

On the Hardness of Simple Stochastic Games

Brendan Juba

Abstract

We examine the problem of finding optimal strategies in *simple stochastic games*, and the equivalent problem of finding stable configurations of MIN/MAX/AVG circuits. The problem can be seen as a nontrivial extension to linear programming, but no polynomial-time algorithm is known, despite significant efforts to find such an algorithm. Our investigation will be concerned with the relative computational power of certain variants of these circuits. We will begin by defining the computational problems associated with these games and circuits, and stating in which senses the two are equivalent. We will show that restrictions of the problem are in \mathbf{P} , that the simple stochastic game problem is in $\mathbf{NP} \cap \mathbf{coNP}$, that the stable circuit problem is in $\mathbf{PLS} \cap \mathbf{PPAD}$, and that the problem of deciding whether a MIN or MAX gate in a MIN/MAX/AVG circuit will take the left or right input is \mathbf{PSPACE} -hard. We will conclude by discussing how these results fit together.

1 Introduction

Simple stochastic games are a class of games originally introduced and studied by Shapley in 1953 [14]. The computational problem of determining whether or not a player has a probability of winning a simple stochastic game of greater than $1/2$ was studied extensively by Condon much later [4]. At that point, the problem was widely conjectured to be solvable in polynomial-time, but many early attempts at algorithms were either shown to be incorrect or have known exponential counterexamples [5]. Since then, significant effort has been put into finding algorithms, with little success, so we conjecture that the problem is neither in \mathbf{P} nor in \mathbf{BPP} , and thus we wish to show that the problem is hard for some natural class.

Of course, since the problem has been shown to be in $\mathbf{NP} \cap \mathbf{coNP}$ [4], we have restricted our attention to classes relatively close to \mathbf{P} . This suggests examining classes such as \mathbf{PLS} or \mathbf{PPAD} for hardness results – we show in Section 3 that an equivalent function problem is contained in both – but as we will discuss later, hardness results may not be possible. Certainly, if it is hard for one of these classes, it will have some interesting consequences, as we explore in Section 4.

To get a handle on the hardness of simple stochastic games, we were interested in seeing what sorts of gadgets we could build out of instances of simple stochastic games. One of the interesting features of this problem is that it has a variety of equivalent formulations. Condon [4] observed that the problem is equivalent to the decision problem for languages accepted by logspace alternating randomized Turing machines. In this paper, in Section 2, we develop yet another natural equivalent formulation of the problem: finding the stable configurations of circuits of MIN, MAX, and AVG gates with fixed 0/1 inputs. In particular, we show that an instance of one yields an instance of the other in quite a natural way, and that the computational problems associated with each are polynomial-time equivalent. This result is not new; rather, it is a bit of folklore which we have fleshed out [1]. It is then using this circuit formulation of the problem that we began attempting to build gadgets.

In particular, we define a variant on the problem, examining the circuits under a “Leapfrog” model, in which we allow ourselves to read a bit of information from the MIN or MAX gates—we

informally say that the circuit accepts if a particular gate chooses a different input (the LEFT vs. the RIGHT input) during its computation. We present a more formal definition and a brief discussion of the motivation behind this definition in Section 2.6. This model is the focus of Section 3, where we show that these circuits can solve **PSPACE**-hard problems, specifically TQBF; along the way, we give an approach to developing a binary counter in an analog monotone circuit, a seemingly counter-intuitive feat. As such, while these “Leapfrog” circuits are interesting in their own right, they are probably not suitable for use in a reduction to some class contained in $\mathbf{NP} \cap \mathbf{coNP}$, under standard complexity-theoretic assumptions. Section 3.4 defines an alternative decision problem for Leapfrog circuits, which is also **PSPACE**-hard for much the same reasons, and examines where precisely the difficulty lies in relating the solutions to these Leapfrog problems to the stable solutions of the circuit instances.

We wrap up by discussing all of these considerations in greater detail in Section 4, particularly emphasizing the difficulties we encountered in designing gadgets in MIN/MAX/AVG circuits. We further develop the relationships between the stable circuit problem and the classes **PLS** and **PPAD**, and finally close by suggesting some interesting open questions raised by this work.

2 Definitions and Preliminaries

Sections 2.1 and 2.2 introduce the usual definitions for simple stochastic games and MIN/MAX/AVG circuits, respectively. We define our simple stochastic games as done by Condon [4] and define our MIN/MAX/AVG circuits largely as in [2]. In Section 2.3, we make explicit the implicit equivalences hinted at in [2], using some preliminary results from Condon [4]. Finally in Section 2.6, we formally state a new variant on the problem which we will consider in-depth. In Section 2.7, we will state the definitions of certain subclasses of the class of total functions that can be verified in polynomial time, **TFNP**, which were defined by Papadimitriou et al. [8] [12]. In particular, we will restate Papadimitriou’s result that a computational version of Brouwer fixed-points is in his class **PPAD** [12].

2.1 Simple Stochastic Games

Suppose we have a directed graph $G = (V, E)$ in which each vertex has outdegree of at most two, with the exception of two special vertices, $\mathbb{O}, \mathbb{I} \in V$ called the *0-sink* and the *1-sink*, respectively, each of which have out-degree 0. Further suppose G has its vertex set partitioned as follows:

$$V = MIN \cup MAX \cup AVG \cup \{\mathbb{O}\} \cup \{\mathbb{I}\}$$

A *simple stochastic game* is played on this graph by two players, known as the *0-player* and the *1-player*. A token is placed on a designated *start vertex*, $start \in V$, and it is moved across the edges of the graph in the following manner: if the token is at $v \in MAX$, the 1-player chooses which of the neighbors of v to move the token across to. Similarly, from $v \in MIN$, the 0-player chooses which neighbor of v to move the token to, while at $v \in AVG$, the two players flip a fair coin to decide which neighbor will be moved to. The objective of the 1-player is to reach the 1-sink, whereas the objective of the 0-player is simply to confound the 1-player (e.g. either by arriving at the 0-sink *or* by keeping the vertex moving in some loop).

We may consider a *strategy* of some player to be a subset of the edges; specifically, a strategy of the 1-player is defined by $\sigma \subseteq E$ containing for each $v \in MAX$ precisely one edge out of v , and strategies for the 0-player are defined similarly, containing one edge for each $u \in MIN$. From a vertex $v \in MAX$, the 1-player chooses to move across the unique $e \in \sigma$ out of v .




Figure 1: Example MIN/MAX/AVG circuit

Given a pair of strategies, σ and τ we can assign a *value* to each vertex, $value(v, \sigma, \tau)$, the probability that the 1-sink is reached from v given the 1-player plays by σ and the 0-player plays by τ . We can now define the *value of the game* starting from $start \in V$ to be

$$value(start) = \max_{\sigma} \min_{\tau} value(start, \sigma, \tau)$$

e.g., the probability that the 1-player wins in an optimal strategy against a best-response of the 0-player. These values were studied extensively by Condon [4], and we follow her definition of the problem SSG-VALUE($G, start$): “in the simple stochastic game on G , is $value(start) \geq 1/2$?”

2.2 MIN/MAX/AVG Circuits

Consider a circuit C of MIN/MAX/AVG gates, with feedback *allowed*, each gate having fan-in two, and *inputs* to the circuit hard-wired to either constant 0 or constant 1. (see figure 1) We let $|C|$ denote the number of gates in C , and we assume that the output of each gate is some real number from the closed interval $[0, 1]$. We will frequently discuss *value vectors*, $\vec{x} \in [0, 1]^{|C|}$, which represent a setting for each gate output— $x_i \in [0, 1]$ denotes the output of the i th gate. Notice that this is an implicit ordering of the gates of C . We *do not* require that the output of the i th gate be the respective function of its inputs, but when this is so, e.g. if the i th gate is a MAX gate with inputs from gates j and k , and $x_i = \max\{x_j, x_k\}$, then we say that the gate is *satisfied* by the assignment \vec{x} .

We will frequently discuss *update functions*, $F : [0, 1]^{|C|} \rightarrow [0, 1]^{|C|}$, particularly the following two varieties: *circuit-wide* update functions and *single-gate* update functions. A circuit-wide update function F takes \vec{x} to $F(\vec{x}) = \vec{y}$, where if the i th gate has inputs from the j th and k th gates, then

$$\vec{y}_i = \begin{cases} \max\{x_j, x_k\} & \text{if the } i\text{th gate is a MAX gate} \\ \min\{x_j, x_k\} & \text{if the } i\text{th gate is a MIN gate} \\ (x_j + x_k)/2 & \text{if the } i\text{th gate is an AVG gate} \end{cases}$$

A single-gate update function, F_i is the identity function on each component $j \neq i$, but acts like the circuit-wide update function on the i th component. We will call iterating over the single gate update functions “gate-by-gate update,” and will iterate over them in the order of the gate indices.




Figure 2: Two steps into updating from $F(\vec{v})$ to $F(F(\vec{v}))$ in the original, gate-by-gate version (left) and in our circuit-wide construction (right). The dashed lines would represent self-loops in the circuit if present, and the labels represent the values output by the gates.

We have introduced these two notions primarily for convenience. In some ways, the circuit-wide update function is more natural, as it does not require an update order to be imposed on the circuit. On the other hand, gate-by-gate update simplifies our constructions, and makes verification of their correctness much easier.

As an aside, one should note that it is easy to convert an n -gate construction that works under gate-by-gate update into an n^2 gate construction that is correct under circuit-wide update: consider the n^2 gates arranged in a square array, with the i th gate along the diagonal of the same type as the i th gate in the original circuit, and the off-diagonal gates of arbitrary types. If gate i in the original circuit had inputs from gates j and k , then gate (i, i) in our new circuit should take inputs from gates $(i - 1 \bmod n, j)$ and $(i - 1 \bmod n, k)$. Each other gate (i, j) should take both its inputs from gates $(i - 1 \bmod n, j)$, so that it acts as a one-step delay. See figure 2 for an example. For any value vector on the original circuit that followed from some round of gate-by-gate update, there is a corresponding value vector on this new circuit, with the i th gate along the diagonal taking the same value as the i th gate in the original circuit, every gate above the i th gate in its column taking the same value as the i th gate had on the previous round, and every gate below the i th gate taking the same value as the i th gate. It is easy to see that after we apply the circuit-wide update function k times, the value vector we would obtain from the k th row of this array is the same as if we had applied single-gate update functions 1– k to the original configuration.

We say that a value vector is *stable* iff every gate is satisfied; equivalently, the stable vectors are the fixed points of our circuit-wide update function F . Since each component of F is clearly a continuous function, F itself is continuous and maps $[0, 1]^{|C|}$ into itself, and hence by Brouwer's fixed-point theorem, fixed points (stable vectors) exist. It should be clear that the fixed points of the circuit-wide update function and the gate-by-gate update function are the same. We will be particularly interested in the *minimum stable solution*, \vec{m} , which we can define as

$$\vec{m} = \lim_{n \rightarrow \infty} (F|_{C|} \circ \dots \circ F_1)^n(\vec{0})$$

where $\vec{0}$ is the all-zero vector—that is, gate-by-gate update from $\vec{0}$ yields \vec{m} in the limit.

Claim 1. \vec{m} is well-defined.

Proof. Consider the sequence of vectors, $\vec{v}^i = (F_{|C|} \circ \dots \circ F_1)^i(\vec{0})$. Since $\forall i, k, v_k^i \in [0, 1]$, it follows by induction that $\forall k \forall i > j, v_k^i \geq v_k^j$. Then, $\forall i, k$, since $v_k^i \leq 1$ and $\{v_k^i\}_{i=1}^\infty$ is monotone, it has a limit. Thus, $\{v^i\}_{i=1}^\infty$ has a limit as well. \square

\vec{m} is a minimum in the following strong sense

Lemma 1. $\forall \vec{s} \in [0, 1]^{|C|}$, if \vec{s} is stable, then $\forall k, m_k \leq s_k$.

Proof. Consider the sequence of vectors, $\vec{v}^i = (F_{|C|} \circ \dots \circ F_1)^i(\vec{0})$. Since $\forall i, k, v_k^i \in [0, 1]$, it follows by induction that $\forall k \forall i > j, v_k^i \geq v_k^j$. Then, $\forall i, k$, since $v_k^i \leq 1$ and $\{v_k^i\}_{i=1}^\infty$ is monotone, it has a limit. Thus, $\{v^i\}_{i=1}^\infty$ has a limit as well. \square

We now define the function problem STABLE-CIRCUIT(C) to be “given a circuit C , find the minimum stable solution \vec{m} .” We will soon see that it is sufficient to specify answers as either a vector of $4|C|$ -bit approximations to the components of \vec{m} or as a mapping from the MIN/MAX gates to LEFT/RIGHT projection gates, projecting the input wire from which the MIN or MAX gates take their output in \vec{m} .

2.3 Equivalence of Simple Stochastic Games and MIN/MAX/AVG Circuits

We show that the definitions given in the previous two sections are equivalent in several ways. In particular, suppose we are given a graph G defining a Simple Stochastic Game, and we label the vertices in $V(G) - \{\mathbb{O}, \mathbb{I}\}$ by natural numbers. We will show that such a simple stochastic game defines an equivalent MIN/MAX/AVG circuit. Moreover, we can also show that the decision problem SSG-VALUE is polynomial-time equivalent to the function problem STABLE-CIRCUIT.

2.3.1 Translation Between Circuits and Games

Given a MIN/MAX/AVG circuit C , we can construct the equivalent Simple Stochastic Game, $G(C)$ as follows:

$$\begin{aligned}
 MIN &= \{i : \text{gate } i \text{ is a MIN gate}\} \\
 MAX &= \{i : \text{gate } i \text{ is a MAX gate}\} \\
 AVG &= \{i : \text{gate } i \text{ is a AVG gate}\} \\
 E &= \{(i, j) : \text{gate } i \text{ has an input from gate } j\} \\
 &\cup \{(i, \mathbb{O}) : \text{gate } i \text{ has a constant input of } 0\} \\
 &\cup \{(i, \mathbb{I}) : \text{gate } i \text{ has a constant input of } 1\}
 \end{aligned}$$

Likewise, it is clear that given a Simple Stochastic Game G , if we label every $v \in V(G) - \{\mathbb{O}, \mathbb{I}\}$ by the integers $1, 2, \dots, |V(G) - \{\mathbb{O}, \mathbb{I}\}|$, then we can construct the circuit by creating a gate for each node (where the type of each gate is specified by which of MIN , MAX , or AVG contains the node) and assigning inputs for each edge leaving the node. The circuit we obtain in this way from $G(C)$ is equivalent up to relabeling to the original circuit C . These functions should easily be logspace-computable.

2.3.2 Equivalence of Values and Minimum Stable Solutions

We define \overrightarrow{value} such that if $value(i)$ denotes the value of the simple stochastic game G starting from i , then $value_i = \overrightarrow{value}(i)$. We will also let $\overrightarrow{value}^{\sigma,\tau}$ be the vector such that $value_i^{\sigma,\tau} = value(i, \sigma, \tau)$. We now further introduce “strategies” into our circuits—a strategy σ for the MAX (or MIN) gates of a circuit C is a strategy of the 1-player (or 0-player, respectively) in $G(C)$. We apply the strategy to the circuit C as follows: suppose the i th gate is a MAX gate, and $(i, j) \in \sigma$. Then, we define the single-gate update function under σ , F_i^σ , such that $(F_i^\sigma(v))_i = v_j$. F_i^σ acts as the identity on all other gates, and for any gate k that is not a MAX gate, $F_k^\sigma = F_k$. In this way, we have informally assigned our MAX gates to work as LEFT or RIGHT projection gates according to σ . We can now define gate-by-gate update under σ as $F_{|C|}^\sigma \circ \dots \circ F_1^\sigma$, and the minimum stable solution under σ , \vec{m}^σ as

$$\vec{m}^\sigma = \lim_{n \rightarrow \infty} (F_{|C|}^\sigma \circ \dots \circ F_1^\sigma)^n(\vec{0})$$

just as before. The circuit-wide update function under σ , F^σ , is also defined in the natural way. Similarly, we can define the update functions for a *pair* of strategies, σ and τ for the 1-player and 0-player, respectively. Notice that in this case, we assign every MAX and MIN gate to act as a LEFT or RIGHT projection gate. Consequently, we can write

$$F^{\sigma,\tau}(\vec{v}) = Q\vec{v} + \vec{b}$$

where $Q \in \{0, 1/2, 1\}^{|C| \times |C|}$ and $\vec{b} \in \{0, 1/2, 1\}^{|C|}$. To be precise,

$$Q(i, j) = \begin{cases} 1 & \text{if gate } i \text{ is acting as a projection of gate } j \\ 1/2 & \text{if gate } i \text{ is an AVG gate that takes gate } j \text{ as an input} \\ 0 & \text{otherwise} \end{cases}$$

and

$$b_i = \begin{cases} 1 & \text{if gate } i \text{ acts as a projection of an input of constant 1 or has both inputs of constant 1} \\ 1/2 & \text{if gate } i \text{ is an AVG gate taking one input of constant 1} \\ 0 & \text{otherwise} \end{cases}$$

We will need to consider separately gates in the circuit which take input (directly or indirectly) from hard-wired constants and gates in the circuit which do not.

Claim 2. *For any gate i which is not directly or indirectly connected to a hard-wired input, $m_i = 0$.*

Proof. First, consider the set $D = \{i : \text{gate } i \text{ is not connected to a hard-wired input}\}$. Observe that $\forall i \in D$, if i takes its inputs from gates j and k , then if $j \notin D$ or $k \notin D$, i would be connected to a hard-wired input, which would be a contradiction. Observe that since \vec{m} is obtained by gate-by-gate update from $\vec{0}$, initially every gate in D outputs zero.

Now, let i be the gate in D whose output first rises above zero under gate-by-gate update. Since the inputs to i , j and k , are also in D , they must still update zero on the application of F_i which sets the output of i above zero. But now regardless of the type of gate, i will still output zero after the application of F_i , a contradiction. \square

This reduces the difficulty of computing \vec{m} to finding m_i for those gates i which are connected to a hard-wired input. Observe that, if we replace wires from gates in D with hard-wired zero inputs, we do not change \vec{m} . It should therefore be clear that WLOG we can consider every gate in the circuit to be connected to a hard-wired input, by making the above modification and then removing the gates in D from the circuit.

Similarly, in the case of a simple stochastic game G ,

Claim 3. For any $i \in V(G)$ which has no path to \mathbb{O} or \mathbb{I} , $value_i = 0$.

Proof. Observe that if we consider the set of edges chosen during a simple stochastic game, these edges form a connected path. Thus, given the token is at vertex i , there must exist a path from i to \mathbb{I} for the token to reach \mathbb{I} . Since no such path exists, there are no outcomes for the game in which the token reaches \mathbb{I} , and thus by definition, $value_i = 0$. \square

We can perform a modification on games that is similar to what we have done for circuits above, removing those vertices not connected to any sink and replacing edges to them by edges to the 0-sink. Again, it is clear that this does not alter the values of the remaining nodes. Moreover, the reader should note that for any pair of an equivalent circuit and game (as described in the previous section) the circuit and game we obtain by removing nodes not connected to any hard-wired input and vertices with no path to any sink (respectively) are still equivalent.

The following important lemma is due to Condon [4]:

Lemma 2. Let strategies σ and τ be given, and let Q and \vec{b} be as stated above for the circuit with all gates connected to a hard-wired input.

Then, $\overrightarrow{value}^{\sigma, \tau}$ is the unique solution to the equation $\overrightarrow{value}^{\sigma, \tau} = Q\overrightarrow{value}^{\sigma, \tau} + \vec{b}$. Also, $I - Q$ is invertible, all entries of $(I - Q)^{-1}$ are non-negative, and the entries along the diagonal are strictly positive.

Proof. We first observe that for any AVG vertex i with edges to vertices j and k ,

$$\overrightarrow{value}_i^{\sigma, \tau} = \frac{1}{2}(\overrightarrow{value}_j^{\sigma, \tau} + \overrightarrow{value}_k^{\sigma, \tau})$$

since whenever the token is at vertex i , it then reaches vertex j with probability $1/2$ and similarly reaches vertex k with probability $1/2$. Likewise, for any MIN or MAX vertex i for which (i, j) is the unique edge leaving i in $\sigma \cup \tau$, observe that

$$\overrightarrow{value}_i^{\sigma, \tau} = \overrightarrow{value}_j^{\sigma, \tau}$$

since whenever the token reaches vertex i , under the strategies σ and τ , the token subsequently moves to vertex j . Since we have defined Q and \vec{b} via equations matching these conditions, it should be clear that $\overrightarrow{value}^{\sigma, \tau}$ satisfies

$$\overrightarrow{value}^{\sigma, \tau} = Q\overrightarrow{value}^{\sigma, \tau} + \vec{b}$$

Thus, we can show that $\overrightarrow{value}^{\sigma, \tau}$ is unique by showing that $(I - Q)$ is invertible. We will proceed to this end by showing that $\lim_{\ell \rightarrow \infty} Q^\ell = 0$. Then, the result follows directly since 1 cannot be an eigenvalue of Q , and hence $(I - Q)$ is invertible.

Observe that $Q(i, j)$ gives the probability of reaching vertex i from vertex j in one step; it is then clear that $(Q^k)(i, j)$ is the probability of reaching vertex i from vertex j in precisely k steps. Now, we have required that there be a path to some sink from every vertex in G , and clearly this path has length at most n . Hence, since some sequence of at most n steps reaches a sink from vertex i , the sum of any row of Q^n is at most $1 - 1/2^n$.

It follows by induction that the sum of any row of Q^{nk} is at most $(1 - 1/2^n)^k$; it also follows by induction that the entries of Q^{nk} must be non-negative. Thus, we see that for any (i, j) , $\lim_{\ell \rightarrow \infty} Q^\ell(i, j) = 0$, which yields the desired result.

Now, since $(I - Q)^{-1}$ exists, note that it must be given by $\sum_{\ell=0}^{\infty} Q^\ell$. As we observed above, every $Q^\ell(i, j)$ is non-negative, and hence the sum must be non-negative as well. Also, in $Q^0 = I$, we observe that the entries along the diagonal are positive, and hence the entries in the sum must be positive as well. \square

Together with claim 3, this implies that given a pair of strategies σ and τ , in general $\overrightarrow{value}^{\sigma,\tau}$ is unique, since we know the values of the vertices with no path to a sink to be zero, and we have shown that the rest of the vertices have unique values.

Observe now that we can recover a “strategy” of the MAX gates of the circuit from \vec{m} , $\sigma(\vec{m})$, in the natural way—the 1-player in $G(C)$ takes an edge $(i, j) \in \sigma(\vec{m})$ iff for the MAX gate i , $m_i = m_j$ (we can resolve ambiguities by using the first such j for σ), or takes an edge to \circledast or I iff the MAX gate i takes its value from a 0 or 1 input, respectively. Clearly, this puts an edge for each MAX vertex in σ , so it is a valid strategy of the 1-player in $G(C)$. We can similarly recover strategies of the MIN player from \vec{m} . It should be clear that by lemma 2 and claim 2, $\vec{m} = \vec{m}^{\sigma(\vec{m}),\tau(\vec{m})}$.

Lemma 3. *Given σ , a strategy of the 1-player for $G(C)$, and a value vector for any strategy of the 0-player (when the 1-player plays according to σ), \vec{u} , at each i th gate of C , $m_i^\sigma \leq u_i$.*

Proof. First, observe that in every i th component, $0 \leq u_i$. It follows easily by induction that for all n , $((F_{|C|}^\sigma \circ \cdots \circ F_1^\sigma)^n(\vec{0}))_i \leq u_i$: we assumed that the 1-player played by σ in \vec{u} and the value of an AVG vertex in $G(C)$ must be the average of the values at the vertices it has edges to in any valid value vector. Thus, the only values which might be altered by any F_j^σ might be at MIN gates, which could only decrease from u_j . Therefore, now, we see that

$$m_i^\sigma = \left(\lim_{n \rightarrow \infty} (F_{|C|}^\sigma \circ \cdots \circ F_1^\sigma)^n(\vec{0}) \right)_i \leq u_i$$

as promised. □

Of course, $\tau(\vec{m}^\sigma)$ achieves $\overrightarrow{value}^{\sigma,\tau(\vec{m}^\sigma)} = \vec{m}^\sigma$ by lemma 2 since both are fixed points of $F^{\sigma,\tau(\vec{m}^\sigma)}$ in which the disconnected entries are all set to zero. It follows easily that for every i th gate, $m_i^\sigma = \min_\tau \text{value}(i, \sigma, \tau)$.

Lemma 4. $\vec{m} = \overrightarrow{value}$

Proof. Let any σ , a strategy of the 1-player, be given. Now, since in every i th component, $0 \leq m_i$, we find by induction on n that for every i th component, $((F_{|C|}^\sigma \circ \cdots \circ F_1^\sigma)^n(\vec{0}))_i \leq m_i$ since σ could do no better than to select the maximum at each MAX gate, and otherwise $F_j^\sigma(\vec{m}) = m_j$. Thus, $m_i^\sigma \leq m_i$. Of course, again, if the 1-player uses $\sigma(\vec{m})$, then $\vec{m} = \vec{m}^{\sigma(\vec{m}),\tau(\vec{m})}$ by lemma 2 and claim 2.

Now, for each i , consider σ^* such that $m_i^{\sigma^*} = \max_\sigma m_i^\sigma$. Since we have argued that $m_i^{\sigma^*} \leq m_i$, it should be clear that $\sigma(\vec{m})$ achieves the same value as σ^* , and hence $m_i = \max_\sigma m_i^\sigma$. Moreover, any $m_i^\sigma = \min_\tau m_i^{\sigma,\tau}$ by lemma 3. Thus, by lemma 3 and our definition of \overrightarrow{value} , we further find

$$m_i = \max_\sigma \min_\tau \text{value}_i^{\sigma,\tau} = \text{value}_i$$

□

Thus, henceforth we shall discuss simple stochastic games and MIN/MAX/AVG circuits interchangeably.

2.3.3 Polynomial-Time Equivalence of SSG-VALUE and STABLE-CIRCUIT

Now that we have identified \vec{m} with \overrightarrow{value} , the reduction from SSG-VALUE(G, i) to STABLE-CIRCUIT is trivial: compute STABLE-CIRCUIT($C(G)$), and decide “yes” iff $m_i \geq 1/2$. The other direction is a little more interesting. We state another lemma from Condon [4], one that we have implicitly been using for some time:

Lemma 5. *The value of a simple stochastic game G with n non-sink vertices is of the form p/q , where p and q are integers, $0 \leq p, q \leq 4^n$.*

Proof. As noted by Condon, this proof is quite similar to a proof originally by Gill. Let any pair of strategies σ and τ be given. Recall that by claim 3 that the values of the vertices which have no path to any sink must be zero, where this can be written as $p/q = 0/1$. Thus, we need only show that the property holds for the vertices which have a path to a sink.

We consider the game G' we obtain by removing the vertices with no path to any sink, and replacing edges to these vertices with edges to the zero sink. We observe that for any remaining vertex i , this does not change $m_i^{\sigma, \tau}$. Now, as argued in lemma 2, the values of G' are the unique solution to $2(I - Q)\vec{value}^{\sigma, \tau} = 2\vec{b}$, where Q and \vec{b} are as discussed in the previous section; note that the entries of $2(I - Q)$ and $2\vec{b}$ are taken from the set $\{0, \pm 1, \pm 2\}$. Observe that if we sum the absolute values of any i th row of $2(I - Q)$ with $2|b_i|$, the total is at most 4. It also follows that the sum of the absolute values of any subset of the entries for any given row is at most 4.

Now, it follows by Cramer's rule that each $value_i^{\sigma, \tau}$ can be written as the quotient of two integers, N/D where $D = \det(2(I - Q))$ and N is the determinant of the matrix obtained by replacing the i th column of $2(I - Q)$ with $2\vec{b}$. We can argue that $N \leq 4^n$ and $D \leq 4^n$ by induction: we observed above that in every row in these matrices, the sum of the absolute values of the entries is at most 4. It's clear that in any 1×1 minor, the determinant is at most 4, whereas in an $\ell \times \ell$ matrix with this property, when we expand along the first row each of the $(\ell - 1) \times (\ell - 1)$ minors also has this property. If, by IH, the minors have determinants bounded by $4^{\ell-1}$, then since the sum of the absolute values of the entries is at most 4, the determinant of the $\ell \times \ell$ matrix is at most $4 \cdot 4^{\ell-1}$, so the IH holds.

Thus, since $|C(G')| \leq n$, we see $value_i^{\sigma, \tau} = N/D$ where we have $N \leq 4^{|C(G')|} \leq 4^n$ and $D \leq 4^{|C(G')|} \leq 4^n$ as needed. \square

We interpret this as stating that $4|C|$ -bit approximations to the entries of \vec{m} are sufficient to recover the precise values of \vec{m} . Notice that any two entries, $m_i = p_i/q_i$ and $m_j = p_j/q_j$ which differ must have $|m_i - m_j| \geq 1/q_i q_j \geq 1/2^{4|C|}$. Hence, we can tell which is greater from the first $4|C|$ bits, allowing us to recover $\sigma(\vec{m})$ and $\tau(\vec{m})$, where by lemma 2, \vec{m} is the unique solution to

$$\vec{m} = F^{\sigma(\vec{m}), \tau(\vec{m})}(\vec{m}) = Q\vec{m} + \vec{b}$$

for an easily computable Q and \vec{b} . Thus, by Gaussian elimination, we can efficiently recover the precise values of \vec{m} by solving for \vec{m} .

The following simple construction will be useful to us:

Claim 4. *Let $0.b_k \dots b_1$ be the "binary decimal" representation of a constant in the interval $[0, 1)$. We can compute this constant using k AVG gates.*

Proof. The constant $0.b_1$ is easy to compute using a single gate: use 0 and $b_1 \in \{0, 1\}$ as the inputs to an AVG gate. Inductively suppose we have computed $0.b_i \dots b_1$ using i AVG gates. Now, suppose we pass the output of the AVG gate that computes this value as the first input to an AVG gate that has b_{i+1} as its second input. Clearly, the output is $0.b_{i+1}b_i \dots b_1$, and we have only used $i + 1$ AVG gates, so the claim holds. \square

Now, supposing that we could decide $\text{SSG-VALUE}(G, i)$ efficiently, consider the following algorithm for STABLE-CIRCUIT :

Algorithm. On input C ,

1. For $i = 1, \dots, |C|$, repeat step 2:




Figure 3: Extracting a bit of m_i

2. For $k = 1, \dots, 4|C|$, repeat steps 3 and 4:
3. If $k > 1$, compute C_k , which is identical to C except that there is an AVG gate, j , which takes the i th gate as its first input and the constant $1 - 1/2^k - 0.v_1 \dots v_{k-1}$ as its second input. If $k = 1$, $C_k = C$, and $j = i$.
4. Put $v_k = 1$ if $\text{SSG-VALUE}(G(C_k), j)$ is a “yes” instance, and put $v_k = 0$ otherwise. Output v_k .

Claim 5. $v_1 \dots v_{4|C|}$ as output during the i th iteration of step 1 are the first $4|C|$ bits of m_i .

Proof. It is now evident that $\lim_{c \rightarrow \infty} |\vec{m} - \vec{m}'|_\infty = 0$, so in particular, for any given ε we may pick a c such that $|\vec{m} - \vec{m}'|_\infty \leq \varepsilon$, as desired. Suppose we consider the binary decimal representation of $m_i = 0.b_1b_2\dots$. When $k = 1$ in the i th iteration of step 1, clearly $v_1 = b_1$, since $\text{SSG-VALUE}(G, i)$ is a “yes” instance iff $b_1 = 1$. Now, suppose inductively that at the k th iteration of step 2, $\forall j < k$, $v_j = b_j$. Then,

$$\begin{aligned} \frac{1}{2}(1 - 1/2^k - 0.v_1 \dots v_{k-1} + 0.b_1b_2 \dots b_{k-1}b_k \dots) &= \frac{1}{2}(0.\overbrace{1 \dots 1}^k + 0.0 \dots 0 \overbrace{b_k \dots}^{k-1}) \\ &= \frac{1}{2}(b_k \dots) \\ &= 0.b_k \dots \end{aligned}$$

so clearly the output of the AVG gate in C_k is greater than or equal to $1/2$ iff $b_k = 1$, and hence $v_k = b_k$. \square

It is also clear that the algorithm runs in polynomial time, given that our algorithm for SSG-VALUE does. Thus, we see that SSG-VALUE is a decision problem that is polynomial-time equivalent to the function problem STABLE-CIRCUIT.

2.4 Discounted Games

Much of the previous work on Simple Stochastic Games is founded on transforming a general instance of a SSG G into a closely related “discounted game” that halts with probability 1. This is so because such discounted games have desirable properties, notably the property that such games have a unique stable value vector. This value vector is, of course, \vec{m} for the discounted game. The remarkable fact is that we may ensure that the minimum stable solution for this game is arbitrarily close to the minimum stable solution of the original game.

Let any SSG G be given. We construct the $1/2^m$ discounted game G' by replacing each edge $(i, j) \in E(G)$ with a sequence of m AVG vertices as seen in figure 4. To be precise, we add




Figure 4: Replacing an edge (i, j) with a sequence of m AVG vertices in the discounted game construction.

vertices $(i, j, 1), (i, j, 2), \dots, (i, j, m)$ to the set AVG . We then include an edge $((i, j, \ell), j)$ for each $\ell \in \{1, \dots, m\}$, an edge $(i, (i, j, 1))$, an edge $((i, j, m), \mathbb{O})$, and an edge $((i, j, \ell), (i, j, \ell + 1))$ for each $\ell \in \{1, \dots, m - 1\}$. It should be clear that since in crossing from any vertex i to any other vertex j the game ends with probability at least $1/2^m$, the token will reach a sink with probability 1.

This construction leads us to the following result due to Shapley:

Lemma 6. *If G' is a $1/2^m$ discounted game corresponding to G , then $C(G')$ has a unique stable solution.*

Proof. First, observe that for any gate in $C(G)$, if we set the output of the corresponding gate in $C(G')$ to be x , then it can be shown by induction that in a stable solution, the output of the ℓ th AVG gate chained to it will be $x \cdot (1 - 1/2^\ell)$.

Now, given the circuit-wide update function for $C(G)$, F , observe that we can define a circuit-wide update function for the gates of $C(G')$ that correspond to gates in $C(G)$ as $F'(\vec{v}) = (1/2^m)F(\vec{v})$, e.g., where $F' : [0, 1]^{|C(G)|} \rightarrow [0, 1]^{|C(G)|}$. Since we observed above that the rest of the gates of $C(G')$ will have their stable values uniquely determined by the gates operated on by F' , where it is clear that any fixed point of F' is equivalent to a stable solution of $C(G')$, we find that $C(G')$ has a unique stable solution iff F' has a unique fixed point.

Suppose for contradiction that \vec{v} and \vec{w} are fixed points of F' such that $\vec{v} \neq \vec{w}$. Consider $c = |\vec{v} - \vec{w}|_\infty = \max_i |v_i - w_i|$; clearly $c > 0$. We also see that

$$|\vec{v} - \vec{w}|_\infty = |F'(\vec{v}) - F'(\vec{w})|_\infty = (1 - 1/2^m)|F(\vec{v}) - F(\vec{w})|_\infty$$

Now, for each gate i with inputs from gates j and k , observe that since $|v_i - w_i| \leq c$,

$$|F(\vec{v})_i - F(\vec{w})_i| = \left\{ \begin{array}{ll} |\max\{v_j, v_k\} - \max\{w_j, w_k\}| & \text{if } i \text{ is a MAX gate} \\ |\min\{v_j, v_k\} - \min\{w_j, w_k\}| & \text{if } i \text{ is a MIN gate} \\ |\frac{1}{2}(v_j + v_k) - \frac{1}{2}(w_j + w_k)| & \text{if } i \text{ is a AVG gate} \end{array} \right\} \leq c$$

This yields

$$c = |\vec{v} - \vec{w}|_\infty = (1 - 1/2^m)|F(\vec{v}) - F(\vec{w})|_\infty \leq (1 - 1/2^m)c$$

which is of course a contradiction, since $c > 0$. □

We now present the most important property of these discounted games. This proof is due to Condon, and is a discrete version of a result by Filar, that the value of G is the limit of the value of the β discounted game corresponding to G , as $\beta \rightarrow 0$.

Lemma 7. $\forall \varepsilon > 0 \exists c > 0$ such that if G is a simple stochastic game with n vertices, the value of G is within ε of the value of the $(1/2^{cn})$ discounted game corresponding to G .

Proof. WLOG, suppose we have removed all vertices with no path to any sink from G ; it should be clear that in both G and G' , these vertices all have value zero. As done in the proof of lemma 6, we will consider F' , the circuit-wide update function for the gates of $C(G')$ that correspond to gates in $C(G)$. It is clear that since $C(G)$ and $C(G')$ have corresponding sets of MIN and MAX gates, the two circuits have corresponding sets of strategies. We will fix $\sigma = \sigma(\vec{m})$ and $\tau = \tau(\vec{m})$. Now, observe that for Q and \vec{b} as in lemma 2, $\vec{m} = F^{\sigma, \tau}(\vec{m}) = Q\vec{m} + \vec{b}$. We will now put \vec{m}' as the unique solution to $\vec{m}' = F'^{\sigma, \tau}(\vec{m}') = (1 - 1/2^{cn})Q\vec{m}' + (1 - 1/2^{cn})\vec{b}$, e.g., the minimum stable solution of $C(G')$ under strategies σ and τ , restricted to gates with corresponding gates in $C(G)$. Now, since Q is invertible, we find $\vec{m} = (I - Q)^{-1}\vec{b} = \sum_{\ell=0}^{\infty} Q^{\ell}\vec{b}$, and similarly that $\vec{m}' = \sum_{\ell=0}^{\infty} (1 - 1/2^{cn})^{\ell+1} Q^{\ell}\vec{b}$. We then find that

$$\vec{m} - \vec{m}' = \sum_{\ell=0}^{\infty} (1 - (1 - 1/2^{cn})^{\ell+1}) Q^{\ell}\vec{b}$$

It is well known that in particular for the ℓ_{∞} -norm, given any matrix M and vector \vec{v} , $|M\vec{v}|_{\infty} \leq |M|_{\infty}|\vec{v}|_{\infty}$. Of course, in this case, the entries of \vec{b} are taken from the set $\{0, 1/2, 1\}$ and hence $|\vec{b}|_{\infty} \leq 1$. Thus, we find

$$|\vec{m} - \vec{m}'|_{\infty} \leq \sum_{\ell=0}^{\infty} (1 - (1 - 1/2^{cn})^{\ell+1}) |Q^{\ell}|_{\infty}$$

Notice that we can rewrite $(1 - (1 - 1/2^{cn})^{\ell+1})$ as a telescoping sum; it is then evident that $(1 - (1 - 1/2^{cn})^{\ell+1}) \leq (\ell + 1)/2^{cn}$. Recall that in the proof of lemma 2, we showed that $|Q^{nk}|_{\infty} \leq (1 - 1/2^n)^k$ due to our assumption that all vertices have a path to a sink. We therefore find that

$$|\vec{m} - \vec{m}'|_{\infty} \leq \sum_{\ell=0}^{\infty} (1/2^{cn})(\ell + 1)(1 - 1/2^n)^{\lfloor \ell/n \rfloor}$$

We here observe that since (in general) $n \geq 1$,

$$\sum_{\ell=0}^{\infty} (\ell+1)(1-1/2^n)^{\lfloor \ell/n \rfloor} \leq \int_0^{\infty} (x+1)(1-1/2^n)^{x/n-1} dx = \frac{1}{1-1/2^n} \left(\frac{n}{\ln(1-1/2^n)} \left(\frac{n}{\ln(1-1/2^n)} - 1 \right) \right)$$

so we have

$$|\vec{m} - \vec{m}'|_{\infty} \leq (1/2^{cn}) \frac{1}{1-1/2^n} \left(\frac{n}{\ln(1-1/2^n)} \left(\frac{n}{\ln(1-1/2^n)} - 1 \right) \right)$$

It is now evident that $\lim_{c \rightarrow \infty} |\vec{m} - \vec{m}'|_{\infty} = 0$, so in particular, for any given ε we may pick a c such that $|\vec{m} - \vec{m}'|_{\infty} \leq \varepsilon$, as desired. Moreover, since the above argument holds for any pair of strategies σ and τ , by picking c sufficiently large, it is clear that the corresponding value vector $\vec{m}'^{\sigma, \tau}$ is optimal only if $\vec{m}^{\sigma, \tau}$ is, and hence \vec{m}' must be the minimum stable solution of $C(G')$. \square

Now, the following is also discussed at length in the introductory papers [4] [2], and briefly mentioned in Condon's SSG Algorithms paper [5]:

Claim 6. Any instance of a simple stochastic game G can be efficiently transformed into a game G' such that $C(G')$ has a unique stable solution that is within any desired polynomial number of bits of the minimum stable solution of $C(G)$. $|C(G')|$ is also polynomially related to $|C(G)|$.

Proof. It is clear that we can compute the value c such that the $1/2^{cn}$ stopping game G' has values within ℓn^k bits of those of G via the formula

$$c = \left\lceil \frac{1}{n} \left(\lg \left(\frac{n^2}{\ln^2(1 - 1/2^n)} - \frac{n}{\ln(1 - 1/2^n)} \right) - \lg(1 - 1/2^n) \right) + \ell n^{k-1} \right\rceil$$

in time polynomial in the size of G .

It is likewise clear that once we have computed c , we could easily output the $1/2^{cn}$ stopping game corresponding to G in polynomial time. By lemma 6, we know that this stopping game has a unique stable solution, and by lemma 7, this stable solution is within ℓn^k bits of that of G . Finally, we observe that our construction adds polynomially many vertices on each edge, where there were polynomially many edges before, and hence there are polynomially many vertices in the stopping game. \square

2.5 Known Algorithms

A variety of algorithms for solving SSG-VALUE have been proposed; some of these algorithms are correct. See the SSG Algorithms paper[5] to get a better sense of the difficulties. It is perhaps more compelling to note that whenever we only permit two of the three types of vertices (gates), the problem is solvable in polynomial time.

The case where AVG vertices are omitted was observed to be solvable in polynomial time by Condon [4]. It is trivial, given Lemma 4: simply compute \vec{m} by definition. Since the circuit only contains MIN and MAX gates and every gate has inputs from the set $\{0, 1\}$, and thus will have an output from this set as well. It is easy to see by induction that this will hold on every round. Since the sequence of vectors obtained under gate-by-gate update starting from $\vec{0}$ is monotone, observe that we will fix the output of any gate at most once; thus, we are guaranteed to find a stable vector in at most $|C|$ rounds.

The cases where MIN or MAX gates are omitted are similar, and originally due to Derman [6]. Consider the following linear program for the case without MAX gates: for the circuit corresponding to the discounted game G' that has values within $1/2^{4|C|}$ of G , maximize $\sum_i x_i$ subject to the constraints that for each gate i with inputs from gates j and k ,

$$\begin{aligned} x_i &\leq \frac{1}{2}(x_j + x_k) && \text{if } i \text{ is an AVG gate} \\ x_i &\leq x_j && \text{if } i \text{ is an MIN gate} \\ x_i &\leq x_k && \text{if } i \text{ is an MIN gate} \\ x_i &\leq 1 \end{aligned}$$

(where x_j or x_k is replaced by a constant 0 or 1 if gate i takes a constant input of 0 or 1)

Claim 7. \vec{m} is the unique optimal solution to Derman's linear program.

Proof. First, note that \vec{m} is a feasible solution to Derman's linear program. Now, suppose for contradiction that the optimal solution to the linear program is some vector \vec{v} such that for some gate i with inputs j and k , one of the following cases holds:

Case i. i is a MIN gate and $v_i < v_j$ and $v_i < v_k$.

Consider now the vector \vec{v}' in which $\forall \ell \neq i, v'_\ell = v_\ell$ and $v'_i = \min\{v_j, v_k\}$. Clearly, v'_i still satisfies the constraints where it appears on the LHS, and similarly, since $v'_i > v_i$, all constraints where v'_i appears on the RHS are still satisfied.

Case ii. i is an AVG gate and $v_i < \frac{1}{2}(v_j + v_k)$.

Here, we consider the vector \vec{v}' in which $\forall \ell \neq i, v'_\ell = v_\ell$ and $v'_i = \frac{1}{2}(v_j + v_k)$. Again, note that \vec{v}' still satisfies all of the constraints, and $v'_i > v_i$.

We now observe that in both cases, there is a feasible vector \vec{v}' such that $\sum_i v'_i > \sum_i v_i$, so \vec{v} was not optimal, contradicting the assumption it was. Thus, in the optimal vector \vec{v} , for each AVG gate i with inputs j and k , we have $v_i = \frac{1}{2}(v_j + v_k)$ and for each MIN gate i with inputs j and k , we have $v_i = \min\{v_j, v_k\}$. Since G' is a stopping game, we know that the minimum stable solution is the unique stable solution by lemma 6. Moreover, it is within $4|C|$ bits of the minimum stable solution of G by lemma 7, where we know that finding \vec{m} to this precision is sufficient by lemma 5. \square

Now, since we know that linear programs may be solved in polynomial time due to Khachiyan [9], the discounted game is easily constructed and polynomial in the size of the original game by claim 6, and the transformation from the game to the linear program can clearly be carried out in logspace, we have a polynomial-time algorithm for the case where MAX gates have been omitted. It is clear that the case of the omission of MIN gates is similar.

Observe that Derman's linear program shows how, given a strategy σ , one can compute \vec{m}^σ , and thus how one can find $\tau(\vec{m}^\sigma)$, which we recall that as shown in lemma 3, is the best response to the strategy σ . It turns out that by (cleverly! [5]) choosing a counter-strategy to $\tau(\vec{m}^\sigma)$, one can repeatedly improve the 1-player's strategy until it is stable, as discovered by Hoffman and Karp [7].

Algorithm. On input G , compute the $1/2^{4|V(G)|}$ stopping game G' and pick an arbitrary pair of strategies, σ and τ .

1. While $\vec{m}^{\sigma, \tau}$ is not stable, repeat steps 2–4.
2. For each MAX vertex i with edges to j and k , repeat step 3.
3. If $\vec{m}_i^{\sigma, \tau} < \max\{\vec{m}_j^{\sigma, \tau}, \vec{m}_k^{\sigma, \tau}\}$, switch the edge included in σ .
4. Put $\tau \leftarrow \tau(\vec{m}^\sigma)$ by solving Derman's linear program

Claim 8. Suppose σ' is the strategy obtained from σ in step 2 of the algorithm and $\tau' = \tau(\vec{m}^{\sigma'})$, e.g., as computed in step 4. Then either $|\vec{m}^{\sigma, \tau}|_1 < |\vec{m}^{\sigma', \tau'}|_1$ (where $|\cdot|_1$ is the L_1 norm) or $\vec{m}^{\sigma, \tau}$ was stable.

Proof. Since we are operating on the stopping game G' , clearly lemma 2 applies directly since every vertex connected to the rest of the graph has a path to a sink (and WLOG, G was connected). Hence, $\vec{m}^{\sigma, \tau}$ is the unique solution to $\vec{m}^{\sigma, \tau} = Q\vec{m}^{\sigma, \tau} + \vec{b}$ and $\vec{m}^{\sigma', \tau'}$ is the unique solution to $\vec{m}^{\sigma', \tau'} = Q'\vec{m}^{\sigma', \tau'} + \vec{b}'$. (Where Q and \vec{b} are obtained from σ and τ , and Q' and \vec{b}' are obtained from σ' and τ' .)

Put $\vec{\Delta} = \vec{m}^{\sigma', \tau'} - \vec{m}^{\sigma, \tau}$. We now show that $\forall i \Delta_i \geq 0$, and $\exists i \Delta_i > 0$. From this, it will follow immediately that $|\vec{m}^{\sigma, \tau}|_1 < |\vec{m}^{\sigma', \tau'}|_1$.

Add and subtract $Q'\vec{m}^{\sigma, \tau} + \vec{b}'$ to $\vec{\Delta}$:

$$\vec{\Delta} = (Q'\vec{m}^{\sigma', \tau'} + \vec{b}') - (Q'\vec{m}^{\sigma, \tau} + \vec{b}') + (Q'\vec{m}^{\sigma, \tau} + \vec{b}') - (Q\vec{m}^{\sigma, \tau} + \vec{b})$$

so, if we further put $\vec{\delta} = (Q'\vec{m}^{\sigma,\tau} + \vec{b}') - (Q\vec{m}^{\sigma,\tau} + \vec{b})$, we find $\vec{\Delta} = Q'\vec{\Delta} + \vec{\delta}$.

Now, recall that from lemma 2, we know $(I - Q')$ is invertible, the entries of $(I - Q')^{-1}$ are all non-negative and the entries along the diagonal are strictly positive. Thus, $\vec{\Delta} = (I - Q')^{-1}\vec{\delta}$, and it suffices to show that $\forall i \delta_i \geq 0$ and $\exists i \delta_i > 0$.

Observe that the pairs Q, Q' and \vec{b}, \vec{b}' differ only in rows i where the edge in σ or τ was switched to obtain σ' or τ' . Notice that if the i th rows are equal in Q and Q' and the i th entries of \vec{b} and \vec{b}' are equal, then $\delta_i = 0$.

Since the algorithm would have terminated if $\vec{m}^{\sigma,\tau}$ were stable (where τ is the best response to σ), we see that there must be some MAX vertex i such that (i, k) in σ' has replaced (i, j) in σ , e.g., where $\vec{m}_k^{\sigma,\tau} > \vec{m}_j^{\sigma,\tau}$. Notice that for every such vertex i , $\delta_i = \vec{m}_k^{\sigma,\tau} - \vec{m}_j^{\sigma,\tau} > 0$.

Now, consider any MIN vertex i such that (i, k) in τ' has replaced (i, j) in τ . Since τ was the best response to σ , observe that $\vec{m}_j^{\sigma,\tau} \leq \vec{m}_k^{\sigma,\tau}$ held. Now, we see again that $\delta_i = \vec{m}_k^{\sigma,\tau} - \vec{m}_j^{\sigma,\tau} \geq 0$, and the claim holds. \square

Thus, since the inequality is strict, we see that we cannot examine any pair of strategies more than once during the algorithm. Therefore, since there are only a finite number of pairs of strategies, the algorithm is guaranteed to halt on a stable solution. We observe that by claim 6, this is sufficient since transformed the given game G into a sufficiently close discounted game where there is a unique stable solution.

Notice that we did not argue that this algorithm completes in polynomial time; while it is clear that steps 2–4 in the body of the algorithm all complete in polynomial time, it is not known how many iterations may be required by the algorithm. The best known worst case bound on the Hoffman-Karp algorithm is $O(2^n/n)$, as recently shown by Kumar and Tripathi [10].

2.6 Leapfrog Circuits

Observe that a MIN or MAX gate always acts as a projection of one of its two inputs. It is natural to consider whether some particular gate is ever “forced” to take its left (or right) input; we could imagine watching the circuit during its computation (either under circuit-wide or gate-by-gate update), waiting to see if (say) the value of the left input to a MAX gate is above the right input. Such a state could be considered an “accepting state” for the circuit.

To make use of this idea, it will be helpful to think of the right input as a threshold: suppose that for each gate x in the circuit, there is a threshold gate $th(x)$. Let us say that the gate x outputs T iff the value output by x is greater than the value output by $th(x)$. If we were to then feed x as the left input to a MAX gate and $th(x)$ as the right input in the scenario previously described, we would find that the circuit enters an accepting state iff x encodes T. It should be clear that the gate itself serves no purpose except to motivate this model, so henceforth we shall disregard it. We will call circuits in this model “Leapfrog” circuits, for reasons that we will try to elucidate later.

It will be critical for us to allow different gates of the circuit to be considered relative to different thresholds. Suppose $th(C)$ denotes the set of gates that are used as thresholds in our circuit C . We now define for each $i \in th(C)$ the set of gates $\{x \in C : th(x) = i\}$ to be the i th region of C .

Leapfrog circuits are an attractive model because they allow our monotone MIN/MAX/AVG circuits to exhibit non-monotone behavior. We will see that under the assumption that $\mathbf{NP} \neq \mathbf{PSPACE}$, that we can compute significantly more using these circuits than we could if we only examined stable configurations of the circuits (i.e., solutions to the stable circuit problem). Inspired by the definition of \vec{m} , we define $\text{LEAPFROG}(C, i, th)$ to be the following decision problem: “ $\exists t$ such that for $\vec{v}^t = (F_{|C|} \circ \dots \circ F_1)^t(\vec{0})$, $v_i^t > v_{th}^t$?” where F_i is the i th single-gate update function, as before.

2.7 Function Complexity Classes

We now state the definitions of some relevant subclasses of **TFNP**. We will show, in Section 3.2, that **STABLE-CIRCUIT** is contained in each of the following classes. In the case of **PPAD**, we will require a few preliminary results.

A problem Π is in **PLS**, introduced in [8], if for any instance x of Π , there is an associated set of *feasible solutions*, a *neighborhood structure*, and a *potential function* to be optimized, defined on the feasible solutions, where we solve the instance x by finding a feasible solution for which no neighbor has a “better” value under the potential function. These objects further must have the following characteristics:

1. An initial feasible solution can be computed in polynomial time
2. The value of the potential function for any feasible solution may be computed in polynomial time
3. Given a locally optimal feasible solution, we can determine that the solution is locally optimal in polynomial time
4. Given a feasible solution that is not locally optimal, we can find a neighbor with a “better” value in polynomial time.

Papadimitriou [12] has defined the class **PPAD** to be the set of function problems polynomial-time reducible to a search problem of the following form: suppose there is some Turing machine M such that for any strings x and c where $|c|$ is bounded by polynomial function of $|x|$, $p(|x|)$, $M(x, c)$ outputs in polynomial time a list of at most two strings, whose lengths are also bounded by $p(|x|)$. M on x implicitly defines a directed graph $G(x)$ on the strings of length at most $p(|x|)$ —an edge $(c, c') \in G(x)$ iff c is the first string in $M(x, c')$ and c' is the second string in $M(x, c)$. We are further guaranteed that a particular string is known to be a leaf, where our objective is to find another leaf (source or sink) in the graph—one is known to exist since “all graphs contain an even number of odd-degree nodes.”

The first problem demonstrated by Papadimitriou to be in **PPAD** is k D-SPERNER for $k \geq 2$, (in fact, he has shown that 3D-SPERNER is **PPAD**-complete, but we will not require this) a computational problem associated with Sperner’s lemma. The problem is defined on the k -dimensional hypercube, subdivided into n^k “cubelets” for a given n , where the vertices of the cubelets, $(x_1/n, x_2/n, \dots, x_n/n)$ for $x_1, x_2, \dots, x_n \in \{0, 1, \dots, n\}$, are each assigned one of the colors $0, 1, \dots, k$ by a polynomial-time algorithm M , subject to the following constraints:

1. Let e^1, e^2, \dots, e^n be the vertices such that the j th component of e^i is 1 iff $j = i$ and otherwise the j th component is 0—these are called the *principal vertices*. We require that the vertex e^i takes the i th color and the vertex $e^0 = (0, 0, \dots, 0)$ takes the color 0.
2. Any point on a face containing the vertex e^0 cannot obtain the color of the principal vertex absent from that face.
3. Any point on any face not containing the vertex e^0 cannot obtain the color 0.

We assume that the colors assigned by M are legal, though they are easily checked, and if the color output does not satisfy some constraint, we can easily consider it to be replaced by some standard legal color for the face. The problem is to then find a cubelet that contains all $k + 1$ colors.

We will have cause to “simplicize” the k -dimensional cubelets in a standard way (Papadimitriou engages in a rather thorough discussion of the difficulties and some alternatives, motivating the




Figure 5: “Simplicizing” the square (2-dimensional hypercube)

definition [12]). Informally, this means to carve each k -dimensional cubelet into k -dimensional simplices, while applying a mapping to “pull” the principal vertices outwards, transforming the k -dimensional cube into a k -dimensional simplex: each of the faces of the cube containing the point e^0 forms one face of the simplex, and the remaining faces of the cube (not containing e^0) form the final face of the simplex. See figure 5 for an example of the $k = 2$ case. Naturally, we will never compute this simplicization explicitly; instead, we will generate a simplex along with its adjacent neighbors as-needed, one dimension at a time. Since the k -dimensional simplex will have at most $k + 1$ neighboring simplices, algorithms for performing this generation in time polynomial in k are easily obtained from known techniques [11] by extruding the current set of simplices one dimension at a time, and throwing away any simplices that are not adjacent to the simplex in question after each step.

The fact that some cubelet will contain all $k + 1$ colors will follow from the following **PPAD** algorithm for k D-SPERNER due to Papadimitriou (where a solution is thus guaranteed to exist by the parity argument).

Theorem 1. k D – SPERNER \in **PPAD** for $k \geq 2$.

Proof. Consider the standard k -dimensional simplicization of our k -dimensional cubelets. We next add edges from the vertex e^0 to every other vertex on the $0 - 1 - \dots - (k - 1)$ face. Note that this creates a unique “external” $k - 1$ -dimensional face of the simplex guaranteed to contain the colors $0 - (k - 1)$: the face consisting of the vertices $\{e^0, e^1, \dots, e^{k-1}\}$.

Our graph is implicitly represented by these $k - 1$ -dimensional faces containing the colors $0 - (k - 1)$ or $0 - (k - 2)$, as well as interior k -dimensional simplices containing the colors $0 - k$ or $0 - (k - 1)$, with the former corresponding to leaves of the graph. Notice that a $k - 1$ -dimensional simplex which contains the colors $0 - (k - 2)$ must share a face with some other $k - 1$ -dimensional simplex which also contains those colors. Similarly, each k -dimensional simplex with the colors $0 - (k - 1)$ must share a face with another k -dimensional simplex with those colors. In each case, there is an arc from the former simplex to the latter where the direction of the arc depends on the orientation of the face when the vertices are taken in the order naturally specified by the colors. Furthermore, we provide an arc from each $k - 1$ -dimensional simplex containing the colors $0 - (k - 1)$ to the k -dimensional simplex it resides on. (Clearly, this k -dimensional simplex also contains the colors $0 - (k - 1)$) We define the standard leaf to be the unique “external” $k - 1$ -dimensional face containing the colors $0 - (k - 1)$.

Clearly, since the only vertices of the graph described which are not guaranteed to have both an incoming edge and an edge to another vertex in the graph are the “external” face and the k -dimensional simplices containing the colors $0 - k$, the latter are precisely the solutions to the **PPAD** problem. Moreover, since some k -dimensional simplex carved from some k -dimensional cubelet has the colors $0 - k$, the cubelet itself must contain all $k + 1$ colors as well. \square

Papadimitriou has also defined BROUWER [12], the following computational version of Brouwer fixed-points: “given an integer n and a Turing machine which computes a function $f : [0, 1]^k \rightarrow [0, 1]^k$, with the input given as a vector in coordinate multiples of $1/n$, the output given as a displacement vector of magnitude no greater than $1/n^2$, and the rest of the map implicitly defined by interpolating the values across any cubelet, find an input x such that $f(x) = x$.” He is then able to prove that BROUWER is **PPAD**-complete. We will only require the following:

Theorem 2. BROUWER \in **PPAD**.

Proof. We will provide a reduction from BROUWER to k D-SPERNER; the desired result will then follow from theorem 1.

Suppose we provide the following coloring: for each point p , we color the vertices in any arbitrary manner satisfying the following constraints:

1. A point \vec{p} may take color i only if the projection of $f(\vec{p}) - \vec{p}$ onto the vector defined by the i th principal vertex, \vec{e}^i , is negative.
2. A point \vec{p} may take color 0 only if the projection of $f(\vec{p}) - \vec{p}$ on the vector defined by each principal vertex, \vec{e}^i , is positive.

with the exception that, if $f(\vec{p}) - \vec{p} = 0$, we permit the vertex to be colored arbitrarily. Observe that unless we stumble upon such a \vec{p} , one of the constraints must hold.

Observe that on a face not containing the i th principal vertex, since f maps the cube into itself, the projection on the i th principal vertex cannot be negative, since this would map to a point outside the cube. Likewise, on the faces not containing e^0 , f cannot have a positive displacement in all components.

The solution to this instance of k D-SPERNER must contain a point such that $f(\vec{x}) = \vec{x}$: observe that we could increase n , subdividing the cubelet containing all $k + 1$ colors. Note that, up to some permutation of the colors, this subdivided cubelet is colored in some legal way due to the definition of our map, so this smaller cubelet (under some permutation and some easy transformation back to the cube $[0, 1]^k$) is another valid instance of k D-SPERNER. Consider the sequence defined by the centers of these cubelets, $\vec{x}^1, \vec{x}^2, \dots$. It is clear this sequence is converging to some limit, \vec{x}' .

Suppose now, for contradiction, that $f(\vec{x}') \neq \vec{x}'$. Then, either $f(\vec{x}') - \vec{x}'$ has a negative projection on some principal vector, or else it has a positive projection on all principal vectors. In the former case, since f is continuous, there must exist a δ such that within the ball of radius δ around \vec{x}' , f is within $|(f(\vec{x}') - \vec{x}') \cdot \vec{e}^i|/3$ of $f(\vec{x}')$, so clearly any cubelet contained within this δ -radius ball but around \vec{x}' must not have had any vertex colored 0, a contradiction. Similarly, in the latter case, we choose the δ such that within the ball of radius δ around \vec{x}' , f is within $\min_i \{|(f(\vec{x}') - \vec{x}') \cdot \vec{e}^i|/3\}$ of $f(\vec{x}')$, where it then follows that a sufficiently close cubelet must have all vertices colored 0, which is also a contradiction. Thus, we see that this cubelet contains a fixed-point, as desired. \square

3 Complexity Results

In Section 3.1, we will state one important result due to Condon [4], that the simple stochastic game problem is in $\mathbf{NP} \cap \mathbf{coNP}$. We will next state our classifications for the stable circuit problem in section 3.2, showing it to be in $\mathbf{PPAD} \cap \mathbf{PLS} \subseteq \mathbf{FNP}$. Finally, we will examine our ‘‘Leapfrog’’ variant, and will show that LEAPFROG is \mathbf{PSPACE} -hard, in section 3.3.

3.1 Prior Results for SSG-VALUE

The following important result is again due to Condon [4]:

Theorem 3. $\text{SSG-VALUE} \in \mathbf{NP} \cap \mathbf{coNP}$.

Proof. We first describe a nondeterministic Turing machine M that decides SSG-VALUE. On input graph G and vertex $start$, M constructs the stopping game G' whose minimum stable solution is within $4|V(G)|$ bits of that of G , as described in claim 6. Observe that by lemma 5, the value of the i th vertex in G is less than $1/2$ iff the value of the i th vertex in G' is less than $1/2$. Thus, M nondeterministically guesses the unique stable solution to G' , verifies that it is stable, and decides ‘‘yes’’ iff the value of vertex $start$ is greater than or equal to $1/2$.

Similarly, we can describe a nondeterministic Turing machine \bar{M} to decide the complement of SSG-VALUE. On input graph G and vertex $start$, \bar{M} also constructs the same sufficiently close stopping game G' , guesses the unique stable solution to G' , verifies that it is stable, and decides “yes” iff the value of vertex $start$ is less than $1/2$.

By claim 6, we can construct G' in time polynomial in G . It is further clear that we can guess the unique stable solution and verify that it is stable in time polynomial in G' , so M and \bar{M} both finish in polynomial time. Thus, SSG-VALUE \in $\mathbf{NP} \cap \mathbf{coNP}$, as desired. \square

3.2 Classifying Stable Circuit

We now present some new classifications of STABLE-CIRCUIT.

Theorem 4. STABLE – CIRCUIT \in **PLS**.

Proof. Consider the definition of a problem in **PLS**, and consider the Hoffman-Karp algorithm.

1. Clearly, we can choose an arbitrary pair of strategies σ and τ in polynomial time.
2. Define our potential function to be $\Phi(\sigma, \tau) = |\vec{m}^{\sigma, \tau}|_1$ for the sufficiently close discounted game; clearly, the L_1 norm can be computed in polynomial time, and $\vec{m}^{\sigma, \tau}$ can be found via Gaussian elimination in polynomial time. Thus, Φ can be computed in polynomial time.
3. We can clearly decide whether $\vec{m}^{\sigma, \tau}$ is stable in polynomial time.
4. By claim 8, since steps 2–4 of the Hoffman-Karp may clearly be completed in polynomial time, given a pair of strategies that are not locally optimal, we can clearly find a better pair in polynomial time.

Thus, if we simply define our neighborhood structure to be that defined by the body of the Hoffman-Karp algorithm, we see that we have met the conditions to conclude that STABLE – CIRCUIT \in **PLS**. \square

This result, while interesting, is much more interesting when one considers the following:

Theorem 5. STABLE – CIRCUIT \in **PPAD**.

Proof. (Sketch) We will reduce STABLE-CIRCUIT to BROUWER: Given a circuit C we first construct the circuit C' as described in claim 6, which has a unique stable solution that is within $4|C|$ bits of the minimum stable solution of C , and let F' be the circuit-wide update function for C' . Define $g(x) = \left(1 - \frac{1}{|C'|2^{8|C|}}\right)x + \frac{1}{|C'|2^{8|C|}}F'(x)$. Now, clearly $|x - F'(x)| \leq |C'|$, so $|x - g(x)| \leq 2^{-8|C|}$. Hence, if we put $n = 2^{4|C|}$, our displacement vectors $\mu(x)$ have length $|\mu(x)| \leq 1/n^2$, satisfying the conditions for BROUWER, where we can easily construct a Turing machine to compute these displacements. Now, since g and F' have the same fixed points, where the unique fixed point of F' is within $4|C|$ bits of the solution to STABLE-CIRCUIT(C), the first $4|C|$ bits of the fixed point to g is our solution. Hence, by Theorem 2, STABLE-CIRCUIT \in **PPAD**. \square

This makes STABLE-CIRCUIT the first problem outside **P** to be contained in **PLS** \cap **PPAD**, to our knowledge.




Figure 6: NOT Gadget

3.3 PSPACE-hardness of Leapfrog Circuits

In this section, we will build up to a reduction from TQBF to LEAPFROG in several steps, introducing one important gadget or circuit at each step. These results are from joint work with Manuel Blum and Ryan Williams. Before we begin, suppose that the i th region of the circuit, $th^{-1}(i) = \{x_0, \dots, x_k\}$. Suppose that the gate i outputs some value th_i and $\forall x_j : j = 0, \dots, k, x_j \in \{T_i, F_i\}$ where $0 \leq F_i < th_i < T_i \leq 1$ —that is to say, in the i th region, there are unique values for truth and falsehood, which are distinct from the threshold value. Observe that, within the i th region now, a MAX gate functions as an OR gate, and a MIN gate functions as an AND gate.

3.3.1 NOT Gadget

We now demonstrate that, by changing regions of the circuit, we can obtain the negation of an arbitrary gate’s output. We have taken to calling circuits in this model “Leapfrog” circuits because on this transition, the threshold and the value being negated “leapfrog” over one another. The reader should be aware that, in order to maintain the desired invariants outlined in the introduction to this section, it will be necessary to update every value from the previous region of the circuit when we change regions. We will do this in such a way as to preserve the encodings of each value, except for the one being negated.

Claim 9. *Suppose, in the gadget shown in figure 6, that the wire labeled th entering from above holds th_i , the threshold for x_0, \dots, x_k , where $\forall x_j : j = 0, \dots, k, x_j \in \{T_i, F_i\}$ such that $0 \leq F_i < th_i < T_i \leq 1$. Then, if the gates are updated in the order indicated by their labeling and we take the output of the gate labeled 2 to be the threshold th in a new region of the circuit, there will be values $0 \leq F' < th < T' \leq 1$ such that for each wire leaving the circuit labeled x_j ,*

1. $x_j \in \{F', T'\}$
2. The output $\sim x_0 = F'$ iff the input $x_0 = T_i$
3. $\forall j > 0$ the output $x_j = F'$ iff the input $x_j = F_i$.

Proof. We observe that the gates labeled 1, 2, 3, and 4 are updated first. Let th be the output of gate 2, T' be the output of gate 3, and F' be the output of gate 4. Notice that the output of gate 2 is

$$\frac{1}{2} \left(x_0 + \frac{1}{2}(x_0 + th_i) \right) = \frac{3}{4}x_0 + \frac{1}{4}th_i$$

While the output of gate 3 is $\max\{\frac{1}{2}(x_0 + th_i), x_0\}$ and the output of gate 4 is $\min\{\frac{1}{2}(x_0 + th_i), x_0\}$.

Case i. $th_i < x_0$. Clearly,

$$0 \leq F_i < F' < th < T' = T_i \leq 1$$

Now, $x_0 = T_i$, and we see $\sim x_0 = F'$. Moreover, for any x_j with $j > 0$, if the input $x_j = T_i$, then clearly the output $x_j = T_i = T'$. Likewise, if the input $x_j = F_i$, then in the MAX gates (e.g. gates 7, 8, etc. in figure 6) the output $x_j = \max\{F', F_i\} = F'$.

Case ii. $th_i > x_0$ holds, so

$$0 \leq F_i = F' < th < T' < T_i \leq 1$$

This time, $x_0 = F_i$ and $\sim x_0 = T'$. Now, given any x_j with $j > 0$, if the input $x_j = T_i$, then in the MIN gates (gates 5, 6, etc.) we find that the output is $\min\{T', T_i\} = T'$, and then the output labeled x_j will be T' as well. If, on the other hand, the input $x_j = F_i = F'$, then clearly the output labeled x_j will be F' as well.

□

We claim that the construction in figure 6 extends in an obvious way to an arbitrary number of wires. Clearly, if there are w wires excluding the wire carrying the threshold, then the gadget has size $2w + 2 = \Theta(w)$ gates. The reader is here advised to reconcile the seemingly contradictory set of circumstances—that we have constructed a non-monotone gadget using circuits of monotone gates! Of course, under gate-by-gate update, the sequence of outputs for each gate is still monotone; it is only by interpreting the values relative to one another that non-monotonicity is exhibited.

3.3.2 Boolean Formulas

Now, since we had gates that were equivalent to AND and OR, and now we have a construction for NOT gates, it is easy to build circuits that are equivalent to boolean formulas. It is particularly easy if the formula is given in, say, CNF: suppose we are given n input wires x_1, \dots, x_n and we wish to evaluate a CNF formula $\phi(x_1, \dots, x_n)$. After n stages of NOT gadgets we can, for each index i , obtain wires carrying x_i and $\sim x_i$. Since each NOT gadget has size $\Theta(n)$, where $|\phi| > n$ this portion has size $O(|\phi|^2)$. Now, for each j th clause C_j , clearly we can use $|C_j| - 1$ MAX gates to compute the OR of the literals in the clause (one new MAX gate for each literal after the first), and to compute the AND of the clauses takes one MIN gate for each clause after the first, so clearly by using only $\Theta(|\phi|)$ MAX and MIN gates, the final MIN gate will output T relative to the threshold iff $\phi(x_1, \dots, x_n)$ would evaluate to truth on the assignment encoded on the input wires. Hence, it is clear that we can evaluate any formula ϕ using circuits of size $O(|\phi|^2)$, which should be computable from the input formula ϕ in logspace.

3.3.3 Binary Counter

We next show how to construct a more intricate device using the NOT gadgets, one which has n gate outputs in the same region of the circuit that, on the k th round of gate-by-gate update

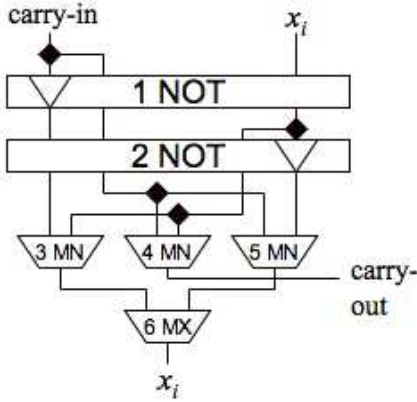


Figure 7: Circuitry for one bit of a binary counter

starting from $\bar{0}$, encode the integer $k \pmod{2^n}$ in binary. That is, when used in a circuit in an instance of LEAPFROG, it behaves as a binary counter in this way.

We will denote the least significant bit of the counter by x_0 . The construction controlling this bit is particularly easy, since we know that it flips between T and F on every round of gate-by-gate update. Thus, if we pass x_0 to a NOT gadget, we will obtain the value of x_0 for the next round, and may ultimately feed this value back to the first region of the circuit.

The remaining bits of the counter are not so simple. We take a moment now to state that, when we say that there is a “carry out” from the i th bit of the counter, we mean that if we were to increment the value displayed on the counter in the current round, then there would be a carry from the i th position. The reader should pause for a minute to note that the “carry out” value from x_0 is simply the value of the bit from the previous round. Keeping this definition in mind, we consider the circuit diagrammed in figure 7. For simplicity’s sake, we have omitted the threshold and any other wires that are “along for the ride” in the region transitions over the NOT gadgets from this diagram. The reader should also note that, when we label the NOT gadget with an integer, we mean that every gate of the NOT gadget should be updated in the order specified when we constructed the gadget, and these gates should be inserted into the indexing in the position indicated on the gate. When we refer to the “output” of a NOT gadget, we mean the value that was negated during the transition. In the diagram, this is the wire passing through the downward triangle.

Claim 10. *If the incoming wire x_i encodes the i th bit of the counter on the current round and the wire “carry in” encodes whether there would be a carry from the $i - 1$ st bit into the i th bit if the integer on the counter were incremented, then the outgoing wire x_i encodes the i th bit of the counter on the next round, and the wire “carry out” encodes whether there would be a carry from the i th bit into the $i + 1$ st bit of the counter.*

Proof. It will be helpful to recall that we assumed in our encoding that $T > F$.

Case i. carry-in encodes T, x_i encodes T.

Then, since the output of NOT 1 is F, the output of MIN 3 is F, and since the output of NOT 2 is F, the output of MIN 5 is F. Hence, the output of MAX 6 is F, as x_i should be on the next round. Likewise, it is clear that there would have been a carry out of the i th bit, and

since both carry-in and x_i encode T, MIN 4 will output T, which is the “carry out” value desired.

Case ii. carry-in encodes T, x_i encodes F.

Here, since carry-in encodes T and the output of NOT 2 will encode T, MIN 5 will output T, and hence MAX 6 will output T, which should be the value of x_i on the next round. Also, since x_i was F, notice MIN 4 will output F, where there would not have been a carry out.

Case iii. carry-in encodes F, x_i encodes T.

This time, x_i encodes T and the output of NOT 1 will encode T, so the output of MIN 3 will be T, and hence again MAX 6 must output T, as it should. Since carry-out encoded F, MIN 4 outputs F, and this is consistent with the fact that there is again no carry out.

Case iv. carry-in encodes F, x_i encodes F.

Again, since neither encodes T, clearly MAX 4 will output F, where this is correctly indicating that no carry will occur out of the i th bit. Since x_i encodes F, MIN 3 will output F, and since carry-in encodes F, MIN 5 also outputs F, and hence MAX 6 must output F as well, which is the value of x_i on the next round.

□

We now only need to ensure that our starting encoding has the desired properties, to have a working binary counter gadget. Unfortunately, $\bar{0}$ gives every gate output the same value, so we cannot immediately interpret the encoded truth values relative to thresholds. Suppose now, we include an AVG gate which has one input of constant 1 and its other input from the threshold in the final region of the circuit. If we give this gate the first index, then when it is updated, the gate outputs will be as in the vector $(1/2, 0, 0, \dots)$, so any gate assigned 1 as its threshold will, on the first iteration of the circuit, encode F.

Claim 11. *Suppose the gate i whose output is a threshold has its output passed through an AVG gate in which the second input is a 1. Then, if we take this gate’s output to be the threshold th' in a new region of the circuit, for each value $x_j \in \{T_i, F_i\}$ in the i th region of the circuit that is also passed through an AVG gate in which the second input is a 1, there will be values $T' > th' > F'$ such that the AVG gate outputs T' iff $x_j = T_i$ and otherwise the AVG gate outputs F' .*

Proof. Recall again that $T_i > th_i > F_i$. Put $th' = (th_i + 1)/2$, $T' = (T_i + 1)/2$, and $F' = (F_i + 1)/2$. Now, $T' > th' > F'$. Moreover, since $x_j \in \{T_i, F_i\}$, if $x_j = T_i$, then $AVG(x_j, 1) = T'$ and otherwise since $x_j = F_i$, then $AVG(x_j, 1) = F'$. □

Thus, we may use these AVG gates as a region transition. We choose to give the AVG gate for the threshold gate index 1 and the AVG gates for x_0, \dots, x_{n-1} indices $|C| - n + 1, \dots, |C|$. In this way, across rounds of gate-by-gate update, the gates for this region transition are in effect updated consecutively, but the threshold is updated at the beginning of the first step of a round where the rest are updated during the final steps of the previous round. Putting this all together, we can obtain a circuit that acts as a two-bit binary counter, as shown in figure 8. We have collapsed the circuitry controlling x_1 together with the circuitry controlling x_0 , since the latter (as discussed earlier) requires no more than a NOT gadget, and we also noted that the previous value of x_0 encoded the “carry out” value. Gates 1, 7, and 8 are as discussed in claim 11, with gate 1 specifically is being used to push the threshold above 0.




Figure 8: A 2-bit binary counter circuit

It is important to note that, although we have omitted the carry-out wires from this diagram, it is also possible to pass the “carry out” values back to the first region of the circuit. It should be clear that, given n , we could easily output a n -bit counter with these features using only logspace. Notice that, even if we pass each counter bit x_i and its “carry out” value through the entire circuit, this is still $\Theta(n)$ values, so each NOT gadget has size $\Theta(n)$. Since there are two NOT gadgets for each bit of the counter past the first, where clearly the rest of the circuitry for each bit of the counter is constant, each counter bit has size $\Theta(n)$ overall. Hence, the n -bit counter circuit has size $\Theta(n^2)$ overall. As an easy consequence, we find the following:

Theorem 6. $\text{SAT} \leq_{\ell} \text{LEAPFROG}$.

Proof. On input ϕ , we observed in Section 3.3.2 that we could output a circuit that computes $\phi(x_1, \dots, x_n)$ using $O(|\phi|^2)$ gates. It should be clear that we can output an n -bit binary counter using $\Theta(n^2)$ gates, with the $O(|\phi|^2)$ size formula circuits taking their inputs from the first region of the binary counter (e.g. the gate outputs encoding the binary integers), for a circuit of size $O(|\phi|^2)$ overall. Our LEAPFROG instance is this circuit, examining the output of the formula circuit relative to its threshold.

Since the counter places the binary encoding of $k \pmod{2^n}$ on the inputs to the formula circuits during the k th round of gate-by-gate update, if the formula has some satisfying assignment, then on the round such that the satisfying assignment interpreted as a binary integer is the round number, the output of the formula circuit will be strictly above its threshold, so the LEAPFROG instance will be a “yes.” On the other hand, if no such assignment exists, then on every round, the output of the formula circuit will be strictly below threshold, so the LEAPFROG instance will be a “no.” Thus, since the computation of these polynomial-sized circuits could clearly be carried out in logspace, we have logspace-reduced SAT to LEAPFROG. \square

So, we already see that under the assumption that $\text{NP} \neq \text{coNP}$, by Theorem 3, LEAPFROG




Figure 9: Tree of assignments to two bits

$\not\leq$ SSG-VALUE. We can and will demonstrate a stronger result next, though.

3.3.4 Universal Quantifier Gadgets

Before we discuss the remaining pieces needed to reduce TQBF to LEAPFROG, we will consider the following interpretation of the construction from Theorem 6 for motivation. Suppose we organize the assignments to x_0, \dots, x_{n-1} in a binary tree of depth $n + 1$ as follows: on a path from the root at the i th level, if we take the left branch we put $x_{i-1} = 0$, and if we take the right branch, we put $x_{i-1} = 1$. If the root of the tree is taken to be the n th level and the level above the leaves is taken to be the 1st level, then our n -bit binary counter is “walking” across the leaves of the tree from left to right; at the i th level of the tree, each time there is a carry out from the $i - 2$ nd bit, the assignment to x_{i-1} changes, and we switch directions, entering the right subtree of the current node if there is no carry out from x_{i-1} , and entering an entirely different subtree from a higher level if there is a carry out from x_{i-1} .

We may think of $\text{SAT}(\phi)$ as $\exists x_{n-1} \dots \exists x_0 \phi(x_0, \dots, x_{n-1})$, where in this tree, if ϕ evaluates to \top under the assignment at any leaf of the tree, we accept. To move to TQBF, we would like to switch one or more of these existential quantifiers to universal quantifiers, yielding a tree such as in figure 9. We will see that this can be done by granting each universally quantified variable one bit of memory (realized by a value being passed through the circuit in a loop) so that it can remember its value on the left branch, adding $\Theta(n)$ gates per quantifier.

Claim 12. *Suppose, in the gadget shown in figure 10, that the input $x_i = \text{F}$ (the current value of the i th bit of the counter), the input “carry-out: x_i ” holds the “carry out” value for x_i , that the input A will hold \top on some k th round of gate-by-gate update iff the quantified boolean formula $A(x_{n-1}, \dots, x_{i+1}, \text{F}) = \top$, and that the input v_i^0 holds F when the counter evaluates the assignment $\vec{0}$ and then holds the encoding of the same value output by gate 6 on the previous round. Then, the output of gate 7 is F and the output of gate 6 will hold \top after the counter evaluates all settings to x_0, \dots, x_{i-1} iff $A(x_{n-1}, \dots, x_{i+1}, \text{F}) = \top$.*

Proof. First observe that, when the counter enters the subtree rooted at the current assignment to x_{n-1}, \dots, x_i , since $x_i = \text{F}$, if we are not at the assignment $\vec{0}$, then there was a carry out from x_i . Hence, in this case the output of NOT 1 will be F , and under the assignment $\vec{0}$ we assumed the input $v_i^0 = \text{F}$, so in either case we find that the output of MIN 2 is initially F .




Figure 10: Universal quantifier gadget

Now, notice that since $x_i = \text{F}$, the output of NOT 3 is T . Hence, if $A(x_{n-1}, \dots, x_{i+1}, \text{F}) = \text{T}$ on some k th round, then clearly the output of gate 5 will be T , so the output of gate 6 is T . Likewise, if on some previous iteration while $x_i = \text{F}$, gate 6 output T , then since by assumption, v_i^0 is passed back through the circuit and there will not be another carry out of x_i until $x_i = \text{T}$, the output of gate 2 will be T , and so the output to gate 6 will continue to be T .

On the other hand, if $A(x_{n-1}, \dots, x_{i+1}, \text{F}) = \text{F}$, then the output to gate 5 will always be F , and since we saw earlier that the output of gate 2 is initially F , the output to gate 6 will be F . This will hold on the following iterations, as we pass this value for v_i^0 back around, since the output to gate 2 is again F .

To see that the output to gate 7 stays F , we need only observe that $x_i = \text{F}$, so the output of gate 4 will be F , and hence gate 7 will output F as well. \square

Claim 13. *Suppose, in the gadget shown in figure 10, that the input $x_i = \text{T}$, the input “carry-out: x_i ” holds the “carry out” value for x_i , and that the input A is set to T on some round of gate-by-gate update iff the quantified boolean formula $A(x_{n-1}, \dots, x_{i+1}, \text{T}) = \text{T}$. Also assume that the input v_i^0 holds $A(x_{n-1}, \dots, x_{i+1}, \text{F})$. Then, the output of gate 7 is set to T during some round of gate-by-gate-update iff $\forall x_i A(x_{n-1}, \dots, x_{i+1}, x_i) = \text{T}$.*

Proof. Notice that since $x_i = \text{T}$, the output of NOT 3 is F , so the output of gate 5 is F , and hence the output of gate 6 is v_i^0 . Since, by assumption, v_i^0 holds the evaluation of $A(x_{n-1}, \dots, x_{i+1}, \text{F})$, if this is F , then clearly $\forall x_i A(x_{n-1}, \dots, x_{i+1}, x_i)$ evaluates to F as well—where since the output of gate 6 is then F , the output of gate 7 will always be F , as needed.

Now suppose instead $v_i^0 = \text{T}$. If $A(x_{n-1}, \dots, x_{i+1}, \text{T}) = \text{T}$, then by assumption on some round during evaluation of x_0, \dots, x_{i-1} , the input A will be T , and since $x_i = \text{T}$ as well, the output to gate

4 will be \top . Moreover, since we assumed $v_i^0 = \top$, this means $A(x_{n-1}, \dots, x_{i+1}, \top)$ should evaluate to \top , and consequently $\forall x_i A(x_{n-1}, \dots, x_{i+1}, x_i)$ should also evaluate to \top . Of course now, both inputs to gate 7 are \top , so gate 7 outputs \top on this round, as needed.

On the other hand, if $A(x_{n-1}, \dots, x_{i+1}, \top) = \text{F}$ then by hypothesis gate 4 will always output F , and therefore gate 7 will output F as well. Naturally, since $\forall x_i A(x_{n-1}, \dots, x_{i+1}, x_i)$ should evaluate to F in this case, this is correct. \square

There are now, for each variable x_i , at most four wires being passed through the circuit—one for the current value of x_i , one for its negation, one for its “carry out” value, and now potentially one for its left-branch bit, if it is universally quantified. Still, the NOT gadgets have size $\Theta(n)$, since this is just a constant factor. This brings us to the main result:

Theorem 7. $\text{TQBF} \leq_\ell \text{LEAPFROG}$

Proof. Suppose our instance is of the form $Q_{n-1}x_{n-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_0)$. We then build the circuit as in Theorem 6 for $\phi(x_{n-1}, \dots, x_0)$, additionally passing the “carry out” values and an extra value v_i^0 for each universally quantified x_i through the counter. As observed above, this only increases the number of wires being passed between regions of the circuit by a constant factor, so the size of the circuit thus far is still $O(|\phi|^2)$.

Now, for each $Q_i = \forall$, we chain a quantifier gadget onto the output of the formula circuit for $\phi(x_{n-1}, \dots, x_0)$ —we pass the evaluation of this formula to the first quantifier gadget, and pass the output of gate 7 from the previous quantifier gadget as the wire marked A in the next. Each quantifier $Q_i = \forall$ is passed the current value from the counter of x_i , the “carry out” value for x_i , and the value v_i^0 in the other labeled inputs with its outputs labeled x_i and v_i^0 , of course, taken to be the current values for these wires subsequently. As usual, we pass $\Theta(n)$ values between regions of the circuit through the two NOT gadgets used in each quantifier gadget where there are clearly $O(n)$ quantifiers, so this chain of quantifier gadgets takes $O(n^2)$ gates, leaving our total for the circuit at $O(|\phi|^2)$. We take our LEAPFROG instance to be this circuit, asking about the output of the copy of gate 7 in the final quantifier gadget, relative to its threshold. As usual, this construction can clearly be done in logspace.

Assume inductively that the j th quantifier gadget (on x_i) will output \top iff

$$\forall x_i Q_{i-1}x_{i-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_0)$$

is true under the assignment specified by x_{n-1}, \dots, x_{i+1} on the counter while evaluating all settings of the lower $i + 1$ bits—for the base case, we simply observe that the unquantified formula is evaluated by our formula gadget.

Now, suppose the $j + 1$ st quantifier gadget is on x_k . Then, since during evaluation of the first k bits, the counter will evaluate

$$\forall x_i Q_{i-1}x_{i-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_0)$$

on all settings of x_{k-1}, \dots, x_{i+1} , we see that the input A for the $j + 1$ st quantifier gadget will be set to \top iff under the current assignment to x_{n-1}, \dots, x_k ,

$$\exists x_{k-1} \cdots \exists x_{i+1} \forall x_i Q_{i-1}x_{i-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_0) = \top$$

where, since x_k is the next universally quantified variable, this formula is clearly

$$Q_{k-1}x_{k-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_0).$$




Figure 11: Register circuit

Now, by claim 12, we see that v_k^0 will be set to T while $x_k = F$ and the counter evaluates all settings of the lower k bits iff

$$Q_{k-1}x_{k-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_{k+1}, F, x_{k-1}, \dots, x_0)$$

evaluates to T, and hence by claim 13, when x_k is set to T, the output of the $j + 1$ st quantifier gadget will be set to T at some point while the counter evaluates all settings of the lower k bits iff

$$\forall x_k Q_{k-1}x_{k-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_0)$$

evaluates to T under the current, fixed setting to x_{n-1}, \dots, x_{k+1} .

Hence, the quantifier gadgets evaluate the formulas as promised, up to the final universal quantifier. Now we only observe that, if the final universally quantified variable is x_k , then since the counter tries all settings to the existentially quantified x_{n-1}, \dots, x_{k+1} , the final quantifier gadget outputs T and is hence a “yes” instance of LEAPFROG iff the formula $Q_{n-1}x_{n-1} \cdots Q_0x_0\phi(x_{n-1}, \dots, x_0)$ is in TQBF. Since otherwise the final quantifier gadget always outputs F, making this a “no” instance of LEAPFROG, we see that we have successfully logspace-reduced TQBF to LEAPFROG. \square

Corollary 1. *If $\mathbf{NP} \neq \mathbf{PSPACE}$, $\mathbf{LEAPFROG} \not\leq \mathbf{SSG-VALUE}$.*

Proof. By Theorem 3 $\mathbf{SSG-VALUE} \in \mathbf{NP}$, and hence $\mathbf{LEAPFROG} \leq \mathbf{SSG-VALUE}$ would imply $\mathbf{NP} = \mathbf{PSPACE}$, as advertised. \square

3.4 Extending Leapfrog to the Limit

The reader may, at this point, feel that the result for LEAPFROG is rather unremarkable; after all, given that the problem asks for the existence of some index t such that the value output by the i th gate sneaks above the value output by the gate th . It might seem that the hardness of LEAPFROG is due only to the complicated pattern of values rising and falling with respect to their thresholds, which is clearly irrelevant to the limiting values. While it is true that we have shown that the pattern of rising and falling is irrelevant to the limiting values, it turns out that for not much more effort, we can obtain a result that suggests more about the relationship between Leapfrog circuits and their minimum stable solutions.

We define $\mathbf{LEAPFROG-LIMIT}(C, i, th)$ to be the decision problem “ $\exists N$ such that $\forall n > N$, if $\vec{v}^n = (F_{|C|} \circ \cdots \circ F_1)^n(\vec{0})$, $v_i^n > v_{th}^n$?” where F_i is the i th single-gate update function, as before. The reader should take note that if $\mathbf{LEAPFROG-LIMIT}(C, i, th)$ is a “yes” instance, then in the minimum stable solution, \vec{m} , $m_i \geq m_{th}$. Likewise, if it is a “no,” then $m_i \leq m_{th}$. Moreover, if $m_i > m_{th}$, then $\mathbf{LEAPFROG-LIMIT}(C, i, th)$ is a “yes” instance, and if $m_i < m_{th}$, then clearly $\mathbf{LEAPFROG-LIMIT}(C, i, th)$ is a “no” instance—that is, the only ambiguity occurs when $m_i = m_{th}$. Keeping these considerations in mind, observe the following:

Theorem 8. $\mathbf{TQBF} \leq_\ell \mathbf{LEAPFROG-LIMIT}$

Proof. Given ϕ , an instance of TQBF, construct the circuit as in Theorem 7, but make an extra copy of the circuitry for the threshold. In this copy, insert a MAX gate in the region containing the last quantifier gadget, and pass the gadget’s output into the second input of the MAX gate, constructing the register circuit shown in Figure 11. Suppose we otherwise treat this copy exactly as we treat the threshold—in particular, suppose we update its AVG gate for the transition into the first region of the circuit immediately after the AVG gate for the threshold in the first region (i.e. make it gate index 2). Denote the circuit constructed in this way by C and consider the problem LEAPFROG-LIMIT($C, 2, 1$).

Suppose $\phi \notin \text{TQBF}$. Then, as we argued before, on every round the output of the final quantifier gadget is below threshold. Thus, on every round the input to the register will be below threshold, so the register circuit’s internal wire will stay at threshold, and hence the register will continue to mimic the threshold circuitry precisely. Therefore, LEAPFROG-LIMIT($C, 2, 1$) is a “no.”

Now, suppose instead $\phi \in \text{TQBF}$. Then, on some t th round, the output of the final quantifier gadget is above threshold. Hence, the input to the register circuit will also be above threshold, so a value greater than threshold will be placed on the internal wire. Then, since the action of the AVG gates during translation between regions is strictly monotone, the internal register wire will stay strictly above threshold. Thus, clearly for all $n > t$, $v_2 > v_1$, and hence LEAPFROG-LIMIT($C, 2, 1$) is a “yes” instance, and since this extended construction should also be easily performed in logspace, we have logspace-reduced TQBF to LEAPFROG-LIMIT. \square

Corollary 2. *If $\text{NP} \neq \text{PSPACE}$, LEAPFROG – LIMIT $\not\leq$ SSG – VALUE.*

Proof. LEAPFROG-LIMIT is **PSPACE**-hard, so if LEAPFROG-LIMIT \leq SSG-VALUE, we would find **NP** = **PSPACE** by Theorem 3. \square

Thus, we see that under the assumption **NP** \neq **PSPACE**, it is *necessarily* the case that, in some Leapfrog circuits, $m_i = m_{th}$, e.g., that the encodings of T and F must approach each other, and this collapse of the encodings cannot be halted by means of some clever gadget.

4 Discussion and Open Problems

The primary motivation for this work was the wealth of work that has been put into finding a polynomial-time algorithm for SSG-VALUE, and the apparent lack of successes [5]. We felt that the problem was ripe for a hardness result, and set about trying to devise gadgets to this end. Leapfrog circuits seemed to be a natural model to build gadgets under, and particularly after discovering how to build a NOT gadget, we tried to combat the growth of the bits required by the encodings.

To get a sense of this, the reader is advised to quickly trace through the 2-bit binary counter in figure 8. Observe that each time the values pass through a NOT gadget, the AVG gates cause the number of bits of precision required to represent the values to increase by two. The AVG gates used in the transition from the final region of the circuit back to the first region of the circuit then add another required bit of precision *per iteration*. Hence, observe that the number of bits of precision needed to represent the values in a distinguishable way as increasing by five in each round. It should be clear that, since an exponential number of rounds is required in general for the counter to have stepped through all configurations and the number of bits by which the required precision grows is proportional to the number of NOT gadgets – a linear number of bits *per round* – the number of bits per value required by the 2^n th round is $\Theta(n2^n)$.

We had tried to fight this growth in hopes of obtaining a reduction to SSG-VALUE, which would provide us with our desired hardness result. Once we realized what these NOT gadgets permitted

us to do (e.g., Theorem 7), it became evident that we would be unlikely to keep the encodings distinguishable from some finite bit position. Theorem 8 and the related discussion in Section 3.4 frame this difficulty in a much more precise way. Consequently, although Leapfrog circuits are quite an interesting model in their own right, their relationship to the stable configurations of the circuits (i.e. and therefore to SSG-VALUE) is more distant than we had hoped.

Returning to STABLE-CIRCUIT, the recent oracle separations of **PLS** and **PPAD** by Morioka and Buresh-Oppenheim [3] raise the possibility that the classes are distinct. Clearly, a hardness result for STABLE-CIRCUIT relative to either of these classes would provide us with a containment one way or the other, but given that absolute separations between **P** and **NP** are (to put it mildly) quite rare, it seems possible that we may be left wondering where exactly the problem lies indefinitely.

4.1 Open Problems

Of course, we would still like to know where STABLE-CIRCUIT lies. In particular, if it is hard for either of **PLS** or **PPAD**, this would have the added value of demonstrating a containment between these two classes. Barring this, we would like to know if it is interreducible with some other problem in **PLS** or **PPAD** that has not been shown to be complete for either class (e.g. the Nash Equilibrium problem for two players where payoff matrices are given is a good candidate, an approach originally advocated by Schoenebeck [13]).

Returning to LEAPFROG now, for a moment, recall that we observed that the number of bits required to represent the encodings in a distinguishable way throughout the sequence of all 2^n configurations of an n -bit binary counter was exponential in n . Although we stopped after showing the **PSPACE**-hardness of LEAPFROG since this was the most powerful result we could show under our AND/OR/NOT interpretation (after all, AND/OR/NOT circuits can be evaluated in **PSPACE**), it is not clear that this is the most powerful result possible. That the circuits we used repeat configurations after 2^n steps is not sufficient to conclude that this is so in general. In fact, it is clear that the time and space needed to simulate the Leapfrog circuits “into the limit” in the straightforward way are not bounded. We wonder whether or not there is some more clever way to evaluate the circuits, giving us some sort of bound on the resources required.

It is clear that $\text{LEAPFROG} \in \mathbf{RE}$, since we can always perform the simulation until the finite index where the i th gate’s output rises above its threshold’s output, but this is quite a weak statement. It is currently unclear whether or not Leapfrog circuits can be made to function as a UTM, which would naturally imply that **RE** is as good as we can do. It would be interesting to see if this is the case, or if we can place LEAPFROG in some more modest class, such as **EXP**.

Acknowledgements

The authors would like to thank Anupam Gupta, Christos Papadimitriou, and Grant Schoenebeck for their contributions.

References

- [1] Avrim Blum. Communication.
- [2] Manuel Blum, Rachel Rue, and Ke Yang. On the complexity of MAX/MIN/AVRG circuits. Technical Report CMU-CS-02-110, Department of Computer Science, Carnegie Mellon University, 2002.

- [3] Joshua Buresh-Oppenheim and Tsuyoshi Morioka. Relativized NP search problems and propositional proof systems. *ECCC*, 84, 2004.
- [4] Anne Condon. The complexity of stochastic games. *Inf. Comput.*, 96(2):203–224, 1992.
- [5] Anne Condon. On algorithms for simple stochastic games. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 13:51–73, 1993.
- [6] Cyrus Derman. *Finite State Markovian Decision Processes*. Academic Press, Inc., 1970.
- [7] Alan J. Hoffman and Richard Karp. On nonterminating stochastic games. *Management Science*, 12:359–370, 1966.
- [8] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37(1):79–100, 1988.
- [9] L. G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Math. Dokl.*, 20:191–194, 1979. English translation of original article.
- [10] V. S. Anil Kumar and Rahul Tripathi. Algorithmic results in simple stochastic games. Technical Report URCS-TR-2004-855, Department of Computer Science, University of Rochester, 2004.
- [11] A. Paoluzzi, F. Bernardini, C. Cattani, and V. Ferrucci. Dimension-independent modeling with simplicial complexes. *ACM Trans. Graph.*, 12(1):56–102, 1993.
- [12] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.
- [13] Grant Schoenebeck. Communication.
- [14] Lloyd S. Shapley. Stochastic games. *Proc. National Academy of Sciences*, 39:1095–1100, 1953.