

Another Sub-exponential Algorithm for the Simple Stochastic Game

Decheng Dai · Rong Ge

Received: 23 February 2010 / Accepted: 4 May 2010 / Published online: 18 May 2010
© Springer Science+Business Media, LLC 2010

Abstract We study the problem of solving simple stochastic games, and give both an interesting new algorithm and a hardness result. We show a reduction from fine approximation of simple stochastic games to coarse approximation of a polynomial sized game, which can be viewed as an evidence showing the hardness to approximate the value of simple stochastic games. We also present a randomized algorithm that runs in $\tilde{O}(\sqrt{|V_R|!})$ time, where $|V_R|$ is the number of RANDOM vertices and \tilde{O} ignores polynomial terms. This algorithm is the fastest known algorithm when $|V_R| = \omega(\log n)$ and $|V_R| = o(\sqrt{\min\{|V_{\min}|, |V_{\max}|\}})$ and it works for general (non-stopping) simple stochastic games.

Keywords Simple stochastic game · Subexponential algorithm

1 Introduction

1.1 Simple Stochastic Games

Simple stochastic games are games played by two players on a graph, it is a restricted version of general stochastic games introduced by Shapley [5]. In a simple stochastic game, two players (MAX and MIN) move a pebble along directed edges in a graph.

D. Dai was supported by the National Natural Science Foundation of China Grant 60553001 and the National Basic Research Program of China Grant 2007CB807900, 2007CB807901.

D. Dai (✉)
Tsinghua University, Beijing, China
e-mail: ddc02@mails.tsinghua.edu.cn

R. Ge
Princeton University, Princeton, USA
e-mail: rongge@cs.princeton.edu

The vertices in the graph can have one of the three labels: MAX, MIN or RANDOM. If the pebble is on a vertex labeled MAX (or MIN), then MAX (or MIN) player decides through which outgoing edge the pebble should move; if the pebble is on a vertex labeled RANDOM, then the pebble moves along a randomly chosen edge. The graph also has a special vertex called the “1-sink”. If the pebble is moved to 1-sink, the game ends and the MAX player wins; if the game never ends (i.e. the MIN player forces it never to go to 1-sink), the MIN player wins.

Simple Stochastic Games (SSG for short) have many interesting applications. In complexity theory, SSGs are used in the analysis of space bounded computations with alternations and randomness [1] (the MAX and MIN vertices correspond to universal and existential quantifiers and RANDOM vertices correspond to coin tosses). It is also a special case of general stochastic games. In complexity view, the decision version of SSG is one of the few famous problems in $\mathbf{NP} \cap \mathbf{coNP}$, whether it is in \mathbf{P} is one of the major open questions in theoretical computer science. In practice, SSGs are used to model reactive systems. In such systems, RANDOM vertices are used to model stochastic environmental changes, MAX vertices are used to model adversary or arbitrary behaviors, MIN vertices are used to model choices of the system. The 1-sink vertex represents a failure. The goal of the system is thus minimizing the probability of failure (reaching 1-sink vertex).

Finding the optimal strategies for SSGs has been an interesting open problem for a long time. A lot of algorithms have been proposed. Condon [1] proved the decision version of SSG is in $\mathbf{NP} \cap \mathbf{coNP}$, and later in 1993, she showed several iterative algorithms for SSG in [2], but all of these algorithms require exponential time. She also suggested an approximation version of SSG problem, but there are no polynomial time algorithms known. Our gap amplification result gives an evidence on why the approximation problem is also difficult. Ludwig gave a sub-exponential ($\tilde{O}(2^{\sqrt{n}})$, \tilde{O} hides polynomial terms) time randomized algorithm for SSGs in [4], which uses local search techniques. Somla [6] purposed a new iterative algorithm in 2004 which might be better than previous algorithms, however there’s no evidence that shows the algorithm runs in polynomial time. Recently, Gimbert and Horn [3] presented a new non-iterative algorithm that runs in time $\tilde{O}(|V_R|!)$. This highlights one of the main reasons the problem has exponential complexity: the existence of random vertices.

1.2 Our Results

In this paper, we investigate the SSG problem in both hardness and algorithmic aspects. On the hardness side, by constructing a polynomial time reduction, we show that a coarse approximation of SSGs is as hard as a fine approximation. This is done by constructing a new game G' from a game G , such that G' has polynomial size and a coarse approximation to G' would give a fine approximation of game G .

At the algorithmic side, we present an algorithm based on the algorithm of Gimbert and Horn [3]. They considered a set of strategies called \mathbf{f} -strategies, and showed at least one of \mathbf{f} -strategies is optimal. However they were not able to distinguish “good” \mathbf{f} -strategies and “bad” \mathbf{f} -strategies. By finding a way to evaluate the “correctness” of \mathbf{f} -strategies, we are able to apply local search algorithms to find the optimal \mathbf{f} -strategy, and reduce the running time to $\tilde{O}(\sqrt{|V_R|!})$. Our algorithm is the

fastest known randomized algorithm for solving SSGs when $|V_R| = \omega(\log n)$ and $|V_R| = o(\sqrt{\min\{|V_{\max}|, |V_{\min}|\}})$.

In Sect. 2 we give definitions for Simple Stochastic Games and strategies. Then we describe the reduction from fine approximation to coarse approximation in Sect. 3. After that, we give a brief introduction to **f**-strategies and then present our algorithm.

2 Basic Definitions

There are many variations of SSGs, we define the game formally as follows

Definition 2.1 (Simple Stochastic Games) A simple stochastic game is specified by a directed graph $G = \langle V, E \rangle$ and a starting vertex $v_{\text{start}} \in V$. Each vertex $v \in V$ has 2 outgoing edges and a label (MAX, MIN or RANDOM). V_{\min} , V_{\max} , V_R are the sets of vertices with label MIN, MAX and RANDOM respectively. There's a special vertex v_1 (the 1-sink) in the graph.

Initially the pebble is at v_{start} . If the pebble is at a MAX/MIN vertex, then the corresponding player moves the pebble along one of the outgoing edges. If the pebble is at a RANDOM vertex, then the pebble moves along a random outgoing edge (both edges are chosen with probability 1/2). If the pebble reaches v_1 then MAX player wins, otherwise MIN player wins.

Solving SSGs means calculating the winning probabilities for the players if they all follow optimal strategy. Informally, the strategy of a player decides which edge should the pebble follow in the game. Although a strategy can decide the edge by considering history and using random coins, it's well known that *positional* optimal strategies exist for simple stochastic games [1, 5]. A positional strategy makes the decision only by the current position of the pebble. Formally, a positional strategy for MAX player α is a function from V_{\max} to V , for any vertex $v \in V$, $(v, \alpha(v))$ is an edge and it is the outgoing edge that the MAX player would choose if the pebble is currently at vertex v . Similarly, a positional strategy for MIN player β is a function from V_{\min} to V . Hereafter when we mention strategy we mean positional strategy. We define the value of a vertex to be the winning probability of the MAX player if initially the pebble is at this vertex, and denote this by $\text{val}(v)$, the value of the game is $\text{val}(v_{\text{start}})$. When it's not clear which game we are talking about, we use $\text{val}[G](v)$ to specify the value of v in game G .

In simple stochastic games, MIN player wins the game by forcing the pebble to move infinitely many steps without reaching the 1-sink. Sometimes it's easier to consider the situation that the game has two sinks: a 0-sink and a 1-sink. The game guarantees no matter what strategies the players use, with probability 1 the game will reach one of the sinks in finitely many steps. The goal of MAX player is to reach the 1-sink and the goal of MIN player is to reach the 0-sink. This variation of SSG is called stopping simple stochastic games (stopping-SSG). Condon showed in [1] that any SSG can be converted to a stopping-SSG in polynomial time while the change in the value of the game is exponentially small.

3 Coarse Approximation is As Hard As Fine Approximation

Since no polynomial time algorithms has been discovered for exactly solving SSGs, Condon [1] purposed the following “approximation” version of the problem. Consider the following sets,

$$L_{\text{yes}} = \left\{ G : \text{the value of } G \text{ is at least } \frac{1}{2} + \epsilon \right\},$$

$$L_{\text{no}} = \left\{ G : \text{the value of } G \text{ is at most } \frac{1}{2} - \epsilon \right\}.$$

An ϵ -gap SSG decision problem is to determine whether G is in L_{yes} or L_{no} given it is in one of them. Intuitively it might seem for some large enough ϵ this problem is easy to solve. However, we give a gap amplification reduction showing that when enlarging ϵ from $(1/\text{poly}(n))$ to $(1/2 - e^{-n^\rho})$ for any $\rho < 1$, the problem does not become easier. This reduction is analogue to the hardness amplification results for clique and chromatic number problem.

Theorem 3.1 *For any fixed constant $0 < \rho < 1$ and $c > 0$, if the $(1/2 - e^{-n^\rho})$ -gap SSG decision problem is in \mathbf{P} , then the (n^{-c}) -gap SSG decision problem is in \mathbf{P} .*

Proof First we prove the theorem for stopping SSG.

Now let’s assume $G = \langle V, E \rangle$ is a stopping-SSG with n vertices. There are 3 special vertices in a stopping SSG: v_{start} , the starting vertex; v_1 , the 1-sink vertex; v_0 , the 0-sink vertex. We construct another game $G' = \langle V', E' \rangle$ of size N (which is polynomial in n) such that,

- G' has value larger than $(1 - e^{-N^\rho})$ if G has value larger than $(1/2 + n^{-c})$.
- G' has value less than e^{-N^ρ} if G has value less than $(1/2 - n^{-c})$.

Informally, G' “repeats” G for polynomially many times, if the MAX(MIN) player wins all these games, he wins G' ; otherwise the pebble goes back to v_{start} and restart G' . Since G is a stopping game, G' is also a stopping game.

Let $\{G_0, G_{i,j} | i \in \{0, 1\}, j \in \{1, \dots, K\}\}$ be $2K + 1$ copies of G . We replace the out-going edges for v_0 and v_1 in each of these games to connect them together in the following way (their two outgoing edges will point to the same location, so it doesn’t matter what label they have)

- Connect v_0 in G_0 to v_{start} in $G_{0,1}$, connect v_1 in G_0 to v_{start} in $G_{1,1}$.
- Connect v_1 in $G_{0,j}$ ($1 \leq j \leq K$) to v_{start} in G_0 , connect v_0 in $G_{0,j}$ ($j < K$) to v_{start} in $G_{0,j+1}$.
- Connect v_0 in $G_{1,j}$ ($1 \leq j \leq K$) to v_{start} in G_0 , connect v_1 in $G_{1,j}$ ($j < K$) to v_{start} in $G_{1,j+1}$.
- The starting vertex in G' is v_{start} in G_0 , and the 0-sink vertex is v_0 in $G_{0,K}$ and the 1-sink vertex is v_1 in $G_{1,K}$.

In this constructed game G' (as showed in Fig. 1), the MIN(MAX) player must win G_0 and all $G_{0,j}(G_{1,j})$ to win G' . Let p to be the value in G . By induction, it is easy to prove the probability to reach v_1 in $G_{1,K}$ is $(p^K)/(p^K + (1 - p)^K)$.




Fig. 1 An example on constructing G' from G , for $K = 3$. Every ellipse is a copy of G . Three solid vertices are v_{start}, v_0, v_1 . It is easy to check that the probability to reach the v_0 in G' is exactly $p^3 / (p^3 + (1 - p)^3)$, in which $p = \text{val}[G](v_0)$

Let $K = n^d$ where $d = (c + 1) / (1 - \rho)$, then $N = (2K + 1)n = O(n^{d+1})$ when $p \leq 1/2 - n^{-c}$ we have,

$$\begin{aligned} \frac{p^K}{p^K + (1 - p)^K} &\leq \frac{(\frac{1}{2} - n^{-c})^K}{(\frac{1}{2} - n^{-c})^K + (\frac{1}{2} + n^{-c})^K} \\ &\leq (1 - 2n^{-c})^K \leq e^{-n^{d-c}} \\ &\leq e^{-N^\rho}, \end{aligned}$$

so the value of G' is less than e^{-N^ρ} in this case. Similarly, when $p \geq 1/2 + n^{-c}$, we have the value of G' is larger than $1 - e^{-N^\rho}$. That is,

$$\begin{aligned} \text{val}(G) \leq \frac{1}{2} - n^{-c} &\Rightarrow \text{val}(G') \leq e^{-N^\rho}, \\ \text{val}(G) \geq \frac{1}{2} + n^{-c} &\Rightarrow \text{val}(G') \geq 1 - e^{-N^\rho}. \end{aligned}$$

By applying the algorithm for $(1/2 - e^{-n^\rho})$ -gap SSG decision problem on G' , the algorithm would be able to distinguish between $\text{val}(G) > 1/2 + n^{-c}$ and $\text{val}(G) < 1/2 - n^{-c}$.

For general (possibly non-stopping) SSG, we use Condon’s reduction in [1] that transforms a SSG G to a stopping SSG G' whose value is arbitrarily close to the value of G . The constructed stopping game G' adopts all the vertices of G and inserts cnm new vertices ($m = |E|$ in G). For any vertex v , $|\text{val}[G](v) - \text{val}[G'](v)| \leq 2^{(2-c)n}$. By combining these two constructions, we can reduce solving the (n^{-c}) -gap decision problem to $(1/2 - e^{-n^\rho})$ -gap SSG decision problem. \square

4 Fast Algorithm for SSGs with Few Random Vertices

An interesting case for solving simple stochastic games is when there are a few random vertices. Gimbert and Horn [3] found an algorithm that runs in $\tilde{O}(|V_R|!)$ time. Their algorithm is based on enumerating a special kind of strategies called **f**-strategies. To avoid simple and time consuming enumerations, our algorithm relies on the following Lemma:

Lemma 4.1 (Main Lemma) *There’s a partial order in **f**-strategies such that the following holds:*

1. Any maximal **f** corresponds to a pair of optimal strategies.
2. Two **f**-strategies can be compared in polynomial time.
3. If **f** is not maximal, then in polynomial time we can find **g** which is better than **f**.

The **f**-strategies are first introduced in [3], and they proved a theorem (Lemma 4.3 in this paper) on testing whether the **f**-strategy is optimal or not. The brut-force idea is to enumerate all possible $O(n!)$ **f**-strategies and use Lemma 4.3 to find the optimal one. Our major contrubution is this main lemma, which reduces the problem to a local maximal searching problem and thus enabled us to design faster algorithms.

4.1 **f**-strategies

In this section we’ll first briefly describe ideas in [3] on what are **f**-strategies and how to test their optimality; this is first introduced in [3] and we mention it again for completeness. Then we show how the partial order in Lemma 4.1 is defined and prove the lemma. Finally we use existing randomized algorithms for local search problems to improve the expected running time to $\tilde{O}(\sqrt{|V_R|!})$.

Let $\mathbf{f} = \langle r_1, \dots, r_m \rangle$ (for simplicity let $r_0 = v_1$) be a permutation of the random vertices, where $m = |V_R|$ is the number of the random vertices. A **f**-strategy is a pair of positional strategies associated to **f**.

Let R_i be the first i random vertices in the permutation **f**. The *consuming set* C_i is a set of vertices from which player MAX has a strategy $\sigma_{\mathbf{f}}$ for moving the pebble to R_i and at the same time avoid touching any other random vertices, no matter what strategy player MIN chooses. Similarly, there’s also a strategy $\tau_{\mathbf{f}}$ for player MIN, such that no matter what player MAX does, vertices outside C_i can never reach a vertex in R_i without touching other random vertices. Obviously $C_0 \subseteq C_1 \subseteq \dots \subseteq C_m$. The consuming sets and the strategies $\alpha \beta$ can be constructed using the following algorithm:

1. $C_{-1} = \Phi$.
2. FOR $i = 0$ to $|V|$.
3. $C_i = C_{i-1} \cup \{r_i\}$.
4. REPEAT until no more vertices can be added.
5. ADD a MAX vertex v to C_i if v has an edge (v, w) where $w \in C_i$, let $\alpha(v) = w$.

6. ADD a MIN vertex v to C_i if all remaining outgoing edges of v points to a vertex in C_i , let $\beta(v)$ be the endpoint of any remaining outgoing edge of v .
7. REMOVE all outgoing edges from a MIN vertex in $V \setminus C_i$ to C_i .
8. For remaining vertices choose an arbitrary edge as the strategy.

For any strategy of MIN player, the vertices in $C_i \setminus C_{i-1}$ can always reach r_i by strategy α , this can be proved using simple induction on the order of entering set C_i . For any MAX strategy, the vertices outside C_i can never reach a vertex in $\{r_0, \dots, r_i\}$, because there's no edge between $V \setminus C_i$ and C_i in the remaining graph. This pair of strategies $(\sigma_{\mathbf{f}}, \tau_{\mathbf{f}})$ is called the \mathbf{f} -strategy regarding to the permutation \mathbf{f} . For any permutation \mathbf{f} , let $\text{val}_{\mathbf{f}}(r_i)$ be the probability for player MAX to win if the game starts at vertex r_i , when players follow the \mathbf{f} -strategies $(\sigma_{\mathbf{f}}, \tau_{\mathbf{f}})$. The following lemmas are first proved in [3].

Lemma 4.2 (f-strategy) *Given any permutation \mathbf{f} , the corresponding $\sigma_{\mathbf{f}}$ and $\tau_{\mathbf{f}}$ always exist and can be found in polynomial time.*

Informally, a permutation $\mathbf{f} = \langle r_1, \dots, r_m \rangle$ represents the preference of player MAX and MIN over the random vertices: MAX prefer to begin the play from a random vertex of index as high as possible, i.e., the higher index random vertices have larger probability to reach the 1-sink. There exists a permutation \mathbf{f} , such that the \mathbf{f} -strategy is an optimal strategy. Moreover, there are simple rules to check whether the \mathbf{f} -strategy is optimal or not.

Lemma 4.3 *If \mathbf{f} satisfies the Consistency and Progressive conditions, then the \mathbf{f} -strategy is an optimal strategy for the game.*

Consistency: $\text{val}_{\mathbf{f}}(r_1) \geq \dots \geq \text{val}_{\mathbf{f}}(r_m)$.

Progressive: *For any random vertex r_i ($i > 0$) with $\text{val}_{\mathbf{f}}(r_i) > 0$, at least one of its outgoing edges points to a vertex in C_{i-1} .*

There always exists a permutation \mathbf{f} that satisfy both conditions.

For constructing $\{C_i\}$ in polynomial time and more discussions about the *Consistency* and *Progressive* conditions, see [3]. The *Consistency* constraint is easy to understand, however, the *Progressive* constraint is also important. If a permutation \mathbf{f} satisfies *Consistency* but does not satisfy the *Progressive* condition, its corresponding \mathbf{f} -strategy may not be optimal. Consider the case in Fig. 2. The permutation $\langle r_2, r_1 \rangle$ is consistency and progressive. And in this permutation, the MIN's strategy is to go to r_2 , and the values are 0 for all the vertices. Consider the permutation $\langle r_1, r_2 \rangle$, in this case MIN's strategy is to go to r_1 , and so the values are 1 for every vertices, in this case the ordering is *Consistent* but not *Progressive* and the strategies are not optimal.

However, the *Progressive* constraint is not a necessary constraint for optimality. Consider two random vertices r_1 and r_2 , all outgoing edges of r_2 points to r_1 . Obviously the value of the two vertices are the same. If in the ordering, $v^2 \geq v^1$ and their value are greater than 0, then the ordering cannot be consistent because r_2 can only

Fig. 2 Progressive is an important condition. $\langle r_1, r_2 \rangle$ is a non-progressive but consistent permutation. According to this permutation, player MIN prefers r_1 over r_2 , and his strategy will be totally wrong

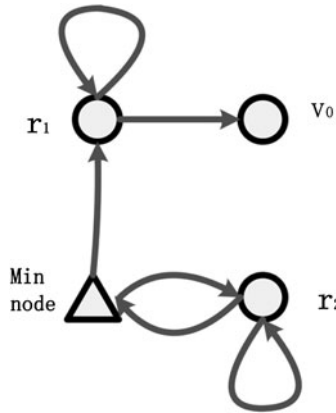



Fig. 3 Values cannot tell whether the f -strategy is good or bad



reach r_1 , which is lower in the ordering. However the f -strategy corresponding to the ordering can be optimal because r_1 and r_2 are essentially the same vertex (we can merge them without changing the value of other vertices).


4.2 The Partial Order for f -strategies

Gimbert and Horn [3] solve SSGs by enumerating all possible permutation f (and f -strategies). A natural way to improve the algorithm by Gimbert and Horn would be smartly updating f when it is not *Consistent* or *Progressive*. However, it is hard to tell two permutations which is better by simply looking at the values. This is not trivial because by simply looking at the values we have no way to tell whether an ordering is “good” or “bad”. A wrong ordering would “fool” both the MAX player and the MIN player, and the result is the value based on f -strategy can be greater than, less than or even equal to the real value. Consider the example in Fig. 3. The consistent and progressive permutation is $f = \langle r_1, r_2, r_3 \rangle$ and their corresponding values are $\text{val}_f(r_1) = 1/4, \text{val}_f(r_2) = 0.5, \text{val}_f(r_3) = 3/4$. The permutation $g = \langle r_3, r_2, r_1 \rangle$ is neither consistent nor progressive, and the corresponding values are $\text{val}_g(r_1) = 0, \text{val}_g(r_2) = 0.5, \text{val}_g(r_3) = 1$, in which r_1 's value is less than its true value but r_3 's value is larger than the true value.

To estimate whether a particular ordering is good or not, we construct a new SSG with respect to the ordering.

Definition 4.1 (Value measure $H(f)$) Let G be a SSG and f be an ordering of random vertices, G_f is a new SSG. G_f has all the vertices and edges in G and m new vertices

Fig. 4 (a) The game G , Δ are MIN vertices, \square are MAX vertices, \circ are RANDOM vertices. (b) The graph $G_{\mathbf{f}}$, in which $\mathbf{f} = \langle r_1, r_2, r_3, v_1 \rangle$. 4 MAX vertices are added. The dashed lines are the original edges and the solid lines are the added edges



u_1, u_2, \dots, u_m , all of them are MAX vertices. The two outgoing edges of u_i go to r_i and r_{i+1} (the two outgoing edges of u_m both go to r_m). All edges of the form (v, r_i) in G are replaced by (v, u_i) in $G_{\mathbf{f}}$. Let $H(\mathbf{f}) \triangleq \sum_{i=1}^m \text{val}[G_{\mathbf{f}}](u_i)$.

Let $H_{\text{OPT}} = \sum_{i=1}^m \text{val}[G](r_i)$. An example on how to compute $H(\mathbf{f})$ is shown in Fig. 4. In G , the values are $\text{val}(r_1) = 0, \text{val}(r_2) = 0.5, \text{val}(r_3) = \text{val}(v_1) = 1$. In $G_{\mathbf{f}}$, the values are $\text{val}(r_1) = \text{val}(r_2) = \text{val}(r_3) = \text{val}(v_1) = 1$. So $H(\mathbf{f}) = 4 > H_{\text{OPT}}$.

Lemma 4.4 For any permutation \mathbf{f} , $H(\mathbf{f}) \geq H_{\text{OPT}}$. When \mathbf{f} is both Consistent and Progressive, $H(\mathbf{f}) = H_{\text{OPT}}$.

Proof Consider a permutation \mathbf{f} and its corresponding $G_{\mathbf{f}}$, assume α, β is a pair of optimal strategies for the original game G . Now we construct a strategy α' for player MAX in Game $G_{\mathbf{f}}$: $\alpha'(v) = \alpha(v)$ for all $v \in G$; $\alpha'(u_i) = r_i$ for all $u_i, 1 \leq i \leq m$. When player MAX takes this strategy, it is easy to check β is the also best response for player MIN in G' . So for every $v \in G$, $\text{val}[G_{\alpha, \beta}](v) = \text{val}[G_{\mathbf{f}, \alpha', \beta}](v)$ and $\sum_{i=1}^m \text{val}(u_i) = \sum_{i=1}^m \text{val}[G_{\mathbf{f}, \alpha', \beta}](v) = H_{\text{OPT}}$. However, MAX may have better strategies in $G_{\mathbf{f}}$, so $H(\mathbf{f}) \geq H_{\text{OPT}}$.

When \mathbf{f} is Consistent and Progressive in G , we first prove that \mathbf{f} is also Consistent and Progressive in $G_{\mathbf{f}}$. Let C_i be the consuming sets in G regarding to \mathbf{f} . By analyzing the structure of graph $G_{\mathbf{f}}$, we have $C'_i = C_i \cup \{u_1, u_2, \dots, u_i\}$ are consuming sets in $G_{\mathbf{f}}$. Consider the strategies (α', β) as defined in the former case. Using the definition of \mathbf{f} -strategy, it is easy to verify that (α', β) are \mathbf{f} -strategy for $G_{\mathbf{f}}$. So $\text{val}[G_{\mathbf{f}}](r_i) = \text{val}[G](r_i)$, which means \mathbf{f} is still Consistent and Progressive for $G_{\mathbf{f}}$. Therefore $H(\mathbf{f}) = H_{\text{OPT}}$. \square

This lemma shows that the minimum value of $H(\mathbf{f})$ is the corresponding value in the optimal strategy. To compute the optimal strategy and values for $G_{\mathbf{f}}$, we use the following Lemma.

Lemma 4.5 For any permutation \mathbf{f} and the \mathbf{f} -strategy (σ, τ) in G , there is an optimal strategy (σ', τ') for $G_{\mathbf{f}}$ such that for all $v \in G$, $\tau(v) = \tau'(v)$ and $\sigma(v) = \sigma'(v)$.

Proof Let the permutation $\mathbf{f} = \langle r_0, \dots, r_m \rangle$ and the corresponding \mathbf{f} -strategy $(\sigma_{\mathbf{f}}, \tau_{\mathbf{f}})$. Now we construct strategy for $G_{\mathbf{f}}$ satisfies the conditions.

Denote the consuming sets for \mathbf{f} as $\{C_i\}$. Let $\mathbf{g} = (r_0, r_{t_1}, r_{t_2}, \dots, r_{t_m})$ be a permutation for $G_{\mathbf{f}}$ which is consistent and progressive (this ordering always exists by Lemma 4.3). Since G and $G_{\mathbf{f}}$ have the same random vertices, \mathbf{f} and \mathbf{g} are permutations over the same set. Denote the consuming sets for \mathbf{g} as $\{C'_i\}$, we have $\bigcup_i C_i = \bigcup_i C'_i$. By the construction of $G_{\mathbf{f}}$, we have $C_i = \bigcup_{j=1}^i (C_j \cup \{u_j\})$ and $C'_i = \bigcup_{j=1}^{\max_{t_1, t_2, \dots, t_i} t_j} (C_j \cup \{u_j\})$. This is because in the \mathbf{f} -strategy in G , player MAX's strategy ensures C_i can reach R_i while MIN strategy ensures that no other vertices outside C_i can reach R_i .

Now the optimal strategies (σ', τ') are defined as follows.

$$\sigma'(v) = \begin{cases} \sigma(v), & \text{if } v \in G, \\ r_{t_i}, & \text{if } v = u_{t_i} \text{ and } t_i = \max_{j \leq i} t_j, \\ u_{t_i+1}, & \text{if } v = u_{t_i} \text{ and } t_i < \max_{j \leq i} t_j. \end{cases} \tag{1}$$

Since the MIN vertices in $G_{\mathbf{f}}$ are the same with G , we simply let $\tau'(v) = \tau(v)$.

Then for any i , the strategy σ' makes sure that no matter what strategy the MIN player uses, C'_i always reach a vertex in $\{r_0, r_1, \dots, r_{t_i}\}$. Similarly, the strategy τ' makes sure that no matter what strategy that the MAX player uses, vertices outside C'_i can never reach a vertex in $\{r_0, r_1, \dots, r_{t_i}\}$. So (σ', τ') is a valid \mathbf{f} -strategy for the permutation \mathbf{g} . Since \mathbf{g} is consistency and progressive, (σ', τ') is therefore optimal by Lemma 4.3. □

By this lemma we can find the optimal strategy for MIN player in $G_{\mathbf{f}}$ in polynomial time, because we know that the strategy for player MIN in \mathbf{f} -strategy for G is also an optimal strategy for MIN player in $G_{\mathbf{f}}$. By using linear programming we can find the optimal strategy for player MAX in polynomial time.

Definition 4.2 (Progressiveness measure $P(\mathbf{f})$) For an permutation $\mathbf{f} = \langle r_1, \dots, r_m \rangle$, $P(\mathbf{f})$ is the smallest i ($i > 0$) such that r_i does not have an outgoing edge to C_{i-1} . If there's no such i or $\text{val}(r_i) = 0$ then $P(\mathbf{f}) = m + 1$.

Denote the set of all permutations over the random vertices as Π . It takes $\tilde{O}(m!)$ time to search this space and output the consistent and progressive one. However, a partial order over Π may help us to find this ordering. We say $\mathbf{f} > \mathbf{g}$ if (1) $H(\mathbf{f}) < H(\mathbf{g})$ or (2) $H(\mathbf{f}) = H(\mathbf{g})$ and $P(\mathbf{f}) > P(\mathbf{g})$.

Any maximal element in $(\Pi, >)$ corresponds to an permutation that is both consistent and progressive. Therefore we have proved the first 2 parts of Lemma 4.1. To prove part 3 of Lemma 4.1, we use the following lemma as a tool to upperbound the H value.

Lemma 4.6 *If function $f : V \rightarrow [0, 1]$ satisfy the following conditions, then $\text{val}(v) \leq f(v)$ for every vertex v .*

1. For vertex v_1 , $f(v_1) = 1$;
2. For vertex $v \in V_R$, assume the two outgoing edges are $(v, w_1), (v, w_2)$, $f(v) \geq (f(w_1) + f(w_2))/2$;

3. For vertex $v \in V_{\text{MAX}}$, assume the two outgoing edges are $(v, w_1), (v, w_2)$, $f(v) \geq \max(f(w_1), f(w_2))$;
4. For vertex $v \in V_{\text{MIN}}$, assume the two outgoing edges are $(v, w_1), (v, w_2)$, $f(v) \geq \min(f(w_1), f(w_2))$.

Proof We say functions $f \leq g$ if for any RANDOM vertex v $f(v) \leq g(v)$.

Consider a pair of optimal \mathbf{f} -strategy (α, β) according to ordering r of G and the corresponding consuming sets. By the construction of consuming sets we have for any vertex v in $C_i \setminus C_{i-1}$, $f(v) \geq f(r_i)$ (this is just a simple reduction on the order of entering consuming sets). For a random vertex r_i , let its two out-going edges point to vertex w_1, w_2 , where $w_1 \in C_j \setminus C_{j-1}$, $w_2 \in C_k \setminus C_{k-1}$, then $\text{val}(r_i) = 1/2(\text{val}(r_j) + \text{val}(r_k))$. The value of non-zero random vertices are the unique solution of a system of linear equations $\text{val}(r_i) = 1/2(\text{val}(r_j) + \text{val}(r_k))$, where if $j = 0$ $\text{val}(r_j)$ is replaced by constant 1 and if $\text{val}(r_j) = 0$ then it is replaced by constant 0 (the same replacement also applies for $\text{val}(r_k)$). By the definition of f , and for any vertex v in $C_i \setminus C_{i-1}$, $f(v) \geq f(r_i)$, we can see that f values satisfy the same sets of constraints with greater or equal signs: $f(r_i) \geq 1/2(f(r_j) + f(r_k))$ (with same replacements). Define operator $\Gamma f(r_i) = 1/2(f(r_j) + f(r_k))$, then the only fixed point of Γ is val . And Γ will only decrease f , so $f \geq \text{val}$. □

The next lemma shows that we can easily find a better ordering if the current ordering is not maximal.

Lemma 4.7 *If an permutation $\mathbf{f} = \langle r_1, \dots, r_m \rangle$ is not maximal, then there exists an element r_i in \mathbf{f} , by deleting r_i and reinsert it in an appropriate place we get a new ordering \mathbf{g} such that $\mathbf{g} > \mathbf{f}$.*

Proof If the ordering \mathbf{f} is not consistent in $G_{\mathbf{f}}$, then there exists some t such that $\text{val}[G_{\mathbf{f}}](r_t) < \text{val}[G_{\mathbf{f}}](r_{t+1})$. Find a place $q > t$ so that $\text{val}[G_{\mathbf{f}}](u_q) < \text{val}[G_{\mathbf{f}}](r_{t+1})$ (if there's no such place then let $q = m + 1$). Delete r_t and reinsert it right before q (if $q = m + 1$ then insert it at the tail). Define $f(v) = \text{val}[G_{\mathbf{f}}](v)$, then for graph $G_{\mathbf{g}}$ f is a valid value function that satisfy the requirements of Lemma 4.6. Therefore for any vertex v $\text{val}[G_{\mathbf{f}}](v) \geq \text{val}[G_{\mathbf{g}}](v)$. Particularly for the current position of r_t , the corresponding u vertex is u_q in $G_{\mathbf{g}}$, $f(u_q) > \max(f(u_{q+1}), f(r_t))$, so even after reducing $f(u_q)$ to $\max(f(u_{q+1}), f(r_t))$, f is still valid. That is, $H(\mathbf{g}) < H(\mathbf{f})$, $\mathbf{g} > \mathbf{f}$.

If the ordering \mathbf{f} is consistent but not progressive, then assume $P(\mathbf{f}) = t$. Define a graph among the random vertices and r_0 as follows: if an original outgoing edge of r_i goes to a vertex v that is in $C_j \setminus C_{j-1}$, then there's an edge from r_i to r_j . Use breadth first search to find $t' > t$, such that the following holds:

1. There's an edge from $r_{t'}$ to $\{r_0, r_1, \dots, r_{t-1}\}$.
2. There's a path from r_t to $r_{t'}$.

Note that such t' must exist because otherwise following the \mathbf{f} -strategy, starting from r_t , the pebble will never be able to reach r_0 , and therefore the value of r_t is 0, which contradict with the fact that $P(r) \neq m + 1$. Also, $\text{val}(r_{t'}) = \text{val}(r_t)$, because if the path from r_t to $r_{t'}$ is (w_0, w_1, \dots, w_k) ($w_0 = r_t$ and $w_k = r_{t'}$), then

since $\text{val}(w_i) = (\text{val}(w_{i+1}) + \text{val}(r^*)) / 2$, both w_{i+1} and r^* are ranked lower than t , $\text{val}(w_{i+1}) \leq \text{val}(r_t)$, $\text{val}(r^*) \leq \text{val}(r_t)$. But $\text{val}(w_0) = \text{val}(r_t)$, by induction for all i $\text{val}(w_i) = \text{val}(w_0) = \text{val}(r_t)$.

Now delete $r_{t'}$ and insert it back before r_t to get a new ordering \mathbf{g} . Define $f(v) = \text{val}[G_{\mathbf{f}}](v)$, then for graph $G_{\mathbf{g}}$ f is a valid value function that satisfy the requirements of Lemma 4.6. Therefore for any vertex v $\text{val}[G_{\mathbf{f}}](v) \geq \text{val}[G_{\mathbf{g}}](v)$. Either all values are equal, in this case $H(\mathbf{f}) = H(\mathbf{g})$ but $P(\mathbf{g}) > P(\mathbf{f})$ so $\mathbf{g} > \mathbf{f}$; or some values are different, in this case $H(\mathbf{g}) < H(\mathbf{f})$ so $\mathbf{g} > \mathbf{f}$.

Since there are only polynomially many ways to delete and reinsert an element, a better ordering can always be found in polynomial time. □

4.3 The Randomized Algorithm

Now we can use the existing randomized local minimum searching algorithm to solve the simple stochastic game. The following is the algorithm to solve the value of a simple stochastic game $G = (V, E)$:

1. Randomly choose $\sqrt{|V_R|!} \log(|V_R|!)$ permutations, and let \mathbf{f}_0 be the maximal permutation among them;
2. Starting from \mathbf{f}_0 , repeatedly find better permutations $\mathbf{f}_1, \mathbf{f}_2, \dots$ by Lemma 4.7 until a maximal permutation \mathbf{f}_t is found.

By Lemma 4.7, we can always find better permutations unless \mathbf{f} is maximal, and there are only $|V_R|!$ permutations, the algorithm will eventually find a maximal permutation and thus the optimal strategy. Now we analyze the running time of the algorithm:

The first step takes $\tilde{O}(\sqrt{|V_R|!})$ time, after that, each iteration of the loop will take $\text{poly}(|V_R|)$ time, so the key is how many iterations Step 2 needs.

Lemma 4.8 *The probability that Step 2 needs more than $\sqrt{|V_R|!}$ steps is no more than $1/(|V_R|!)$.*

Proof Consider any total ordering of the permutations that agrees with the partial ordering we defined. The probability that none of the $\sqrt{|V_R|!}$ largest elements are chosen is at most $(1 - \sqrt{|V_R|!}/(|V_R|!))^{\sqrt{|V_R|!} \log(|V_R|!)} = e^{-\log(|V_R|!)} = 1/(|V_R|!)$. □

Therefore, the expectation of number of iterations is at most $\sqrt{|V_R|!}$. The running time is $\tilde{O}(\sqrt{|V_R|!})$.

References

1. Condon, A.: The complexity of stochastic games. *Inf. Comput.* **96**(2), 203–224 (1992)
2. Condon, A.: On algorithms for simple stochastic games. In: *Advances in Computational Complexity Theory. DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, vol. 13, pp. 51–73. American Mathematical Society, Providence (1993)
3. Gimbert, H., Horn, F.: Solving simple stochastic games. In: *CiE '08: Proceedings of the 4th conference on Computability in Europe*, pp. 206–209. Springer, Berlin/Heidelberg (2008)

4. Walter, L.: A subexponential randomized algorithm for the simple stochastic game problem. *Inf. Comput.* **117**(1), 151–155 (1995)
5. Shapley, S.: Stochastic games. In: *Proceedings of the National Academy of Sciences* (1953)
6. Somla, R.: New algorithms for solving simple stochastic games. In: *Proceedings of the Workshop on Games in Design and Verification (GDV 2004)*. *Electronic Notes in Theoretical Computer Science*, vol. 119, pp. 51–65. Elsevier, Amsterdam (2005)