

# Pure-Circuit: Strong Inapproximability for PPAD

Argyrios Deligkas  
Royal Holloway, UK  
argyrios.deligkas@rhul.ac.uk

Alexandros Hollender  
University of Oxford, UK  
alexandros.hollender@cs.ox.ac.uk

John Fearnley  
University of Liverpool, UK  
john.fearnley@liverpool.ac.uk

Themistoklis Melissourgos  
University of Essex, UK  
themistoklis.melissourgos@essex.ac.uk

*Abstract*—The current state-of-the-art methods for showing inapproximability in PPAD arise from the  $\varepsilon$ -Generalized-Circuit ( $\varepsilon$ -GCIRCUIT) problem. Rubinstein (2018) showed that there exists a small unknown constant  $\varepsilon$  for which  $\varepsilon$ -GCIRCUIT is PPAD-hard, and subsequent work has shown hardness results for other problems in PPAD by using  $\varepsilon$ -GCIRCUIT as an intermediate problem.

We introduce PURE-CIRCUIT, a new intermediate problem for PPAD, which can be thought of as  $\varepsilon$ -GCIRCUIT pushed to the limit as  $\varepsilon \rightarrow 1$ , and we show that the problem is PPAD-complete. We then prove that  $\varepsilon$ -GCIRCUIT is PPAD-hard for all  $\varepsilon < 0.1$  by a reduction from PURE-CIRCUIT, and thus strengthen all prior work that has used GCIRCUIT as an intermediate problem from the existential-constant regime to the large-constant regime.

We show that stronger inapproximability results can be derived by a direct reduction from PURE-CIRCUIT. In particular, we prove that finding an  $\varepsilon$ -well-supported Nash equilibrium in a polymatrix game is PPAD-hard for all  $\varepsilon < 1/3$ , and that this result is tight for two-action games.

*Index Terms*—TFNP, PPAD, approximation, Nash equilibrium, polymatrix games, generalized circuit

## I. INTRODUCTION

The complexity class PPAD has played a central role in determining the computational complexity of many problems arising in game theory and economics [Pap94]. The celebrated results of Daskalakis, Goldberg, and Papadimitriou [DGP09] and Chen, Deng, and Teng [CDT09] established that finding a Nash equilibrium in a strategic form game is PPAD-complete, and subsequent to this breakthrough many other PPAD-completeness results have been shown [CSVY08], [CDDT09], [VY11], [DQS12], [Das13], [KPR<sup>+</sup>13], [CDO15], [CPY17], [SSB17], [Meh18], [Rub18], [DFS20], [CKK21a], [CKK21b], [PP21], [GH21], [?], [DSZ21], [CCPY22].

These celebrated results not only showed that it is PPAD-hard to find an exact equilibrium, but also that finding approximate solutions is PPAD-hard. The result of Daskalakis, Goldberg, and Papadimitriou [DGP09] showed that finding an  $\varepsilon$ -Nash equilibrium is PPAD-complete when  $\varepsilon$  is exponentially small, while the result of Chen, Deng, and Teng [CDT09] improved this to show hardness for polynomially small  $\varepsilon$ . This

lower bound is strong enough to rule out the existence of an FPTAS for the problem.

The main open question following these results was whether equilibrium computation problems in PPAD were hard for constant  $\varepsilon$ , which would also rule out the existence of a PTAS. Here one must be careful, because some problems do in fact admit approximation schemes. For example, in the case of two-player strategic-form games, a quasipolynomial-time approximation scheme is known [LMM03], meaning that the problem cannot be hard for a constant  $\varepsilon$  unless every problem in PPAD can be solved in quasipolynomial time. But for other types of game such results are not known. This includes *polymatrix games*, which are  $n$ -player games with succinct representation [Jan68].

In another breakthrough result, Rubinstein [Rub18] developed techniques for showing constant inapproximability within PPAD, by proving that there exists a constant  $\varepsilon$  such that finding an  $\varepsilon$ -well-supported Nash equilibrium in a polymatrix game is PPAD-complete. This lower bound is obtained by first showing constant inapproximability for the  $\varepsilon$ -Generalized-Circuit ( $\varepsilon$ -GCIRCUIT) problem introduced by Chen, Deng, and Teng [CDT09], and then utilizing the known reduction from GCIRCUIT to polymatrix games [DGP09].

Rubinstein’s lower bound has since been used to show constant inapproximability for other problems. Rubinstein himself showed constant inapproximability for finding Bayesian Nash equilibria, relative approximate Nash equilibria, and approximate Nash equilibria in non-monotone markets [Rub18]. Subsequent work has shown constant inapproximability for finding clearing payments in financial networks with credit default swaps [SSB17], finding equilibria in first-price auctions with subjective priors [?], finding throttling equilibria in auction markets [CKK21b], finding equilibria in public goods games on directed networks [PP21], and finding consensus-halving solutions in fair division [FRFGZ18], [GHI<sup>+</sup>22].

Rubinstein’s lower bound is an *existential-constant* result, meaning that it shows that there exists some constant  $\varepsilon$  below which the problem becomes PPAD-hard. The fact that such a constant exists is important, since it rules out a PTAS. On the other hand, Rubinstein does not give any concrete lower bound

The second author was supported by EPSRC grant EP/W014750/1 “New Techniques for Resolving Boundary Problems in Total Search”.

on the size of the constant (understandably so, since this was not the purpose of his work). One could of course deduce such a lower bound by a careful examination of his reduction, but it is clear that this would yield an extremely small constant. Due to this, all of the other results that have utilized Rubinstein’s lower bound are likewise existential-constant results, which rule out PTASs but do not give any concrete lower bounds.

Ultimately, this means that existing work does not rule out an efficient algorithm that finds, say, a 0.001-approximate solution for any of these problems, which would likely be more than enough for most practical needs. Arguably, a player would be more than happy to know that her strategy is an optimal best-response, up to a loss of at most 0.001 in her utility value. Moreover, the existing work gives us no clue as to where the threshold for hardness may actually lie. To address these questions one would need to prove a *large-constant* inapproximability result, giving hardness for a known substantial constant.

Rubinstein’s lower bound is the ultimate source of all of the recent existential-constant lower bounds, so if one seeks a large-constant lower bound, then Rubinstein’s result is the bottleneck. Attempting to directly strengthen or optimize Rubinstein’s result does not seem like a promising direction. His proof, while ingenious, is very involved, and does not lend itself to easy optimization. Furthermore, it consists of many moving parts, so that even if one was able to optimize each module, the resulting constant would still be very small.

**Our contribution.** In this paper we introduce the techniques needed to show large-constant inapproximability results for problems in PPAD. Our key technical innovation is the introduction of a new intermediate problem, called PURE-CIRCUIT, which we show to be PPAD-complete.

Then, by reducing onwards from PURE-CIRCUIT, we show large-constant inapproximability results for a variety of problems in PPAD. In this sense, PURE-CIRCUIT now takes on the role that  $\varepsilon$ -GCIRCUIT has taken in the past, as an important intermediate problem from which all other results of this type are derived.

The PURE-CIRCUIT problem itself can be thought of as a version of  $\varepsilon$ -GCIRCUIT that is taken to its limits, and also dramatically simplified. In fact, the problem has only two gates (or, in a different formulation, three gates), which should be compared to  $\varepsilon$ -GCIRCUIT, which has nine distinct gates. Perhaps more importantly, the gates in PURE-CIRCUIT have very weak constraints on their outputs: the gates can be thought of as taking inputs in  $[0, 1]$ , and producing outputs in  $[0, 1]$ , but the gates themselves essentially only care about the values 0 and 1, with all other values being considered to be “bad” or “garbage” values (which we will later simply denote by “ $\perp$ ”, instead of using values in  $(0, 1)$ ). This should be compared to  $\varepsilon$ -GCIRCUIT gates, where, for example, one must output a value in  $[0, 1]$  that is within  $\varepsilon$  of the sum of two inputs.

Combined, these properties make PURE-CIRCUIT a very attractive problem to reduce from when showing a hardness

result, since one only has to implement two (or three) gates, and the constraints that one must simulate are very loose, making them easy to implement. We formally introduce PURE-CIRCUIT, and compare it to  $\varepsilon$ -GCIRCUIT, in Section II.

Our main result is to show that the PURE-CIRCUIT problem is PPAD-complete. It is worth noting that there is no  $\varepsilon$  in this result, and in fact the PURE-CIRCUIT problem does not even take a parameter  $\varepsilon$  in its definition. This is because, in some sense, the PURE-CIRCUIT problem can be viewed as a variant of  $\varepsilon$ -GCIRCUIT in which we have taken the limit  $\varepsilon \rightarrow 1$ . We give further justification of this idea in Section II, but at a high level, this means that there is no loss of  $\varepsilon$  in our main hardness result, with the only losses coming when one reduces onwards from PURE-CIRCUIT. The proof of our main result is presented in Section III, but we present a brief exposition of the main ideas in Section I-A.

Finally, in Section IV we present a number of new large-constant hardness results for problems in PPAD, all of which are shown via reductions from PURE-CIRCUIT. We begin by showing that  $\varepsilon$ -GCIRCUIT is PPAD-hard for all  $\varepsilon < 0.1$ , giving a direct strengthening of Rubinstein’s lower bound. This also implies large-constant inapproximability results for all of the problems that currently have existential-constant lower bounds proved via GCIRCUIT. However, to determine the constant, one would need to determine the amount of  $\varepsilon$  that is lost in each of the onward reductions, and these reductions often did not optimize this, since they were proving existential-constant lower bounds.

We argue that the way forward now is providing *direct* reductions from PURE-CIRCUIT in order to get the best possible hardness results. As evidence of this, we present the first *tight* inapproximability result for additive approximate equilibria in polymatrix games. Via a direct reduction from PURE-CIRCUIT, we show that finding an  $\varepsilon$ -well-supported Nash equilibrium in a polymatrix game (POLYMATRIXWSNE) is PPAD-hard for all  $\varepsilon < 1/3$ , even when every player only has two actions. This is much stronger than the lower bound of 0.05 that we would have obtained by a reduction from our lower bound for GCIRCUIT. It is also a tight result for two-action games: we give a polynomial-time algorithm for finding a  $1/3$ -well-supported Nash equilibrium, and so our lower bound completely characterizes the computational complexity of approximate well-supported equilibria in two-action polymatrix games.

Due to space constraints, the formal statements and proofs of our results for polymatrix games are omitted, but they can be found in the full version. The full version also contains further large-constant inapproximability results obtained through a direct reduction from PURE-CIRCUIT, such as for computing  $\varepsilon$ -Nash equilibria in polymatrix games and for finding approximate equilibria in threshold games [PP21].

### A. Proof Overview for Our Main Result

We begin this proof overview by defining a very weak version of PURE-CIRCUIT. An instance of the problem consists of a Boolean circuit using the standard gates NOT, AND,

and OR, but with the following tweak: the circuit is allowed to have *cycles*. A solution to the problem is an assignment of values to each node of the circuit, so that all gates are satisfied. If we are only allowed to assign values in  $\{0, 1\}$  to the nodes, then it is easy to see that the problem is not a total search problem, i.e., some instances do not have a solution. For example, there is no way to assign consistent values to a cycle of three consecutive NOT gates.

In order to ensure that the problem is total (and can thus be used to prove PPAD-hardness results), we make the value space continuous by extending it to  $[0, 1]$ . We extend the definition of the logical gates NOT, AND, and OR to non-Boolean inputs in the most permissive way: if at least one input to the gate is not a pure bit (i.e., not in  $\{0, 1\}$ ), then the gate is allowed to output any value in  $[0, 1]$ . The attentive reader might observe that this problem is now trivial to solve: just assign arbitrary values in  $(0, 1)$  to all the gates.

It is thus clear that the definition of the problem needs to be extended, by adding extra gates or by strengthening existing gates, so that the problem becomes PPAD-hard. However, in order to discover the least amount of additional structure needed to make the problem hard, it is instructive to proceed with this definition for now, and attempt to prove hardness.

In order to prove the PPAD-hardness of the problem, we cannot follow Rubinfeld’s approach, which goes through the construction of a continuous Brouwer function, because PURE-CIRCUIT only offers very weak gates. Instead, we proceed via a direct reduction from the STRONGSPERNER problem, a discrete problem that is a computational version of Sperner’s Lemma. The problem was shown to be PPAD-hard by Daskalakis, Skoulakis, and Zampetakis [DSZ21] (who called it the HIGHD-BISPERNER problem), and is the “PPAD-analogue” of the STRONGTUCKER problem which was recently used to prove PPA-hardness results [DFHM22]. This approach completely bypasses the continuous aspect of all such existing hardness reductions and enables us to work with the very weak gates that PURE-CIRCUIT offers.

At a high level, our hardness construction works as follows: the PURE-CIRCUIT instance implements the evaluation of the STRONGSPERNER labeling on some input point  $x$  (represented in unary by multiple nodes) and then uses a feedback mechanism to ensure that the circuit is only satisfied if  $x$  is a solution to the STRONGSPERNER instance. The full reduction is presented in Section III, but we mention here the two main obstacles when trying to implement this idea, and how to overcome them.

1. **The input point  $x$  might not be represented by a valid bitstring.** Indeed, since the gates take values in  $[0, 1]$  (and values in  $(0, 1)$  essentially do not carry any information), there is no guarantee that the input  $x$  will be represented by bits  $\{0, 1\}$ . But then the implementation of the STRONGSPERNER labeling (which is given as a Boolean circuit) will also fail. To resolve this issue, we introduce a new gate, the PURIFY gate, which, on any input, outputs two values, with the guarantee that at least one of them is a “pure” bit, i.e., 0 or 1. If the input is

already a pure bit, then both outputs are guaranteed to be copies of the input. Using a binary tree of PURIFY gates, we can now create many copies of  $x$ , such that most of them consist only of pure bits, and then use the logical gates to compute the STRONGSPERNER labeling correctly on these good copies.

2. **How to implement the feedback mechanism?** Given the outputs of the STRONGSPERNER labeling at all the copies of  $x$ , we now need to provide some kind of feedback to  $x$ , so that  $x$  is forced to change if it is not a solution of STRONGSPERNER. It turns out that this step can be performed if we have access to *sorting*: given a list of values in  $[0, 1]$ , sort them from smallest to largest. Unfortunately, this is impossible to achieve with the gates at our disposal, namely standard logical gates and the PURIFY gate. We circumvent this obstacle by observing that: (i) it is sufficient to be able to perform some kind of “weak sorting” (essentially, we only care about pure bits being sorted correctly), and (ii) this weak sorting can be achieved if we make our logical gates *robust*. For example, the robust version of the AND gate outputs 0, whenever at least one of its inputs is 0, irrespective of whether the other input is a pure bit or not.

With these two extensions in hand—namely, the PURIFY gate and the robustness of the logical gates—it is now possible to prove PPAD-hardness of the problem. A very natural question to ask is: Is it really necessary to add both extensions for the problem to be hard? In Appendix A we show that any attempt to weaken the gate-constraints makes the problem polynomial-time solvable. In particular, the introduction of the PURIFY gate is not enough by itself to make the problem PPAD-hard; the robustness of the logical gates is also needed.

The robustness of, say, the AND gate seems like a very natural constraint to impose. It is consistent with the meaning of the logical AND operation, but we also observe in our applications that this “robustness” seems to always be automatically satisfied in all simulations of the AND gate. On the other hand, the PURIFY gate, which might look a bit unnatural or artificial at first, actually corresponds to the simplest possible version of a bit decoder, a crucial tool in all prior works. As mentioned above, we show in Appendix A that these are the minimal gate-constraints that are needed for the problem to be PPAD-hard. In that sense, we argue that PURE-CIRCUIT captures the essence of PPAD-hardness: it consists of the minimal set of ingredients that are needed for a problem to be PPAD-hard.

The attentive reader might have noticed that our gates do not distinguish between different values in  $(0, 1)$ . For this reason, it will be more convenient to use a single symbol to denote such values in the definition of PURE-CIRCUIT (Section II) and in the rest of this paper. As explained in more detail in Section II, the symbol “ $\perp$ ” will be used to denote these “garbage” values. In other words, the nodes of the circuit will take values in  $\{0, 1, \perp\}$  instead of  $[0, 1]$ .

## II. THE PURE-CIRCUIT PROBLEM

In this section we define our new problem PURE-CIRCUIT and state our main result, namely its PPAD-completeness. Before defining PURE-CIRCUIT, we begin by explaining the intuition behind its definition, and how it relates to the Generalized-Circuit (GCIRCUIT) problem.

**The Generalized-Circuit problem.** In the Generalized-Circuit (GCIRCUIT) problem (formally defined in Section IV-A) we are given a circuit and the goal is to assign a value to each node of the circuit so that each gate is computed correctly. Importantly, the circuit is a *generalized* circuit, meaning that cycles are allowed. If cycles were not allowed, then it would be easy to find values satisfying all gates: just pick arbitrary values for the input gates, and then evaluate the circuit on those inputs.

Every node of GCIRCUIT must be assigned a value in  $[0, 1]$ , and the gates are arithmetic gates, such as addition, subtraction, multiplication by a constant (with output truncated to lie in  $[0, 1]$ ), and suitably defined logical gates. Reducing from GCIRCUIT is very useful for obtaining hardness of approximation results, because the problem remains PPAD-hard, even when we allow some error at every gate. In the  $\varepsilon$ -GCIRCUIT problem, the goal is to assign a value in  $[0, 1]$  to each node of the circuit, so that each gate is computed correctly, up to an additive error of  $\pm\varepsilon$ .

The problem was first defined by Chen et al. [CDT09], who proved that it is PPAD-hard for inverse polynomial  $\varepsilon$ , and who used it to prove PPAD-hardness of finding Nash equilibria in bimatrix games. Prior to that, Daskalakis et al. [DGP09] had implicitly proved that it is PPAD-hard for inverse exponential  $\varepsilon$ . Rubinfeld’s [Rub18] breakthrough result proved that there exists some constant  $\varepsilon > 0$  such that  $\varepsilon$ -GCIRCUIT remains PPAD-hard.

**Taking the limit  $\varepsilon \rightarrow 1$ .** In order to get strong inapproximability results, it seems necessary to prove hardness of  $\varepsilon$ -GCIRCUIT for large, explicit, values of  $\varepsilon$ . Ideally, we would like to obtain hardness for the largest possible  $\varepsilon$ . While it is unclear what that value is for GCIRCUIT, in theory, as long as  $\varepsilon < 1$  the output of a gate still carries some information. Namely a gate whose actual output should be 0 cannot take the value 1.

This observation leads us to define a problem to essentially capture the setting  $\varepsilon \rightarrow 1$ . In that case, a node carries information only if its value is 0 or 1. Otherwise, its value is irrelevant. As a result, the natural operations to consider in this setting are simple Boolean operations, such as NOT, AND, OR, NAND, and NOR. We only require these gates to output the correct result when their input is relevant, i.e., 0 or 1. For example, the NOT gate should output 1 on input 0, and output 0 on input 1, but there is no constraint on its output when the input lies in  $(0, 1)$ .

Since values in  $(0, 1)$  do not carry any information, and are as such interchangeable (e.g., a value  $1/2$  can be replaced by  $1/3$  without impacting any of the gates), we will instead use the symbol “ $\perp$ ” to denote any and all values in  $(0, 1)$ . In other

words, instead of assigning a value in  $[0, 1]$  to each gate, we will assign a value in  $\{0, 1, \perp\}$ , where  $\perp$  is interpreted as a “garbage” value, i.e., not corresponding to a pure bit value 0 or 1. With this new notation, the updated description of the NOT gate would be that it must output 1 on input 0, it must output 0 on input 1, and it can output anything (namely, 0, 1, or  $\perp$ ) on input  $\perp$ .

Unfortunately, if we only allow these logical gates, then the problem is trivial to solve: assigning the “garbage” value  $\perp$  (or any value in  $(0, 1)$  if we use the old notation) to every node will satisfy all gates. Thus, we need a gate that makes this impossible.

**The PURIFY gate.** To achieve this, we introduce the PURIFY gate: a gate with one input and two outputs, which, intuitively, “purifies” its input. When fed with an actual pure bit, the PURIFY gate outputs two copies of the input bit. However, when the input is not a pure bit, the gate still ensures that at least one of its two outputs is a pure bit. In more detail:

- If the input is 0, then both outputs are 0.
- If the input is 1, then both outputs are 1.
- If the input is  $\perp$ , then at least one of the outputs is a pure bit, i.e., 0 or 1.

Note that the gate is quite “under-defined”. For example, we do not specify which pure bit the gate should output when the input is  $\perp$ , nor do we specify the output on which this bit appears. This is actually an advantage, because it makes it easier to reduce from the problem, since the less constrained the gates are, the easier it is to simulate them in the target application problem.

**Robustness of the logical gates.** The introduction of the PURIFY gate makes the problem non-trivial: if a PURIFY gate appears in the circuit, then assigning the “garbage” value  $\perp$  to all nodes is no longer a solution. However, it turns out that one more modification is needed to make the problem PPAD-hard: we have to make the logical gates *robust*. For the AND gate, this means the following: if one of its two inputs is 0, then the output is 0, no matter what the other input is (even if it is not a pure bit, i.e., if it is  $\perp$ ). Similarly, for the OR gate we require that the output be 1 when at least one of the two inputs is 1. Robustness is defined analogously for NAND and NOR.

We show that introducing the PURIFY gate and making the logical gates robust is enough to make the problem PPAD-complete. Next, we define the problem formally and state our main result.

**Formal definition.** In the definition below, we use the PURIFY and NOR gates, because these two gates are enough for the problem to already be PPAD-complete. However, the problem remains hard for various other combinations of gates and restrictions on the interactions between nodes, as we detail in Corollaries II.2 and II.3. In Appendix A we discuss the definition in more detail, and explain why any attempt at relaxing the definition (in particular, removing the robustness) makes the problem polynomial-time solvable.

**Definition 1 (PURE-CIRCUIT).** An instance of PURE-CIRCUIT is given by a vertex set  $V = [n]$  and a set  $G$  of gate-constraints (or just *gates*). Each gate  $g \in G$  is of the form  $g = (T, u, v, w)$  where  $u, v, w \in V$  are distinct nodes and  $T \in \{\text{NOR}, \text{PURIFY}\}$  is the type of the gate, with the following interpretation.

- If  $T = \text{NOR}$ , then  $u$  and  $v$  are the inputs of the gate, and  $w$  is its output.
- If  $T = \text{PURIFY}$ , then  $u$  is the input of the gate, and  $v$  and  $w$  are its outputs.

We require that each node is the output of exactly one gate.

A solution to instance  $(V, G)$  is an assignment  $\mathbf{x} : V \rightarrow \{0, 1, \perp\}$  that satisfies all the gates, i.e., for each gate  $g = (T, u, v, w) \in G$  we have:

- if  $T = \text{NOR}$ , then  $\mathbf{x}$  satisfies (top: mathematically; bottom: truth table)

$$\begin{aligned} \mathbf{x}[u] = \mathbf{x}[v] = 0 &\implies \mathbf{x}[w] = 1 \\ (\mathbf{x}[u] = 1) \vee (\mathbf{x}[v] = 1) &\implies \mathbf{x}[w] = 0 \end{aligned}$$

$u$	$v$	$w$
0	0	1
1	$\{0, 1, \perp\}$	0
$\{0, 1, \perp\}$	1	0
Else		$\{0, 1, \perp\}$

- if  $T = \text{PURIFY}$ , then  $\mathbf{x}$  satisfies

$$\begin{aligned} \{\mathbf{x}[v], \mathbf{x}[w]\} \cap \{0, 1\} &\neq \emptyset \\ \mathbf{x}[u] \in \{0, 1\} &\implies \mathbf{x}[v] = \mathbf{x}[w] = \mathbf{x}[u] \end{aligned}$$

$u$	$v$	$w$
0	0	0
1	1	1
$\perp$	At least one	output in $\{0, 1\}$

The following theorem is our main technical result and is proved in Section III.

**Theorem II.1.** *The PURE-CIRCUIT problem is PPAD-complete.*

The most important part of this statement is of course the PPAD-hardness of PURE-CIRCUIT, but let us briefly discuss the other part, namely the PPAD-membership. This is obtained as a byproduct of our results in Section IV, where we reduce PURE-CIRCUIT to various problems that are known to lie in PPAD. However, there is also a more direct way to prove membership in PPAD, and in particular to establish the existence of a solution, and we briefly sketch it here. Indeed, the PURE-CIRCUIT problem can be reduced to the problem of finding a Brouwer fixed point of a continuous function  $F$ , a problem known to lie in PPAD [Pap94], [EY10]. Given an instance of PURE-CIRCUIT with  $n$  nodes, the function  $F : [0, 1]^n \rightarrow [0, 1]^n$  is constructed by letting  $x \in [0, 1]^n$  represent an assignment of values to the  $n$  nodes, and by defining  $F_i(x) \in [0, 1]$  as a continuous function that outputs a

valid value for the  $i$ th node, given that the other nodes have values according to assignment  $x$  (where any value in  $(0, 1)$  is interpreted as “ $\perp$ ”). For every type of gate, it is not hard to construct a continuous piecewise-linear function  $F_i$  (or, in the case of PURIFY, two such functions  $F_i$  and  $F_j$ ) that satisfies the constraints of that type of gate.

#### A. Alternative Gates and Further Restrictions

In this section, we present various versions of the problem that remain PPAD-complete, in particular versions that use alternative gates and have additional restrictions.

**More gates.** We define the following additional gates.

- If  $T = \text{COPY}$  in  $g = (T, u, v)$ , then  $\mathbf{x}$  satisfies

$$\begin{aligned} \mathbf{x}[u] = 0 &\implies \mathbf{x}[v] = 0 \\ \mathbf{x}[u] = 1 &\implies \mathbf{x}[v] = 1 \end{aligned}$$

$u$	$v$
0	0
1	1
$\perp$	$\{0, 1, \perp\}$

- If  $T = \text{NOT}$  in  $g = (T, u, v)$ , then  $\mathbf{x}$  satisfies

$$\begin{aligned} \mathbf{x}[u] = 0 &\implies \mathbf{x}[v] = 1 \\ \mathbf{x}[u] = 1 &\implies \mathbf{x}[v] = 0 \end{aligned}$$

$u$	$v$
0	1
1	0
$\perp$	$\{0, 1, \perp\}$

- If  $T = \text{OR}$  in  $g = (T, u, v, w)$ , then  $\mathbf{x}$  satisfies

$$\begin{aligned} \mathbf{x}[u] = \mathbf{x}[v] = 0 &\implies \mathbf{x}[w] = 0 \\ (\mathbf{x}[u] = 1) \vee (\mathbf{x}[v] = 1) &\implies \mathbf{x}[w] = 1 \end{aligned}$$

$u$	$v$	$w$
0	0	0
1	$\{0, 1, \perp\}$	1
$\{0, 1, \perp\}$	1	1
Else		$\{0, 1, \perp\}$

- If  $T = \text{AND}$  in  $g = (T, u, v, w)$ , then  $\mathbf{x}$  satisfies

$$\begin{aligned} \mathbf{x}[u] = \mathbf{x}[v] = 1 &\implies \mathbf{x}[w] = 1 \\ (\mathbf{x}[u] = 0) \vee (\mathbf{x}[v] = 0) &\implies \mathbf{x}[w] = 0 \end{aligned}$$

$u$	$v$	$w$
1	1	1
0	$\{0, 1, \perp\}$	0
$\{0, 1, \perp\}$	0	0
Else		$\{0, 1, \perp\}$

- If  $T = \text{NAND}$  in  $g = (T, u, v, w)$ , then  $\mathbf{x}$  satisfies

$$\begin{aligned} \mathbf{x}[u] = \mathbf{x}[v] = 1 &\implies \mathbf{x}[w] = 0 \\ (\mathbf{x}[u] = 0) \vee (\mathbf{x}[v] = 0) &\implies \mathbf{x}[w] = 1 \end{aligned}$$

$u$	$v$	$w$
1	1	0
0	$\{0, 1, \perp\}$	1
$\{0, 1, \perp\}$	0	1
Else		$\{0, 1, \perp\}$

**Corollary II.2.** *The PURE-CIRCUIT problem is PPAD-complete, for any of the following choices of gate types:*

- PURIFY and at least one of {NOR, NAND};
- PURIFY, NOT, and at least one of {OR, AND}.

*Proof.* This follows from Theorem II.1 by observing that a NOR gate can always be simulated with the given set of gates. Clearly, NOR can be simulated by first using an OR gate and then a NOT gate. Furthermore, OR can be simulated by NOT and AND by applying De Morgan’s laws. Finally, AND can easily be obtained from NOT and NAND, and NOT can be obtained from NAND and PURIFY as follows: first apply a PURIFY gate, and then use its two outputs as the two inputs to a NAND gate.  $\square$

**More structure.** The hardness result is also robust with respect to restrictions applied to the *interaction graph*. This graph is constructed on the vertex set  $V = [n]$  by adding a directed edge from node  $u$  to node  $v$  whenever  $v$  is the output of a gate with input  $u$ . For example, a NOR gate with inputs  $u, v$  and output  $w$  yields the two edges  $(u, w)$  and  $(v, w)$ . On the other hand, a PURIFY gate with input  $u$  and outputs  $v, w$  gives the edges  $(u, v)$  and  $(u, w)$ . Since any given node is the output of at most one gate, it immediately follows that the in-degree of every node is at most 2. However, the out-degree of a node can *a priori* be arbitrarily large. It is quite easy to show that the problem remains PPAD-complete, even if we severely restrict the interaction graph.

**Corollary II.3.** *The PURE-CIRCUIT problem remains PPAD-complete, for any choice of gates {PURIFY, X, Y}, where  $(X, Y) \in \{\text{NOT}\} \times \{\text{OR, AND, NOR, NAND}\}$  or  $(X, Y) \in \{\text{COPY}\} \times \{\text{NOR, NAND}\}$  and even if we also simultaneously have all of the following restrictions.*

- 1) Every node is the input of exactly one gate.
- 2) In the interaction graph, the total degree of every node is at most 3. More specifically, for every node, the in- and out-degrees,  $d_{in}$  and  $d_{out}$ , satisfy  $(d_{in}, d_{out}) \in \{(1, 1), (2, 1), (1, 2)\}$ .
- 3) The interaction graph is bipartite.

The proof of this corollary is again quite simple, and it can be found in the full version.

*Remark 1.* Using Corollary II.3, it is also possible to show that PURE-CIRCUIT with only two gates (namely, PURIFY and one of {NOR, NAND}) remains PPAD-complete even if the total degree of every node is at most 4 in the interaction graph. Indeed, NOT gates can be implemented by first using a PURIFY gate and then a NOR/NAND gate. The structural properties of Corollary II.3 ensure that this yields an interaction graph where the total degree is at most 4 for each node.

This can be further reduced to degree 3, if one modifies the definition of PURE-CIRCUIT (Definition 1) so that the two inputs to a NOR/NAND gate are no longer required to be two *distinct* nodes  $u$  and  $v$ , but can possibly be the same node  $u = v$ . However, if the definition is modified in that way, then one must be careful when reducing from PURE-CIRCUIT to make sure to take into account the possibility that  $u = v$  when constructing the gadget for a NOR/NAND gate.

### III. PPAD-COMPLETENESS OF PURE-CIRCUIT

This section proves our main technical result, namely that PURE-CIRCUIT is PPAD-complete (Theorem II.1). We note that membership in PPAD follows immediately from the reduction of the problem to GCIRCUIT in Section IV-A. In order to establish the PPAD-hardness, we present a polynomial-time reduction from a PPAD-complete problem to PURE-CIRCUIT. The canonical PPAD-complete problem is the END-OF-LINE problem, but, as is usually the case, we do not reduce directly from END-OF-LINE, but from a problem with topological structure instead, which we introduce next.

#### A. The STRONGSPERNER Problem

We will reduce from the STRONGSPERNER problem, which is based on a variant of Sperner’s lemma [Spe28]. This problem is in essence the same as the HIGHD-BISPERNER problem defined by Daskalakis et al. [DSZ21] and used to prove PPAD-hardness of a problem related to constrained min-max optimization. Furthermore, the corresponding “strong” variant of Tucker’s lemma was used by Deligkas et al. [DFHM22] to provide improved PPA-hardness results for the consensus-halving problem in fair division.

**Definition 2.** The STRONGSPERNER problem:

**Input:** A Boolean circuit computing a labeling  $\lambda : [M]^N \rightarrow \{-1, +1\}^N$  satisfying the following boundary conditions for every  $i \in [N]$ :

- if  $x_i = 1$ , then  $[\lambda(x)]_i = +1$ ;
- if  $x_i = M$ , then  $[\lambda(x)]_i = -1$ .

**Output:** Points  $x^{(1)}, \dots, x^{(N)} \in [M]^N$  that satisfy  $\|x^{(i)} - x^{(j)}\|_\infty \leq 1$  for all  $i, j \in [N]$ , and such that  $\lambda(x^{(1)}), \dots, \lambda(x^{(N)})$  cover all labels, i.e., for all  $i \in [N]$  and  $\ell \in \{-1, +1\}$  there exists  $j \in [N]$  with  $[\lambda(x^{(j)})]_i = \ell$ .

Note that the requirement that a solution should consist of exactly  $N$  points is without loss of generality. If we find less than  $N$  points that cover all labels, then we can simply re-use the same points multiple times to obtain a list of  $N$  points that cover all labels (there is no requirement on them being distinct). If we find more than  $N$  points that cover all labels, then it is easy to see that we can extract a subset of  $N$  points that still cover all labels in polynomial time [DFHM22, Lemma 3.1].

**Theorem III.1** ([DSZ21]). *STRONGSPERNER is PPAD-hard, even when  $M$  is only polynomially large (i.e., given in unary in the input).*

*Remark 2.* This was proven by Daskalakis et al. [DSZ21] by reducing from the SUCCINCTBROUWER problem, which had been proven PPAD-hard by Rubinfeld [Rub16]. The PPAD-hardness can also be proved by a more direct reduction from END-OF-LINE. Indeed, END-OF-LINE can be reduced to STRONGSPERNER with  $N = 2$  and exponentially large  $M$  by using the techniques of Chen and Deng [CD09]. Then, a snake embedding technique [CDT09], [DFHM22] can be used to obtain hardness for the high-dimensional version with small  $M$ , in fact, even for constant  $M$ . For our purposes, the hardness for polynomially large  $M$  is sufficient.

### B. Reduction from STRONGSPERNER to PURE-CIRCUIT

Consider an instance  $\lambda : [M]^N \rightarrow \{-1, +1\}^N$  of STRONGSPERNER, where  $\lambda$  is given as a Boolean circuit and  $M$  is only polynomially large (i.e., given in unary). We will now show how to construct an instance of PURE-CIRCUIT in polynomial time such that from any correct assignment to the nodes, we can extract a solution to the STRONGSPERNER instance in polynomial time. We will make use of the gates PURIFY, AND, OR, NOT, COPY. All these gates can easily be simulated using the two gates PURIFY and NOR, by the arguments in the proof of Corollary II.2.

We begin the construction of the PURE-CIRCUIT instance by creating nodes  $u_{i,1}, \dots, u_{i,M}$  for each  $i \in [N]$ . We call these nodes the *original inputs*, and we think of  $u_{i,1}, \dots, u_{i,M}$  as being the unary representation of an element in  $[M]$ . Of course, this only makes sense when all these nodes are assigned pure bit values, i.e., 0 or 1. In general, this will not be the case. The rest of the instance can be divided into four parts: the *purification* stage, the *circuit* stage, the *sorting* stage, and the *selection* stage.

We begin with a brief overview of the purpose of each stage and how they interact with each other. The purification stage uses the PURIFY gate to create multiple “copies” of the original  $u_{i,j}$  nodes, while ensuring that most of the copies have pure bit values (Purification Lemma, Lemma III.2). Then, the circuit stage evaluates the circuit  $\lambda$  on these copies of the original inputs. Since the purification stage ensures that most copies have pure bits, the circuit stage outputs the correct labels for most copies, and, in particular, most outputs are pure bits (Circuit Lemma, Lemma III.3). Next, for each  $i \in [N]$ , the sorting stage “sorts” the list of all  $i$ th output values computed in the circuit stage (Sorting Lemma, Lemma III.4). Since most of the  $i$ th output values are pure bits, the sorting stage ensures that all non-pure values are close to each other in the sorted list. The selection stage then proceeds to select  $M$  values from the sorted list, but in a careful way, namely such that they are all far away from each other. This ensures that at most one of the  $M$  selected values is not a pure bit. The  $M$  selected values are then fed back into the original inputs  $u_{i,1}, \dots, u_{i,M}$ . As a result, the original input  $u_{i,1}, \dots, u_{i,M}$  contains at most one non-pure bit, and thus the purification stage ensures that all the (pure) copies of  $u_{i,1}, \dots, u_{i,M}$  correspond to unary numbers that differ by at most 1. This means that these copies represent points in the STRONGSPERNER domain that are within  $\ell_\infty$ -

distance 1 (Selection Lemma, Lemma III.5). Finally, using the boundary conditions of the STRONGSPERNER instance, we argue that these points must cover all the labels (Solution Lemma, Lemma III.6).

We now describe each of the stages in more detail. We let  $K$  denote the number of copies that we make. It will be enough to pick  $K = 3NM^2$ . In what follows,  $\mathbf{x}$  always denotes an arbitrary solution to the PURE-CIRCUIT instance we construct.

**Step 1: Purification stage.** For each  $(i, j) \in [N] \times [M]$ , we construct a binary tree of PURIFY gates that is rooted at  $u_{i,j}$  and has leaves  $u_{i,j}^{(1)}, \dots, u_{i,j}^{(K)}$ .

We say that  $k \in [K]$  is a good copy, if  $\mathbf{x}[u_{i,j}^{(k)}]$  is a pure bit for all  $(i, j) \in [N] \times [M]$ . We denote the set of all good copies by  $G$ , i.e.,

$$G := \left\{ k \in [K] : \mathbf{x}[u_{i,j}^{(k)}] \in \{0, 1\} \quad \forall (i, j) \in [N] \times [M] \right\}.$$

For a good copy  $k \in G$  and any  $i \in [N]$ , we can interpret the bitstring  $(\mathbf{x}[u_{i,j}^{(k)}])_{j \in [M]} \in \{0, 1\}^M$  as representing a number in  $[M]$  in unary, which we denote by  $u_i^{(k)} \in [M]$ . In other words,  $u_i^{(k)}$  corresponds to the number of 1’s in the bit-string  $(\mathbf{x}[u_{i,j}^{(k)}])_{j \in [M]}$ . For notational convenience we let the all-zero bit string  $0^M$  correspond to 1 as well, i.e., if  $\mathbf{x}[u_{i,j}^{(k)}] = 0$  for all  $j \in [M]$ , then  $u_i^{(k)} = 1$ . (Alternatively, we could also use  $M - 1$  bits instead of  $M$ .) Finally, we let  $u^{(k)} \in [M]^N$  denote the vector  $(u_1^{(k)}, \dots, u_N^{(k)})$ .

**Lemma III.2** (Purification Lemma). *The following hold.*

- 1) *There are at least  $K - NM$  good copies, i.e.,  $|G| \geq K - NM$ .*
- 2) *If for some  $(i, j) \in [N] \times [M]$  the original input  $\mathbf{x}[u_{i,j}]$  is a pure bit, then all copies have that same bit, i.e.,  $\mathbf{x}[u_{i,j}^{(k)}] = \mathbf{x}[u_{i,j}]$  for all  $k \in [K]$ .*

*Proof.* Observe that in a binary tree of PURIFY gates, if some node has some pure value  $b \in \{0, 1\}$ , then all nodes in the subtree rooted at this node also have value  $b$ . Applying this observation at the root of the tree rooted at  $u_{i,j}$ , we immediately obtain part 2 of the statement.

For part 1, note that for any  $(i, j) \in [N] \times [M]$ , in the binary tree of PURIFY gates rooted at  $u_{i,j}$ , all leaves, except at most one, have a pure bit value. This follows from the definition of the PURIFY gate and the observation about subtrees made in the previous paragraph. As a result, all values  $(\mathbf{x}[u_{i,j}^{(k)}])_{(i,j,k) \in [N] \times [M] \times [K]}$  are pure bits, except for at most  $NM$  of them. But this means that there are at least  $K - NM$  values of  $k \in [K]$  such that  $\mathbf{x}[u_{i,j}^{(k)}] \in \{0, 1\}$  for all  $\forall (i, j) \in [N] \times [M]$ . In other words,  $|G| \geq K - NM$ .  $\square$

**Step 2: Circuit stage.** We assume, without loss of generality, that  $\lambda$  is given as a Boolean circuit  $C : (\{0, 1\}^M)^N \rightarrow \{0, 1\}^N$  using gates AND, OR, NOT, and, on input  $z \in (\{0, 1\}^M)^N$ :

- For each  $i \in [N]$ , the  $i$ th block of input bits  $z_i \in \{0, 1\}^M$  is interpreted by  $C$  as representing a number  $\bar{z}_i \in [M]$  in unary, in the exact same way as  $u_i^{(k)} \in [M]$  is obtained from the bitstring  $(\mathbf{x}[u_{i,j}^{(k)}])_{j \in [M]}$ .

- The circuit  $C$  outputs  $\lambda(\bar{z}_1, \dots, \bar{z}_N) \in \{-1, +1\}^N$ , where a  $-1$  output is represented by a 0, and a  $+1$  output by a 1. To keep things simple, in the rest of this exposition we will abuse notation and think of  $\lambda$  as outputting labels in  $\{0, 1\}^N$ .

If the circuit is not originally in this form, then it can be brought in this form in polynomial time.

In the circuit stage, we construct  $K$  separate copies of the circuit  $C$ , using the AND, OR, and NOT gates. For each  $k \in [K]$ , the  $k$ th copy  $C_k$  takes as input the nodes  $(u_{i,j}^{(k)})_{(i,j) \in [N] \times [M]}$  and we denote its output nodes by  $v_1^{(k)}, \dots, v_N^{(k)}$ . Since the gates always have correct output when the inputs are pure bits, we immediately obtain the following.

**Lemma III.3** (Circuit Lemma). *For all good copies  $k \in G$ , the output of circuit  $C_k$  is correct, i.e.,  $\mathbf{x}[v_i^{(k)}] = [\lambda(u^{(k)})]_i$  for all  $i \in [N]$ .*

**Step 3: Sorting stage.** In this stage, for each  $i \in [N]$ , we would like to have a gadget that takes as input the list of nodes  $v_i^{(1)}, \dots, v_i^{(K)}$  (namely, the list of  $i$ th outputs of the circuits  $C_1, \dots, C_K$ ) and outputs the nodes  $w_i^{(1)}, \dots, w_i^{(K)}$ , such that these output nodes are a sorted list of the values of the input nodes (where we think of the values as being ordered  $0 < \perp < 1$ ). Unfortunately, this is not possible given the gates we have at our disposal. However, it turns out that we can do some kind of “weak” sorting by using the robustness of the AND and OR gates (i.e., the fact that AND on input 0 and  $s$ , always outputs 0, no matter what  $s \in \{0, 1, \perp\}$  is).

For now assume that we consider values in  $[0, 1]$  (instead of  $\{0, 1, \perp\}$ ) and that we have access to a comparator gate that takes two inputs  $s_1$  and  $s_2$  and outputs  $t_1$  and  $t_2$ , such that  $t_1, t_2$  is the sorted list  $s_1, s_2$ . Formally, we can write this as  $t_1 := \min\{s_1, s_2\}$  and  $t_2 := \max\{s_1, s_2\}$ . Using comparator gates, it is easy to construct a circuit that takes  $K$  inputs and outputs them in sorted order. Indeed, we can directly implement a sorting network [Knu98], for example. Even a very naive approach will yield such a circuit of polynomial size, which is all we need. We implement this circuit with inputs  $v_i^{(1)}, \dots, v_i^{(K)}$  and outputs  $w_i^{(1)}, \dots, w_i^{(K)}$  in our PURE-CIRCUIT instance, by replacing every comparator gate by AND and OR gates. Namely, to implement a comparator gate with inputs  $s_1, s_2$  and outputs  $t_1, t_2$ , we use an AND gate with inputs  $s_1, s_2$  and output  $t_1$ , and an OR gate with inputs  $s_1, s_2$  and output  $t_2$ . The robustness of the AND and OR gates allows us to prove that this sorting gadget sorts the pure bit values correctly, in the following sense.

**Lemma III.4** (Sorting Lemma). *Let  $K_0$  and  $K_1$  denote the number of zeroes and ones that the  $i$ th sorting gadget gets as input, i.e.,  $K_b := |\{k \in [K] : \mathbf{x}[v_i^{(k)}] = b\}|$ . Then, the first  $K_0$  outputs of the gadget are zeroes, and the last  $K_1$  outputs are ones. Formally,  $\mathbf{x}[w_i^{(k)}] = 0$  for all  $k \in [K_0]$ , and  $\mathbf{x}[w_i^{(K+1-k)}] = 1$  for all  $k \in [K_1]$ .*

*Proof.* Consider the ideal sorting circuit (that uses comparator gates) with input  $f(\mathbf{x}[v_i^{(1)}]), \dots, f(\mathbf{x}[v_i^{(K)}])$ , where  $f$  maps 0

to 0, 1 to 1, and  $\perp$  to  $1/2$ . In other words, we imagine running the comparator circuit on our list of values, except that the “garbage” value  $\perp$  is replaced by  $1/2$ . Since the comparator circuit correctly sorts the list, its output satisfies the desired property: the first  $K_0$  outputs are 0, and the last  $K_1$  outputs are 1. Thus, in order to prove the lemma, it suffices to prove the following claim: *if a node in the ideal circuit has a pure bit value  $b \in \{0, 1\}$ , then the corresponding node in our PURE-CIRCUIT instance must also have value  $b$ .*

We prove the claim by induction. Clearly, all input nodes satisfy the claim. Now consider some node  $t_1$  that is the min-output of a comparator gate with inputs  $s_1$  and  $s_2$ , that both satisfy the claim. Recall that this gate will be implemented in the PURE-CIRCUIT by an AND gate with inputs  $s_1, s_2$  and output  $t_1$ . If the ideal circuit assigns value  $1/2$  to  $t_1$ , then the claim trivially holds for  $t_1$ . If the ideal circuit assigns value 1 to  $t_1$ , then both  $s_1$  and  $s_2$  must have value 1 in the ideal circuit. Since the claim holds for  $s_1$  and  $s_2$ , they also have value 1 in PURE-CIRCUIT, and so the AND gate will ensure that  $t_1$  also has value 1, thus satisfying the claim. Finally, if the ideal circuit assigns value 0 to  $t_1$ , then it must also have assigned value 0 to at least one of  $s_1$  or  $s_2$ . But then, by the claim, PURE-CIRCUIT also assigns value 0 to at least one of  $s_1$  or  $s_2$ , and the robustness of the AND gate ensures that  $t_1$  also has value 0. The same argument also works with max and OR instead.  $\square$

Note that Lemma III.4 only guarantees a “weak” type of sorting: some parts of the output list might not be correctly ordered, and the list of output values might not be a permutation of the input values (namely, it can happen that there are more 0’s and/or 1’s in the output list than in the input list). However, this “weak” sorting will be enough for our needs as we will see below.

**Step 4: Selection stage.** Since the list  $w_i^{(1)}, \dots, w_i^{(K)}$ , is now “sorted”, we can select  $M$  nodes from it in such a way that at most one node does not have a pure value. Indeed, this can be achieved by selecting nodes that are sufficiently far apart from each other. We thus select the nodes  $(w_i^{(j \cdot 2^{NM})})_{j \in [M]}$  and copy their values onto the original input nodes  $(u_{i,j})_{j \in [M]}$ . Namely, for each  $j \in [M]$  we introduce a COPY gate with input  $w_i^{(j \cdot 2^{NM})}$  and output  $u_{i,j}$ . Recall that  $K = 3NM^2 \geq M \cdot 2^{NM}$ , so this is well defined. This selection procedure has the following nice properties.

**Lemma III.5** (Selection Lemma). *We have:*

- 1) *All good copies are close to each other:  $\|u^{(k)} - u^{(k')}\|_\infty \leq 1$  for all  $k, k' \in G$ .*
- 2) *If all good copies agree that the  $i$ th output is  $b \in \{0, 1\}$ , i.e.,  $\mathbf{x}[v_i^{(k)}] = b$  for all  $k \in G$ , then the original input satisfies  $\mathbf{x}[u_{i,j}] = b$  for all  $j \in [M]$ .*

*Proof.* We begin by proving part 1 of the statement. First of all, note that it suffices to show that, for all  $i \in [N]$ , at most one of the original inputs  $\mathbf{x}[u_{i,1}], \dots, \mathbf{x}[u_{i,M}]$  is not a pure bit. Indeed, in that case, by part 2 of the Purification Lemma (Lemma III.2), it follows that for any  $i \in [N]$  and any  $k, k' \in$



$G$ , the bitstrings  $(\mathbf{x}[u_{i,j}^{(k)}])_{j \in [M]}$  and  $(\mathbf{x}[u_{i,j}^{(k')}])_{j \in [M]}$  differ in at most one bit. But this implies that  $|u_i^{(k)} - u_i^{(k')}| \leq 1$  for all  $i \in [N]$ , and thus  $\|u^{(k)} - u^{(k')}\|_\infty \leq 1$ .

Now consider any  $i \in [N]$ . By part 1 of the Purification Lemma (Lemma III.2), we have that the number of good copies  $|G| \geq K - NM$ . By the Circuit Lemma (Lemma III.3) we know that the corresponding outputs are pure bits, i.e.,  $\mathbf{x}[v_i^{(k)}] \in \{0, 1\}$  for all  $k \in G$ . Thus, the list  $\mathbf{x}[v_i^{(1)}], \dots, \mathbf{x}[v_i^{(K)}]$  contains at least  $K - NM$  pure bits. Using the notation from the Sorting Lemma (Lemma III.4), this means that  $K_0 + K_1 \geq K - NM$ . As a result, by applying the Sorting Lemma, it follows that in the list obtained after sorting,  $\mathbf{x}[w_i^{(1)}], \dots, \mathbf{x}[w_i^{(K)}]$ , all the non-pure bits are contained in an interval of length  $NM$ . Since the selected nodes  $(w_i^{(j \cdot 2NM)})_{j \in [M]}$  are sufficiently far apart (namely  $2NM$ ), at most one such node can fall in the “bad” interval. This means that all selected nodes, except at most one, are pure bits, and since the selected nodes are copied into the original inputs, this also holds for them.

It remains to prove part 2 of the statement. If for some  $i \in [N]$ ,  $\mathbf{x}[v_i^{(k)}] = 0$  for all  $k \in G$ , then this means that  $K_0 \geq |G| \geq K - NM$ . By the Sorting Lemma (Lemma III.4), it follows that  $\mathbf{x}[w_i^{(k)}] = 0$  for all  $k \in [K - NM]$ . In particular, since  $K - NM = 3NM^2 - NM \geq M \cdot 2NM$ , this means that for all  $j \in [M]$ ,  $\mathbf{x}[w_i^{(j \cdot 2NM)}] = 0$ . But these are the nodes we select, and we copy their value into the original inputs, so the statement follows. The case where  $\mathbf{x}[v_i^{(k)}] = 1$  for all  $k \in G$  is handled analogously.  $\square$

**Correctness.** The description of the reduction is now complete. We have constructed a valid instance of PURE-CIRCUIT in polynomial time. In particular, note that every node is the output of exactly one gate. To complete the proof, it remains to prove that from any solution of the PURE-CIRCUIT instance we can extract a solution to STRONGSPERNER in polynomial time. We do this in the following final lemma.

**Lemma III.6** (Solution Lemma). *The points  $\{u^{(k)} : k \in G\} \subseteq [M]^N$  yield a solution to the STRONGSPERNER instance  $\lambda$ .*

*Proof.* First of all, by part 1 of the Selection Lemma (Lemma III.5), we know that the points  $u^{(k)}$ ,  $k \in G$ , are all within  $\ell_\infty$ -distance 1. Thus, it suffices to prove that they cover all labels with respect to  $\lambda$ . Note that we can efficiently extract these points from a solution  $\mathbf{x}$ , since, for each  $k \in [K]$ , we can easily decide whether  $k$  lies in  $G$  or not, and then extract the point  $u^{(k)}$ .

We prove by contradiction that the points  $u^{(k)}$ ,  $k \in G$ , must cover all labels. Assume, on the contrary, that there exists  $i \in [N]$  and  $b \in \{0, 1\}$  such that  $[\lambda(u^{(k)})]_i = b$  for all  $k \in G$ . Then, by the Circuit Lemma (Lemma III.3), all the corresponding circuits must output  $b$ , i.e.,  $\mathbf{x}[v_i^{(k)}] = b$  for all  $k \in G$ . By part 2 of the Selection Lemma (Lemma III.5), it follows that the original input satisfies  $\mathbf{x}[u_{i,j}] = b$  for all  $j \in [M]$ . Finally, by part 2 of the Purification Lemma (Lemma III.2), it must be that for all copies  $k \in G$ , all  $M$

bits  $\mathbf{x}[u_{i,1}^{(k)}], \dots, \mathbf{x}[u_{i,M}^{(k)}]$  are equal to  $b$ . Now, if  $b = 1$ , then this means that  $u_i^{(k)} = M$ , and thus  $[\lambda(u^{(k)})]_i = 0 \neq b$  by the STRONGSPERNER boundary conditions, which contradicts the original assumption. Similarly, if  $b = 0$ , then  $u_i^{(k)} = 1$ , and thus  $[\lambda(u^{(k)})]_i = 1 \neq b$  by the STRONGSPERNER boundary conditions, which is again a contradiction. (We recall here that we have renamed the labels  $\{-1, +1\}$  to  $\{0, 1\}$ , respectively, for the purpose of this proof.)  $\square$

## IV. APPLICATIONS

In this section, we derive strong inapproximability lower bounds for PPAD-complete problems, by reducing from PURE-CIRCUIT.

### A. Generalized Circuit

The  $\varepsilon$ -GCIRCUIT problem was introduced by Chen, Deng, and Teng [CDT09]. In this section, we show that  $\varepsilon$ -GCIRCUIT is PPAD-hard for all  $\varepsilon < 0.1$ .

**The  $\varepsilon$ -GCIRCUIT problem.** The problem is defined as follows.

**Definition 3** (Generalized Circuit [CDT09]). A generalized circuit is a tuple  $(V, T)$ , where  $V$  is a set of nodes, and  $T$  is a set of gates. Each gate  $t \in T$  is a five-tuple  $(G, u, v, w, c)$ , where  $G$  is a gate type from the set  $\{G_c, G_{\times c}, G_{=}, G_{+}, G_{-}, G_{<}, G_{\vee}, G_{\wedge}, G_{-}\}$ ,  $u, v \in V \cup \{\text{nil}\}$  are input variables,  $w \in V$  is an output variable, and  $c \in [0, 1] \cup \{\text{nil}\}$  is a rational constant.

The following requirements must be satisfied for each gate  $(G, u, v, w, c) \in T$ .

- $G_c$  gates take no input variables and use a constant in  $[0, 1]$ . So  $u = v = \text{nil}$  and  $c \in [0, 1]$  whenever  $G = G_c$ .
- $G_{\times c}$  gates take one input variable and a constant. So  $u \in V$ ,  $v = \text{nil}$ , and  $c \in [0, 1]$  whenever  $G = G_{\times c}$ .
- $G_{=}$  and  $G_{-}$  gates take one input variable and do not use a constant. So  $u \in V$ ,  $v = c = \text{nil}$ , whenever  $G \in \{G_{=}, G_{-}\}$ .
- All other gates take two input variables and do not use a constant. So  $u \in V$ ,  $v \in V$ , and  $c = \text{nil}$  whenever  $G \notin \{G_c, G_{\times c}, G_{=}, G_{-}\}$ .
- Every variable in  $V$  is the output variable for exactly one gate. More formally, for each variable  $w \in V$ , there is exactly one gate  $t \in T$  such that  $t = (G, u, v, w, c)$ .

The  $\varepsilon$ -GCIRCUIT problem is defined as follows. Given a generalized circuit  $(V, T)$ , find a vector  $\mathbf{x} \in [0, 1]^{|V|}$  such that for each gate in  $T$  the following constraints are satisfied.

Gate	Constraint
$(G_c, \text{nil}, \text{nil}, w, c)$	$\mathbf{x}[w] = c \pm \varepsilon$
$(G_{\times c}, u, \text{nil}, w, c)$	$\mathbf{x}[w] = \mathbf{x}[u] \cdot c \pm \varepsilon$
$(G_-, u, \text{nil}, w, \text{nil})$	$\mathbf{x}[w] = \mathbf{x}[u] \pm \varepsilon$
$(G_+, u, v, w, \text{nil})$	$\mathbf{x}[w] = \min(\mathbf{x}[u] + \mathbf{x}[v], 1) \pm \varepsilon$
$(G_-, u, v, w, \text{nil})$	$\mathbf{x}[w] = \max(\mathbf{x}[u] - \mathbf{x}[v], 0) \pm \varepsilon$
$(G_{<}, u, v, w, \text{nil})$	$\mathbf{x}[w] = \begin{cases} 1 \pm \varepsilon & \text{if } \mathbf{x}[u] < \mathbf{x}[v] - \varepsilon \\ 0 \pm \varepsilon & \text{if } \mathbf{x}[u] > \mathbf{x}[v] + \varepsilon \end{cases}$
$(G_{\vee}, u, v, w, \text{nil})$	$\mathbf{x}[w] = \begin{cases} 1 \pm \varepsilon & \text{if } \mathbf{x}[u] \geq 1 - \varepsilon \\ & \text{or } \mathbf{x}[v] \geq 1 - \varepsilon \\ 0 \pm \varepsilon & \text{if } \mathbf{x}[u] \leq \varepsilon \text{ and } \mathbf{x}[v] \leq \varepsilon \end{cases}$
$(G_{\wedge}, u, v, w, \text{nil})$	$\mathbf{x}[w] = \begin{cases} 1 \pm \varepsilon & \text{if } \mathbf{x}[u] \geq 1 - \varepsilon \\ & \text{and } \mathbf{x}[v] \geq 1 - \varepsilon \\ 0 \pm \varepsilon & \text{if } \mathbf{x}[u] \leq \varepsilon \text{ or } \mathbf{x}[v] \leq \varepsilon \end{cases}$
$(G_{\neg}, u, \text{nil}, w, \text{nil})$	$\mathbf{x}[w] = \begin{cases} 1 \pm \varepsilon & \text{if } \mathbf{x}[u] \leq \varepsilon \\ 0 \pm \varepsilon & \text{if } \mathbf{x}[u] \geq 1 - \varepsilon \end{cases}$

Here the notation  $a = b \pm \varepsilon$  is used as a shorthand for  $a \in [b - \varepsilon, b + \varepsilon]$ . We will also make use of gates of type  $(G_{>}, u, v, w, \text{nil})$  which enforce the constraint

$$\mathbf{x}[w] = \begin{cases} 0 \pm \varepsilon & \text{if } \mathbf{x}[u] < \mathbf{x}[v] - \varepsilon \\ 1 \pm \varepsilon & \text{if } \mathbf{x}[u] > \mathbf{x}[v] + \varepsilon \end{cases}$$

Gates of type  $G_{>}$  can be easily built by using a  $G_{\neg}$  gate to negate the output of a  $G_{<}$  gate.

In the remainder of this section, we prove the following result.

**Theorem IV.1.**  $\varepsilon$ -GCIRCUIT is PPAD-hard for every  $\varepsilon < 0.1$ .

*Proof.* We will reduce from the PURE-CIRCUIT problem that uses the gates NOR and PURIFY, which we showed to be PPAD-hard in Theorem II.1. We will encode 0 values in the PURE-CIRCUIT problem as values in the range  $[0, \varepsilon]$  in GCIRCUIT, while 1 values will be encoded as values in the range  $[1 - \varepsilon, 1]$ . Then, each gate from PURE-CIRCUIT will be simulated by a combination of gates in the GCIRCUIT instance.

**NOR gates.** A NOR gate  $(\text{NOR}, u, v, w)$  will be simulated by GCIRCUIT gates that compute

$$\mathbf{x}[u] + \mathbf{x}[v] < 5/9,$$

which requires us to use  $G_+$ ,  $G_{<}$  and  $G_c$ . We claim that this gate works for any  $\varepsilon < 1/9$ .

The idea is that if both inputs lie in the range  $[0, \varepsilon]$ , and thus both inputs encode zeros in PURE-CIRCUIT, we will have  $\mathbf{x}[u] + \mathbf{x}[v] \leq 3\varepsilon < 3/9$ , where the extra  $\varepsilon$  error is introduced by  $G_+$ . The  $G_c$  gate outputting  $5/9$  may output a value as small as  $4/9$  after the error is taken into account. Thus, the  $G_{<}$  gate will compare these two values, and provided that  $\varepsilon < 1/9$ , the comparison will succeed, so the gate will output a value greater than or equal to  $1 - \varepsilon$ , which corresponds to a 1 in the PURE-CIRCUIT instance, as required.

If, on the other hand, at least one input is in the range  $[1 - \varepsilon, 1]$ , then we will have  $\mathbf{x}[u] + \mathbf{x}[v] \geq 1 - 2\varepsilon > 7/9$ , where again  $G_+$  introduces an extra  $\varepsilon$  error. The  $G_c$  gate outputting  $5/9$  may output a value as large as  $6/9$  after the error is taken

into account. Thus, the  $G_{<}$  will compare these two values and provided that  $\varepsilon < 1/9$ , the comparison will succeed, so the gate will output a value less than or equal to  $\varepsilon$ , which corresponds to a 0 in the PURE-CIRCUIT instance, which is again as required.

**PURIFY gates.** A PURIFY gate  $(\text{PURIFY}, u, v, w)$  will be simulated by two GCIRCUIT gadgets. The first will set  $\mathbf{x}[v]$  equal to  $\mathbf{x}[u] > 0.3$ , while the second will set  $\mathbf{x}[w]$  equal to  $\mathbf{x}[u] > 0.7$ , where  $G_c$  and  $G_{>}$  gates are used to implement these operations. We claim that this construction works for any  $\varepsilon < 0.1$ .

We begin by considering the first gadget, which sets  $\mathbf{x}[v]$  equal to  $\mathbf{x}[u] > 0.3$ . If  $\mathbf{x}[u] \leq \varepsilon < 0.1$ , then note that the  $G_c$  gate outputting 0.3 may output a value as small as 0.2 once errors are taken into account. Thus, since  $\varepsilon < 0.1$ , the comparison made by the  $G_{>}$  gate will succeed, and so  $\mathbf{x}[v]$  will be set to a value less than or equal to  $\varepsilon$ . On the other hand, if  $\mathbf{x}[u] \geq 0.5$ , then note that the  $G_c$  gate outputting 0.3 may output a value as large as 0.4 once errors are taken into account. So, the comparison made by the  $G_{>}$  gate will again succeed, and  $\mathbf{x}[v]$  will be set to a value greater than or equal to  $1 - \varepsilon$ .

One can repeat the analysis above for the second gadget to conclude that  $\mathbf{x}[w]$  will be set to a value less than or equal to  $\varepsilon$  when  $\mathbf{x}[u] \leq 0.5$ , and a value greater than or equal to  $1 - \varepsilon$  when  $\mathbf{x}[u] \geq 1 - \varepsilon$ . So we can verify that the conditions of the PURIFY gate are correctly simulated.

- If  $\mathbf{x}[u] \leq \varepsilon$ , meaning that the input encodes a zero, then  $\mathbf{x}[v]$  and  $\mathbf{x}[w]$  will be set to values that are less than or equal to  $\varepsilon$ , and so both outputs encode zeros.
- If  $\mathbf{x}[u] \geq 1 - \varepsilon$ , meaning that the input encodes a one, then  $\mathbf{x}[v]$  and  $\mathbf{x}[w]$  will be set to values that are greater than or equal to  $1 - \varepsilon$ , and so both outputs encode ones.
- No matter what value  $\mathbf{x}[u]$  takes, at least one output will be set to a value that encodes a zero or a one. Specifically, if  $\mathbf{x}[u] \leq 0.5$ , then the second comparison gate will set  $\mathbf{x}[w] \leq \varepsilon$ , while if  $\mathbf{x}[u] \geq 0.5$ , then the first comparison gate will set  $\mathbf{x}[v] \geq 1 - \varepsilon$ .

Thus, the construction correctly simulates a PURIFY gate. Note that  $\varepsilon$  cannot be increased beyond 0.1 in this construction, since then there would be no guarantee that an encoding of a zero or a one would be produced when  $\mathbf{x}[u] = 0.5$ .

**The lower bound.** Given a PURE-CIRCUIT instance defined over variables  $V$ , we produce a GCIRCUIT instance by replacing each gate in the PURE-CIRCUIT with the constructions given above. Then, given a solution  $\mathbf{x}$  to the 0.1-GCIRCUIT instance, we can produce a solution  $\mathbf{x}'$  to PURE-CIRCUIT in the following way. For each  $v \in V$

- if  $\mathbf{x}[v] \leq \varepsilon$ , then we set  $\mathbf{x}'[v] = 0$ ,
- if  $\mathbf{x}[v] \geq 1 - \varepsilon$ , then we set  $\mathbf{x}'[v] = 1$ , and
- if  $\mathbf{x}[v] > \varepsilon$  and  $\mathbf{x}[v] < 1 - \varepsilon$ , then we set  $\mathbf{x}'[v] = \perp$ .

By the arguments given above, we have that  $\mathbf{x}'$  is indeed a solution to PURE-CIRCUIT, and thus Theorem IV.1 is proved.  $\square$

APPENDIX A  
ON THE DEFINITION OF PURE-CIRCUIT

In this section, we explore different ways to weaken the definition of PURE-CIRCUIT, and show how, in each case, the problem is no longer PPAD-hard.

**No PURIFY gate.** If we allow all gates, except the PURIFY gate (so only NOT, COPY, OR, AND, NOR, NAND), then the problem becomes polynomial-time solvable. Indeed, it suffices to assign value  $\perp$  to all the nodes.

**No Negation.** If we allow all gates, except the ones that perform some kind of negation (so only PURIFY, COPY, OR, AND), then the problem becomes polynomial-time solvable. Indeed, assigning the value 1 to each node (or, alternatively, the value 0 to each node) always yields a solution. More generally, we can make the following observation: for any set of gates that can be implemented by monotone functions, the problem lies in the class PLS, and is thus unlikely to be PPAD-complete. Indeed, as already mentioned in Section II, we can view any PURE-CIRCUIT instance with  $n$  nodes as a function  $F : [0, 1]^n \rightarrow [0, 1]^n$ , where each gate is replaced by a continuous function that is consistent with the gate-constraint. Then any fixed point of  $F$  yields a solution to the PURE-CIRCUIT instance. If each of the gates can be replaced by a continuous *monotone* function, then the problem of finding a fixed point of  $F$  is an instance of Tarski’s fixed point theorem, which is known to lie in PLS [EPRY20].

**No robustness.** If we allow all gates (PURIFY, NOT, COPY, OR, AND, NOR, NAND), but we drop the robustness requirement from the logical gates OR, AND, NOR, NAND, then the problem can be solved in polynomial time.

Construct the interaction graph  $G$  of the PURE-CIRCUIT instance, as defined in Section II-A. In the first stage of the algorithm, as long as there exists a directed cycle in graph  $G$ , we do the following:

- 1) Pick an arbitrary directed simple cycle of  $G$ .
- 2) For each node  $u$  on the simple cycle  $C$ , assign value  $\perp$  to it, i.e.,  $x[u] := \perp$ .
- 3) Remove all nodes on the simple cycle  $C$  from the graph  $G$ , including all their incident edges.

At the end of this procedure,  $G$  no longer contains any cycles. In the second stage of the algorithm we then repeat the following, until  $G$  is empty:

- 1) Pick any source  $u$  of  $G$  (which must exist, since  $G$  contains no cycles).
- 2) Let  $g$  be the (unique) gate that has  $u$  as output. Since  $u$  does not have incoming edges in  $G$ , all inputs of  $g$  have already been assigned a value. If  $u$  is the only output of  $g$ , then assign a value to  $u$  that satisfies the gate, and remove  $u$  and its edges from  $G$ . If the gate  $g$  also has another output  $v$ , then  $g$  is a PURIFY gate, and there are two cases:
  - If  $v$  has not been assigned a value yet, then assign values to both  $u$  and  $v$  such that the gate is satisfied. Remove  $u$  and  $v$  and their edges from  $G$ .

- If  $v$  has already been assigned a value, then this happened in the first stage of the algorithm, and both  $v$  and the input to  $g$  were assigned value  $\perp$  (if  $v$  lies on a simple cycle  $C$ , then so does the input of gate  $g$ , because  $v$  has a single incoming edge in the original interaction graph). In that case, we assign value 0 or 1 to  $u$  and  $g$  is satisfied. Remove  $u$  and its edges from  $G$  ( $v$  was already removed in the first stage).

Since a node is removed from  $G$  only when it is assigned a value, all nodes have been assigned a value at the end of the algorithm. We argue that all gates are satisfied. Clearly, any gate that has an output node that is still present in  $G$  after the end of the first stage will be satisfied (by construction of the second stage). Thus, it remains to consider any gate  $g$  such that all its output nodes are removed in the first stage. There are three cases:

- $g$  is a NOT or COPY gate: if the output lies on a simple cycle  $C$ , then so does the input. Both are thus assigned value  $\perp$  and the gate is satisfied.
- $g$  is a (non-robust) OR, AND, NOR, or NAND gate: if the output lies on a simple cycle  $C$ , then so does at least one of its inputs. Thus, at least one input is also assigned value  $\perp$ , and the gate is satisfied. Note that we crucially used the non-robustness of the gate here.
- $g$  is a PURIFY gate: if an output  $v$  of  $g$  lies on a simple cycle  $C$ , then the input  $u$  of  $g$  also lies on  $C$ , and so both are removed at the same time from  $G$ . However, the other output  $w$  of  $g$  cannot lie on that same simple cycle  $C$ , and after  $u$  is removed,  $w$  does not have an incoming edge anymore and will thus not be removed in the first stage. Thus,  $g$  cannot be a PURIFY gate.

**Adding Constant Gates.** It is a natural idea to try to add constant gates in an attempt to make the problem hard for a set of gates for which the problem is not PPAD-hard. By constant gates, we mean a 0-gate which has one output and no input, and which enforces that the value of its output node always be 0, and a 1-gate defined analogously. No matter which subset of gates  $S \subseteq \{\text{PURIFY, NOT, COPY, OR, AND, NOR, NAND}\}$  we use, adding the constant gates 0 and 1 does not change the complexity of the problem. Indeed, it is easy to see that the constants can be “propagated” through the circuit. If a constant is an input to a PURIFY, NOT, or COPY gate, then we can replace the output(s) of that gate by constants that satisfy the gate-constraint. If a constant is an input to an X gate, where  $X \in \{\text{OR, AND, NOR, NAND}\}$ , then there are two cases: either we can replace the output by a constant, or we can replace the gate by an X gate with the same input twice (namely, the other input). In order to create a copy of the other input, we can either use the COPY gate, or the NOT gate, or the PURIFY gate. If none of these three gates lies in  $S$ , then PURE-CIRCUIT with gates  $S \cup \{0, 1\}$  is polynomial-time solvable, since it is polynomial-time solvable with gates  $S \cup \{\text{NOT, 0, 1}\}$  (by the propagation argument, and the lack of PURIFY gate).

## ACKNOWLEDGMENT

We thank the anonymous reviewers for comments and suggestions that helped improve the presentation of the paper.

## REFERENCES

- [CCPY22] T. Chen, X. Chen, B. Peng, and M. Yannakakis, “Computational hardness of the Hylland-Zeckhauser scheme,” in *Proceedings of the 33rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2022, pp. 2253–2268.
- [CD09] X. Chen and X. Deng, “On the complexity of 2D discrete fixed point problem,” *Theoretical Computer Science*, vol. 410, no. 44, pp. 4448–4456, 2009.
- [CDDT09] X. Chen, D. Dai, Y. Du, and S. Teng, “Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities,” in *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2009, pp. 273–282.
- [CDO15] X. Chen, D. Durfee, and A. Orfanou, “On the complexity of Nash equilibria in anonymous games,” in *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC)*, 2015, pp. 381–390.
- [CDT09] X. Chen, X. Deng, and S.-H. Teng, “Settling the complexity of computing two-player Nash equilibria,” *Journal of the ACM*, vol. 56, no. 3, pp. 14:1–14:57, 2009.
- [CKK21a] X. Chen, C. Kroer, and R. Kumar, “The complexity of pacing for second-price auctions,” in *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021, p. 318.
- [CKK21b] —, “Throttling equilibria in auction markets,” in *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)*, 2021, p. 551.
- [CPY17] X. Chen, D. Paparas, and M. Yannakakis, “The complexity of non-monotone markets,” *Journal of the ACM*, vol. 64, no. 3, pp. 20:1–20:56, 2017.
- [CSVY08] B. Codenotti, A. Saberi, K. Varadarajan, and Y. Ye, “The complexity of equilibria: Hardness results for economies via a correspondence with games,” *Theoretical Computer Science*, vol. 408, no. 2–3, pp. 188–198, 2008.
- [Das13] C. Daskalakis, “On the complexity of approximating a Nash equilibrium,” *ACM Transactions on Algorithms*, vol. 9, no. 3, pp. 1–35, 2013.
- [DFHM22] A. Deligkas, J. Fearnley, A. Hollender, and T. Melissourgos, “Constant inapproximability for PPA,” in *Proceedings of the 54th ACM Symposium on Theory of Computing (STOC)*, 2022, pp. 1010–1023.
- [DFS20] A. Deligkas, J. Fearnley, and R. Savani, “Tree polymatrix games are PPAD-hard,” in *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, 2020, pp. 38:1–38:14.
- [DGP09] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, “The complexity of computing a Nash equilibrium,” *SIAM Journal on Computing*, vol. 39, no. 1, pp. 195–259, 2009.
- [DQS12] X. Deng, Q. Qi, and A. Saberi, “Algorithmic solutions for envy-free cake cutting,” *Operations Research*, vol. 60, no. 6, pp. 1461–1476, 2012.
- [DSZ21] C. Daskalakis, S. Skoulakis, and M. Zampetakis, “The complexity of constrained min-max optimization,” in *Proceedings of the 53rd ACM Symposium on Theory of Computing (STOC)*, 2021, pp. 1466–1478.
- [EPRY20] K. Etessami, C. Papadimitriou, A. Rubinstein, and M. Yannakakis, “Tarski’s theorem, supermodular games, and the complexity of equilibria,” in *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS)*, 2020, pp. 18:1–18:19.
- [EY10] K. Etessami and M. Yannakakis, “On the complexity of Nash equilibria and other fixed points,” *SIAM Journal on Computing*, vol. 39, no. 6, pp. 2531–2597, 2010.
- [FRFGZ18] A. Filos-Ratsikas, S. K. S. Frederiksen, P. W. Goldberg, and J. Zhang, “Hardness results for Consensus-Halving,” in *Proceedings of the 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 2018, pp. 24:1–24:16.
- [GH21] P. W. Goldberg and A. Hollender, “The Hairy Ball problem is PPAD-complete,” *Journal of Computer and System Sciences*, vol. 122, pp. 34–62, 2021.
- [GHI<sup>+</sup>22] P. W. Goldberg, A. Hollender, A. Igarashi, P. Manurangsi, and W. Suksompong, “Consensus halving for sets of items,” *Mathematics of Operations Research*, vol. 47, no. 4, 2022.
- [Jan68] E. Janovskaja, “Equilibrium points in polymatrix games,” *Lithuanian Mathematical Journal*, vol. 8, no. 2, pp. 381–384, 1968.
- [Knu98] D. E. Knuth, *The Art of Computer Programming, Volume 3: Sorting and Searching*. Addison-Wesley Professional, 1998.
- [KPR<sup>+</sup>13] S. Kintali, L. J. Poplawski, R. Rajaraman, R. Sundaram, and S. Teng, “Reducibility among fractional stability problems,” *SIAM Journal on Computing*, vol. 42, no. 6, pp. 2063–2113, 2013.
- [LMM03] R. J. Lipton, E. Markakis, and A. Mehta, “Playing large games using simple strategies,” in *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*, 2003, pp. 36–41.
- [Meh18] R. Mehta, “Constant rank two-player games are PPAD-hard,” *SIAM Journal on Computing*, vol. 47, no. 5, pp. 1858–1887, 2018.
- [Pap94] C. H. Papadimitriou, “On the complexity of the parity argument and other inefficient proofs of existence,” *Journal of Computer and System Sciences*, vol. 48, no. 3, pp. 498–532, 1994.
- [PP21] C. Papadimitriou and B. Peng, “Public goods games in directed networks,” in *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021, pp. 745–762.
- [Rub16] A. Rubinstein, “Settling the complexity of computing approximate two-player Nash equilibria,” in *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016, pp. 258–265.
- [Rub18] —, “Inapproximability of Nash equilibrium,” *SIAM Journal on Computing*, vol. 47, no. 3, pp. 917–959, 2018.
- [Spe28] E. Sperner, “Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes,” *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, vol. 6, pp. 265–272, 1928.
- [SSB17] S. Schuldenzucker, S. Seuken, and S. Battiston, “Finding clearing payments in financial networks with credit default swaps is PPAD-complete,” in *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, 2017, pp. 32:1–32:20.
- [VY11] V. V. Vazirani and M. Yannakakis, “Market equilibrium under separable, piecewise-linear, concave utilities,” *Journal of the ACM*, vol. 58, no. 3, pp. 10:1–10:25, 2011.