computational Complexity*, 26(2):469–496, 2017.
- 825 5 Eric Allender, Rahul Ilango, and Neekon Vafa. The non-hardness of approximating circuit
826 size. *Electronic Colloquium on Computational Complexity*, TR18-173, 2018.
- 827 6 Eric Allender, Michal Koucký, Detlef Ronneburger, and Sambuddha Roy. The pervasive reach
828 of resource-bounded kolmogorov complexity in computational complexity theory. *Journal of*
829 *Computer and System Sciences*, 77(1):14 – 40, 2011.
- 830 7 Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge
831 University Press, 2009.
- 832 8 Harry Buhrman and Leen Torenvliet. Randomness is hard. *SIAM Journal on Computing*,
833 30:200–1, 2000.
- 834 9 Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova.
835 Learning algorithms from natural proofs. In *Computational Complexity Conference (CCC)*,
836 pages 10:1–10:24, 2016.
- 837 10 Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- 838 11 Alexander Golovnev, Rahul Ilango, Russell Impagliazzo, Valentine Kabanets, Antonina Koloko-
839 lova, and Avishay Tal. AC⁰[p] lower bounds against MCSP via the coin problem. *Electronic*
840 *Colloquium on Computational Complexity*, TR19-018, 2019.
- 841 12 G.H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. Clarendon Press,
842 Oxford, England, 5th edition, 1979.
- 843 13 Shuichi Hirahara. Non-black-box worst-case to average-case reductions within NP. In *Sym-*
844 *posium on Foundations of Computer Science (FOCS)*, pages 247–258, 2018.
- 845 14 Shuichi Hirahara, Igor C. Oliveira, and Rahul Santhanam. NP-hardness of minimum circuit
846 size problem for OR-AND-MOD circuits. In *Computational Complexity Conference (CCC)*,
847 pages 5:1–5:31, 2018.
- 848 15 Shuichi Hirahara and Rahul Santhanam. On the average-case complexity of MCSP and its
849 variants. In *Computational Complexity Conference (CCC)*, pages 7:1–7:20, 2017.
- 850 16 Shuichi Hirahara and Osamu Watanabe. Limits of minimum circuit size problem as oracle. In
851 *Computational Complexity Conference (CCC)*, pages 18:1–18:20, 2016.
- 852 17 John Hitchcock and Aduri Pavan. On the NP-completeness of the minimum circuit size
853 problem. In *Conference on Foundation of Software Technology and Theoretical Computer*
854 *Science (FSTTCS)*, 2015.
- 855 18 Russell Impagliazzo, Valentine Kabanets, and Ilya Volkovich. The Power of Natural Properties
856 as Oracles. In *Computational Complexity Conference (CCC)*, pages 7:1–7:20, 2018.
- 857 19 Valentine Kabanets and Jin-Yi Cai. Circuit minimization problem. In *Proceedings of ACM*
858 *Symposium on Theory of Computing (STOC)*, pages 73–79, 2000.
- 859 20 William J. Masek. Some NP-complete set covering problems. Unpublished Manuscript, 1979.
- 860 21 Colin McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics*,
861 (Norwich, 1989), London Math. Soc. Lecture Note Ser. 141 . Cambridge Univ. Press, Cambridge,
862 (1989), pp. 48–188.
- 863 22 Cody D. Murray and R. Ryan Williams. On the (non) NP-hardness of computing circuit
864 complexity. In *Computational Complexity Conference (CCC)*, pages 365–380, 2015.
- 865 23 Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167,
866 1994.
- 867 24 Igor C. Oliveira, Ján Pich, and Rahul Santhanam. Hardness magnification near state-of-the-art
868 lower bounds. *Electronic Colloquium on Computational Complexity*, TR18-158, 2018.
- 869 25 Igor C. Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit
870 lower bounds, and pseudorandomness. In *Computational Complexity Conference (CCC)*, pages
871 18:1–18:49, 2017.
- 872 26 Igor C. Oliveira and Rahul Santhanam. Hardness magnification for natural problems. In
873 *Symposium on Foundations of Computer Science (FOCS)*, pages 65–76, 2018.

- 874 27 A. A. Razborov. Lower bounds on the size of constant-depth networks over a complete basis
875 with logical addition. *Mathematicheskie Zametki*, 41(4):598–607, 1987.
- 876 28 Ronen Shaltiel and Emanuele Viola. Hardness amplification proofs require majority. *SIAM J.*
877 *Comput.*, 39(7):3122–3154, 2010.
- 878 29 R. Smolensky. On representations by low-degree polynomials. In *FOCS*, pages 130–138, 1993.
- 879 30 Boris Trakhtenbrot. A survey of Russian approaches to perebor (brute-force searches) al-
880 gorithms. *IEEE Ann. Hist. Comput.*, 6(4):384–400, October 1984.
- 881 31 Luca Trevisan. Non-approximability results for optimization problems on bounded degree
882 instances. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages
883 453–461, 2001.
- 884 32 Stanislav Žák. A Turing machine time hierarchy. *Theoretical Computer Science*, 26(3):327–333,
885 1983.

886 **A** MAJORITY reduces to (AC_d^0) -MCSP when Lipchitzness holds

887 Our goal in this section is to find a small (AC_d^0) -MCSP-oracle circuit that computes MAJORITY
888 on n -bits for sufficiently large n . We can do this using the techniques of Golovnev *et al.* [11].
889 In order to make our proof relatively self-contained, we differ slightly from the presentation
890 in [11]. In particular, our presentation follows a method for computing MAJORITY that is
891 described in Shaltiel and Viola [28].

892 At a high-level, this procedure works by using the input string to sample a random
893 variable whose circuit complexity spikes depending on the weight of the input and then
894 using Lipchitzness to prove that this spike occurs with high enough probability that we can
895 derandomize using non-uniformity.

896 Continuing the notation from Section 3, assume that there is an $m \in \{q^{10}, q^{50}\}$ such that
897 $CC_{AC_d^0}$ is $(m^{.25})$ -Lipchitz on inputs of length m .

898 We define the random variable $T_{p,m} \in \{0, 1\}^m$ where each bit in $T_{p,m}$ is independently
899 chosen to be one with probability p and zero with probability $1 - p$.

900 ► **Lemma A.1.** $\mathbb{E}[CC_{AC_d^0}(T_{p,m})] = \tilde{O}(pm)$ if $p \geq m^{-1/3}$

901 **Proof.** By Hoeffding’s inequality we have that the probability that $T_{p,m}$ has greater than
902 k ones is at most $\exp(-2\epsilon^2 m)$. Via computation by DNF, if a truth table $T \in \{0, 1\}^m$ has
903 at most k ones, $CC_{AC_d^0}(T) = k \log m = \tilde{O}(k)$. Similarly, we have that $\max\{C(T) : T \in$
904 $\{0, 1\}^m\} = \tilde{O}(m)$. Hence, we get that

$$905 \quad \mathbb{E}[CC_{AC_d^0}(T_{p,m})] = \tilde{O}(k) + \tilde{O}(\exp(-2\epsilon^2 m)m) = \tilde{O}(pm + pm\epsilon + \exp(-2\epsilon^2 m)m).$$

906 If we set $\epsilon = \sqrt{\frac{\ln m}{2m}}$, then we have

$$907 \quad \mathbb{E}[CC_{AC_d^0}(T_{p,m})] \leq \tilde{O}(pm + p\sqrt{m} \ln m + 1) \leq \tilde{O}(pm + \sqrt{m} \ln m)$$

908 Finally, if $p \geq 1/m^{1/3}$, we have $\mathbb{E}[T_{p,m}] = \tilde{O}(pm)$ as desired ◀

909 We will make use of the following concentration inequality.

910 ► **Theorem A.2** (McDiarmid’s “bounded differences inequality” [21]). *Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$*
911 *be c -Lipchitz. Let X_1, \dots, X_n be independent random variables with values in $\{0, 1\}$. Let*
912 *$\mu = \mathbb{E}_{X_1, \dots, X_n}[f(X_1, \dots, X_n)]$. Then*

$$913 \quad \Pr[|f(X_1, \dots, X_n) - \mu| \geq \epsilon] \leq 2\exp(-\frac{\epsilon^2}{nc^2}).$$

914 For $t \in \mathbb{N}$ and $w_1 \neq w_2 \in [t]$, we say a Boolean function $f : \{0, 1\}^t \rightarrow \{0, 1\}$ computes
 915 $\text{WTDIS}_t[w_1, w_2]$ if $\text{wt}(x) = w_1$ implies $f(x) = 1$ and $\text{wt}(x) = w_2$ implies $f(x) = 0$. (WTDIS
 916 is short for weight distinguishing.)

917 ► **Theorem A.3.** *If n is sufficiently large, then for all $1 \leq b \leq \sqrt{q}/2$, there exists a (non-*
 918 *uniform) NC⁰ oracle circuit C with at most $O(n^{100})$ wires such that $C^{(\text{AC}_d^0)\text{-MCSP}}$ computes*
 919 *$\text{WTDIS}_q[w_1, w_1 + b]$ for some $w_1 \geq \sqrt{q}/2$. Moreover, C has a single gate.*

920 **Proof.** For $w \in [q]$, let $p_w = \frac{w}{2q}$. Let w_0 be the largest integer less than \sqrt{q} such that $q - w_0$
 921 is a multiple of b . (Note that $w_0 \geq \sqrt{q} - b \geq \sqrt{q}/2$.)

922 By Lemma A.1, we have that

$$923 \quad \mathbb{E}[\text{CC}_{\text{AC}_d^0}(T_{p_{w_0}, m})] = \tilde{O}\left(\frac{\sqrt{q}}{q}m\right) = \tilde{O}(m/\sqrt{q}).$$

924 On the other hand, since $p_q = 1/2$, $T_{p_q, m}$ is just a binary string of length m picked uniformly
 925 at random, so the formula size lower bounds of Shannon and Riordan imply

$$926 \quad \mathbb{E}[\text{CC}_{\text{AC}_d^0}(T_{p_q, m})] = \mathbb{E}_{x \in \{0, 1\}^m} [C(x)] = \tilde{\Omega}(m)$$

927 (note that an AC_d^0 formula of size s implies an unrestricted formula of size s). Hence, by an
 928 averaging argument there exists a $w_1 \geq w_0 \geq \sqrt{q}/2$ such that

$$929 \quad \mathbb{E}[\text{CC}_{\text{AC}_d^0}(T_{p_{w_1+b}, m})] - \mathbb{E}[\text{CC}(T_{p_{w_1}, m})] \geq \frac{\tilde{\Omega}(m) - \tilde{O}(m/\sqrt{q})}{q} = \tilde{\Omega}(m/q).$$

930 Let $t = \frac{\mathbb{E}[T_{p_{w_1+b}, m}] + \mathbb{E}[T_{p_{w_1}, m}]}{2}$. Then we have that $\mathbb{E}[T_{p_{w_1+b}, m}] - t = \tilde{\Omega}(m/q)$ and $t -$
 931 $\mathbb{E}[T_{p_{w_1}, m}] = \tilde{\Omega}(m/q)$.

932 We now outline a probabilistic oracle circuit D that we will later make into a deterministic
 933 NC⁰ circuit. D takes as input a string $x \in \{0, 1\}^n$ and takes as its random “inputs” strings
 934 $u_1, \dots, u_m \in \{0, 1\}^{\log q}$ and $v_1, \dots, v_m \in \{0, 1\}$. The reduction then computes the string
 935 $y := y_1 \dots y_m$ where y_i is zero if v_i is zero and y_i is the u_i th bit of x if v_i is one (recall, q is a
 936 power of two). D then outputs $(\text{AC}_d^0)\text{-MCSP}(y, t)$.

937 We now argue for correctness with high probability. Realize each y_i is independent with
 938 probability $\frac{\text{wt}(x)}{2n}$ of being 1. Hence, y is just the random variable $T_{p_w, m}$ where $w = \text{wt}(x)$.
 939 Hence, if $\text{wt}(x) = w_1$, then

$$940 \quad \Pr[R(x) \neq 1] = \Pr[\mathcal{C}(T_{p_{w_1}, m}) > t]$$

941 Recall that $t - \mathbb{E}[T_{p_{w_1}, m}] = \tilde{\Omega}(m/q)$ and, by assumption, $\text{CC}_{\text{AC}_d^0}$ on inputs of length m is
 942 $(m^{.25})$ -Lipchitz, so by Theorem A.2, we have that this probability is bounded by

$$943 \quad 2\exp\left(-2\frac{\tilde{\Omega}(m^2)}{\tilde{O}(q^2 m^{1.5})}\right) \leq \exp\left(-2\frac{\tilde{\Omega}(q^{5 \cdot 10})}{\tilde{O}(q^2)}\right) = O(\exp(-q^3))$$

944 using the fact that $m \geq q^{10}$. A similar analysis shows that the probability D errs if
 945 $\text{wt}(x) = w_1 + b$ is at most $O(\exp(-q^3))$. This completes the analysis of D .

946 We now argue that this reduction can be derandomized using non-uniformity. For each
 947 input of weight either w_1 or $w_1 + b$, we have shown the fraction of random strings which err
 948 on that input is $O(\exp(-q^3))$. Hence, the fraction of random seeds which err on at least one
 949 input of weight w_1 or $w_1 + b$ is at most

$$950 \quad 2^q O(\exp(-q^3)) < 1$$

951 for large enough n . Thus, there exists some fixed u_1, \dots, u_m and v_1, \dots, v_m which work on all
 952 inputs of length q . Once we are (non-uniformly) given these u_1, \dots, u_m and v_1, \dots, v_m which
 953 work on all inputs, we can turn D into an NC^0 oracle circuit C which has just a single gate
 954 (an oracle gate) whose inputs are the fixed number t and the string y where each bit of y is
 955 either a fixed bit of x or zero. This yields a NC^0 oracle circuit with $O(m) = O(q^{50}) = O(n^{100})$
 956 wires. ◀

957 ▶ **Corollary A.4.** *If n is sufficiently large, then for all distinct $w_1, w_2 \in [n]$ there is an NC^0*
 958 *oracle circuit C with at most two gates and $O(n^{100})$ wires such that $C^{(\text{AC}_d^0)\text{-MCSP}}$ computes*
 959 *$\text{WTDIS}_n[w_1, w_2]$.*

960 **Proof.** Fix some $w_1 \neq w_2$. Without loss of generality assume $w_1 < w_2$ (if this is not the case,
 961 then swap the names of w_1 and w_2 in this proof and add a NOT gate to the top of C). Let
 962 $b = w_2 - w_1$. Recall q is the least power of two such that $n \leq \sqrt{q}/2$. Note that $q = \Theta(n^2)$
 963 and $b \leq n \leq \sqrt{q}/2$. Theorem A.3 guarantees there exists an NC^0 oracle circuit D of size
 964 $O(n^{20})$ such that D^{MCSP} computes $\text{WTDIS}_q[w_3, w_3 + b]$ for some $w_3 \geq \sqrt{q}/2 \geq n$. Finally, let
 965 C be the oracle circuit that on input x outputs $D(y)$ where $y = 1^{w_3 - w_1} 0^{q - n - w_3 + w_1} x$. The
 966 correctness of this output is guaranteed by the fact that $\text{wt}(y) = w_3$ if and only if $\text{wt}(x) = w_1$
 967 and $\text{wt}(y) = w_3 + b$ if and only if $\text{wt}(x) = w_2$. ◀

968 ▶ **Corollary A.5.** *If n is sufficiently large, then there exists a depth-4 AC^0 truth table oracle*
 969 *circuit C with $O(n^{102})$ wires such that $C^{(\text{AC}_d^0)\text{-MCSP}}$ computes MAJORITY on strings on length*
 970 *n .*

971 **Proof.** It suffices to show that, for all $w \in [n]$, one can check if a string $x \in \{0, 1\}^n$ has
 972 weight w using a depth-3 AC^0 truth-table oracle circuit C_w of size $O(n^{101})$. If one is able to
 973 do this, then MAJORITY is computed by $\bigvee_{w \geq n/2} C_w(x)$.

974 For $w \in [n]$, let $\text{wt}_w : \{0, 1\}^n \rightarrow \{0, 1\}$ be the Boolean function that outputs one if and
 975 only if its input is a string of weight w . Now fix some $w \in [n]$. We claim that

$$976 \quad \text{wt}_w(x) = \bigwedge_{w' \in [n]: w \neq w'} \text{WTDIS}_n[w, w']$$

977 If x has weight w , then $\text{WTDIS}_n[w, w'](x) = 1$ for all $w' \neq w$, so

$$978 \quad \text{wt}_{w'}(x) = 1 = \bigwedge_{w' \in [n]: w \neq w'} \text{WTDIS}_n[w, w'].$$

979 On the other hand, if x has weight $w' \neq w$, then $\text{WTDIS}_n[w, w'](x) = 0$, so

$$980 \quad \text{wt}_w(x) = 0 = \bigwedge_{w' \in [n]: w \neq w'} \text{WTDIS}_n[w, w'].$$

981 Finally, by Corollary A.4 we have that $\bigwedge_{w \in [n]: w \neq w'} \text{WTDIS}_n[w, w]$ is computable by a
 982 depth-3 AC^0 truth table oracle circuit with $O(n^{101})$ wires. ◀

983 **B** $\text{NP} \leq_{tt}^P \text{MOCSP}$ implies $\text{EXP} \neq \text{ZPP}$

984 The proof of this result follows essentially exactly from Murray and Williams's [22] proof for
 985 MCSP. For completeness, we replicate the proof here (even using their words and structure).

986 ▶ **Proposition B.1.** *If $\text{NP} \leq_{tt}^P \text{MOCSP}$, then $\text{EXP} \subseteq \text{P/poly}$ implies $\text{EXP} = \text{NEXP}$.*

987 **Proof.** Assume $\text{NP} \leq_{tt}^P \text{MOCSP}$ and $\text{EXP} \subseteq \text{P/poly}$. Let $L \in \text{NTIME}(2^{n^c})$ for some $c \geq 1$. It
 988 suffices to show that $L \in \text{EXP}$.

989 We pad L into the $L' = \{x01^{2^{|x|^c}} : x \in L\}$. Note that $L' \in \text{NP}$. Hence there is a
 990 polynomial-time truth table reduction from L' to MOCSP . Composing the reduction from L
 991 to L' with the reduction from L' to MOCSP , we get a $2^{c'n^c}$ -time truth table reduction R
 992 from n -bit instances of L to $2^{c'n^c}$ -bit instances of MOCSP for some constant c' .

993 Let $Q(x)$ denote the concatenated string of all MOCSP queries produced by R in order
 994 on input x . Define the language

$$995 \quad \text{BITS}_Q := \{(x, i) : \text{the } i\text{th bit of } Q(x) \text{ is } 1\}$$

996 BITS_Q is clearly in EXP . Since $\text{EXP} \subseteq \text{P/poly}$, for some $d \geq 1$ there is a circuit family C_n
 997 of size at most $n^d + d$ computing BITS_Q on n -bit inputs.

998 Thus, on a given instance x , we have $\text{CC}(Q(x)) \leq s(|x|)$ where $s(|x|) := (|x| + 2c'|x|^c)^d + d$.
 999 Therefore, every MOCSP query (T, s', T_1, \dots, T_t) produced by the reduction R on input x
 1000 satisfies

$$1001 \quad \text{CC}^{T_1, \dots, T_t}(T) \leq \text{CC}(T) \leq e \cdot \text{CC}(Q(x)) \leq e \cdot s(|x|)$$

1002 for some constant e since T is a substring of $Q(x)$ (see Lemma 2.2 in [22] for a proof of this
 1003 substring fact). This leads to the following exponential time algorithm for L :

1004 On input x , run the exponential-time reduction $R(x)$ by using the following procedure
 1005 for answering each MOCSP oracle query $(T, s'; T_1, \dots, T_t)$. If $s' > e \cdot s(|x|)$, then
 1006 respond YES to the query. Otherwise, cycle through every oracle circuit E of size
 1007 at most s' . If E^{T_1, \dots, T_t} computes T , then respond YES. If no such E is found, then
 1008 respond NO.

1009 It suffices to show the procedure for answering MOCSP oracle queries runs in exponential time.
 1010 Let $n = |x|$. First, we need to count the number of oracle circuits E on $(\log |T| \leq c'n^c)$ -inputs
 1011 with size at most $s(n)$. As shown in Lemma 4.6, the logarithm of the number of oracle circuit
 1012 of size at most $s(|x|)$ on $(c'n^c)$ -inputs with t oracle functions is at most

$$1013 \quad O(s(n) \log(4 + t) + s(|x|) \log(c'n^c) \log(s(|x|) + \log(c'n^c))).$$

1014 Since $t \leq 2^{c'n^c}$ and s is polynomial in n , it is easy to see that the number of such circuits
 1015 E is at most exponential. Second, one can check if an oracle circuit E satisfies E^{T_1, \dots, T_t}
 1016 computes T in time polynomial in $(|E| + |T| + |T_1| + \dots + |T_t|)$ and hence exponential in n .
 1017 As a result, $L \in \text{EXP}$, completing the proof. ◀

1018 ▶ **Theorem B.2.** *If $\text{NP} \leq_{tt}^P \text{MOCSP}$, then $\text{EXP} \neq \text{NP} \cap \text{P/poly}$. Consequently, $\text{EXP} \neq \text{ZPP}$.*

1019 **Proof.** For contradiction, suppose $\text{NP} \leq_{tt}^P \text{MOCSP}$ and $\text{EXP} = \text{NP} \cap \text{P/poly}$. Then by
 1020 Proposition B.1 $\text{NEXP} \subseteq \text{EXP} \subseteq \text{NP}$ contradicting the nondeterministic time hierarchy
 1021 theorem [32]. ◀