



Optimal Strategies in Turn-Based Stochastic Tail Games

Florian Horn

► To cite this version:

| Florian Horn. Optimal Strategies in Turn-Based Stochastic Tail Games. 2009. hal-00370046

HAL Id: hal-00370046

<https://hal.science/hal-00370046>

Preprint submitted on 23 Mar 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Strategies in Turn-Based Stochastic Tail Games^{*}

Florian Horn
f.horn@cwi.nl

Centrum Wiskunde & Informatica
Amsterdam, The Netherlands

Abstract. Infinite stochastic games are a natural model for open reactive processes: one player represents the controller, and the other represents a hostile environment. The evolution of the system depends on the decisions of the players, supplemented by a random function. The problems on such games can be sorted in two categories: the qualitative analysis ponders whether a player can win with probability one (or arbitrarily close to one), while the quantitative analysis is concerned about the maximal (or supremal) value a player can achieve.

In this paper, we establish the existence of optimal strategies in games whose the winning condition does not depend on finite prefixes. We also present a general procedure to derive quantitative results from qualitative algorithms. It also follows from the correctness of this procedure that optimal strategies are no more complex than almost-sure strategies.

1 Introduction

There is a long tradition of using infinite games to model open reactive processes [BL69,PR89]. The system is represented as a game arena, *i.e.* a graph whose vertices belong either to Eve (controller), Adam (non-deterministic environment), or Random (stochastic evolution). The game is played by moving a token on the arena: when it is in one of Eve's vertices, she chooses its next location among the successors of the current vertex; when it is in one of Adam's vertices, he chooses its next location; when it is in a random vertex, its next location is decided by a fixed random function. Playing the game for ω moves results in a play of the game, *i.e.* an infinite path of the graph. The specification of the system is represented by a (Borel) subset of the possible plays, the winning condition. Eve wins a play if it belongs to the winning condition, and Adam wins otherwise.

In this paper, we focus on “tail conditions”, where the winner of a play does not depend on finite prefixes. Our main motivation is that tail

^{*} This work was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme, and was also partially supported by the french ANR AVERISS

conditions subsume parity conditions. Hence, most of our results carry to regular games. From a verification perspective, tail conditions also correspond to cases where local glitches are tolerated in the beginning of a run, as long as the specification is met in the limit, *e.g.* in self-stabilising protocols. Finally, one of the most popular payoff functions in economic games, the mean-payoff function, is a tail condition. Due to lack of space, many proofs are sketched or omitted. Complete proofs can be found in the third chapter of [Hor08].

Outline of the paper. Section 2 recalls the classical notions about simple stochastic games. In Section 3, we show that the different qualitative criteria are equivalent in finite turn-based stochastic tail games, and define a new notion of *qualitative determinacy*. Section 4 takes on the quantitative problems, and shows how a qualitative algorithm can be used to compute the values of a finite turn-based stochastic tail game. The existence of optimal strategies for both players in finite turn-based stochastic tail games also follows from the proofs, as well as the fact that optimal strategies are no more complex than almost-sure strategies.

2 Definitions

We recall here several classical notions about simple stochastic games, and refer the reader to [GTW02] and [dA97] for more details.

Arenas and plays A *simple stochastic arena* \mathcal{A} is a directed graph $(\mathcal{Q}, \mathcal{T})$ without deadlocks, whose vertices are partitioned between Eve’s vertices (\mathcal{Q}_E , represented as \circ ’s), Adam’s vertices (\mathcal{Q}_A , represented as \square ’s), and random vertices (\mathcal{Q}_R , represented as \triangle ’s), and supplemented by a function $\delta : \mathcal{Q}_R \rightarrow \mathcal{D}(\mathcal{Q})$, which is the random law directing the choice of successors in the random vertices: so $\delta(r)(q) > 0 \Leftrightarrow (r, q) \in \mathcal{T}$. A sub-arena $\mathcal{A}|_B$ of \mathcal{A} is the restriction of \mathcal{A} to a subset B of \mathcal{Q} such that each controlled vertex of B has a successor in B , and all the successors of random vertices in B belong to B . A *play* ρ of \mathcal{A} is an (possibly infinite) path in the graph $(\mathcal{Q}, \mathcal{T})$. The set of infinite plays is denoted by Ω , and the set of infinite plays starting in the vertex q is denoted by Ω_q .

Strategies and measures A *pure strategy* σ for Eve is a deterministic way of extending finite plays ending in a vertex of Eve: $\sigma : \mathcal{Q}^* \mathcal{Q}_E \rightarrow \mathcal{Q}$ is such that $(q, \sigma(wq)) \in \mathcal{T}$. Strategies can also be defined as *strategies with memory*. Given a (possibly infinite) set of memory states M , a strategy σ with memory M is defined by two functions: a “next-move” function

$\sigma^n : (\mathcal{Q}_E \times M) \rightarrow \mathcal{Q}$ and a “memory-update” function $\sigma^u : (\mathcal{Q}_E \times M) \rightarrow M$. Notice that any strategy can be represented as a strategy with memory \mathcal{Q}^* . A play ρ is *consistent with a strategy* σ if and only if $\forall i, \rho_i \in \mathcal{Q}_E \Rightarrow \rho_{i+1} = \sigma(\rho_0, \dots, \rho_i)$. The set of plays consistent with σ is denoted by Ω^σ .

Once an initial vertex q and two strategies σ and τ have been fixed, $\Omega_q^{\sigma, \tau}$ can naturally be made into a measurable space $(\Omega_q^{\sigma, \tau}, \mathcal{O})$, where \mathcal{O} is the σ -field generated by the cones $\{\mathcal{O}_w \mid w \in \mathcal{Q}^*\}$: $\rho \in \mathcal{O}_w$ if and only if w is a prefix of ρ . The probability measure $\mathbb{P}_q^{\sigma, \tau}$ is recursively defined by:

$$\forall r \in \mathcal{Q}, \mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_r) = \begin{cases} 1 & \text{if } r = q, \\ 0 & \text{if } r \neq q; \end{cases}$$

$$\forall w \in \mathcal{Q}^*, (r, s) \in \mathcal{Q}^2, \mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_{wrs}) = \begin{cases} \mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_{wr}) \cdot \mathbf{1}_{\sigma(wr)=s} & \text{if } r \in \mathcal{Q}_E, \\ \mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_{wr}) \cdot \mathbf{1}_{\tau(wr)=s} & \text{if } r \in \mathcal{Q}_A, \\ \mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_{wr}) \cdot \delta(r)(s) & \text{if } r \in \mathcal{Q}_R. \end{cases}$$

Winning conditions and values A *winning condition* Φ is a Borel set of $(\Omega_q^{\sigma, \tau}, \mathcal{O})$. An infinite play is *winning for Eve* if it belongs to Φ , and *winning for Adam* otherwise. Finite plays are not winning for either player. A winning condition Φ is a *tail condition* if the winner of a play does not depend on finite prefixes: $\forall w \in \mathcal{Q}^*, \forall \rho \in \mathcal{Q}^\omega, \rho \in \Phi \Leftrightarrow w\rho \in \Phi$.

The *value of* $q \in \mathcal{Q}$ *with respect to the strategies* σ *and* τ *for Eve and Adam* (or $\{\sigma, \tau\}$ -*value*) is defined by: $v_{\sigma, \tau}(q) = \mathbb{P}_q^{\sigma, \tau}(\Phi)$. The *value of* q *with respect to a strategy* σ *for Eve* (or σ -*value*) is the infimum of its $\{\sigma, \tau\}$ -values: $v_\sigma(q) = \inf_\tau v_{\sigma, \tau}(q)$. Symmetrically, the *value of* q *with respect to a strategy* τ *for Adam* (or τ -*value*) is the supremum of its $\{\sigma, \tau\}$ -values: $v_\tau(q) = \sup_\sigma v_{\sigma, \tau}(q)$. By the quantitative determinacy of Blackwell games [Mar98], the supremum of the σ -values is equal to the infimum of the τ -values. This common value is called the *value of* q .

Winning criteria A strategy σ for Eve is *almost-surely winning* (or *almost-sure*) from a vertex q if and only if the σ -value of q is one. It is *positively winning* (or *positive*) from q if and only if for any strategy τ , the $\{\sigma, \tau\}$ -value of q is positive (notice that the σ -value of q may be zero). The *almost-sure region of Eve* (resp. *positive region of Eve*) is the set of vertices from which Eve has an almost-sure (resp. positive) strategy. The *limit-sure region of Eve* (resp. *bounded regions of Eve*) is the set of vertices with value one (resp. with positive value). In general, the limit-sure and almost-sure criteria are different, as are the positive and bounded criteria.

3 Qualitative Regions

In tail games, it is always possible for both players to disregard the history of a play, and consider that the current vertex is the initial one. We can thus use strategy translations to derive the value of a vertex from its owner and the value of its successors:

$$\begin{aligned} \forall q \in \mathcal{Q}_E, \mathbf{v}(q) &= \max\{\mathbf{v}(s) \mid (q, s) \in \mathcal{T}\} \\ \forall q \in \mathcal{Q}_A, \mathbf{v}(q) &= \min\{\mathbf{v}(s) \mid (q, s) \in \mathcal{T}\} \\ \forall q \in \mathcal{Q}_R, \mathbf{v}(q) &= \sum_{(q,s) \in \mathcal{T}} \delta(q)(s) \cdot \mathbf{v}(s) \end{aligned} \tag{1}$$

This is very similar to the case of reachability games, where such systems can directly be used to compute the values. Yet, there are two important differences: in general, tail games do not feature a “target vertex”, whose value is known to be one; and there is no notion of “stopping games”, where (1) has a unique solution. In order to establish our results, we need thus to consider the more complex notion of σ -value of a finite play:

Definition 1. *The σ -value of a finite play w consistent with σ is the infimum of the $\{\sigma, \tau\}$ -values under the assumption that w is a prefix of the play:*

$$v_\sigma(w) = \inf_{\tau} \mathbb{P}_{w_0}^{\sigma, \tau}(\Phi \mid \rho_0 = w_0, \rho_1 = w_1, \dots) .$$

Using the σ -values of the prefixes, we can observe how the prospects of the players evolve during a play. In particular, for any positive real number η , we define the event \mathcal{L}_η^σ , corresponding to the plays where Eve’s chances of winning have dropped below η at some point:

$$\mathcal{L}_\eta^\sigma = \{\exists i, v_\sigma(\rho_0 \dots \rho_i) \leq \eta\} .$$

This event has two interesting characteristics: first, if the σ -value of the initial vertex is greater than η , the probability that the ensuing play belongs to \mathcal{L}_η^σ is bounded away from one (Proposition 2); second, the probability that Adam wins is zero outside of \mathcal{L}_η^σ (Proposition 3).

Proposition 2. *Let q be a vertex of \mathcal{Q} , σ and τ be strategies for Eve and Adam, and $\eta < \nu \leq v_\sigma(q)$ be two positive real numbers. We have:*

$$\mathbb{P}_q^{\sigma, \tau}(\mathcal{L}_\eta^\sigma) \leq \frac{1 - \nu}{1 - \eta} .$$

Proposition 3. *Let q be a vertex of \mathcal{Q} , τ be a strategy for Adam, and η be a positive real number. We have:*

$$\mathbb{P}_q^{\sigma, \tau}(\Phi \mid \neg \mathcal{L}_\eta^\sigma) = 1 \text{ .}$$

These two results suggest a way to improve σ with a “reset” procedure with respect to a given real number η . Assume that Eve plays σ and, at some point, the σ -value of the prefix drops below η , while the σ -value of the current vertex is greater than η . She can improve her chances to win by forgetting the past, and restart playing σ as if the play just started.

Definition 4. *The strategy σ reset with respect to η , denoted by $\sigma_{\downarrow \eta}$, is a strategy with memory, whose memory states are the plays of \mathcal{A} consistent with σ . Its memory-update and next-move function are defined as follows:*

$$\begin{aligned} \sigma_{\downarrow \eta}^n(w, q) &= \begin{cases} \sigma(q) & \text{if } v_\sigma(wq) \leq \eta \wedge v_\sigma(q) > \eta \\ \sigma(wq) & \text{otherwise} \end{cases} \\ \sigma_{\downarrow \eta}^u(w, q) &= \begin{cases} q & \text{if } v_\sigma(wq) \leq \eta \wedge v_\sigma(q) > \eta \\ wq & \text{otherwise} \end{cases} \end{aligned}$$

If Eve plays according to $\sigma_{\downarrow \eta}$, it follows from Proposition 2 that the number of resets in the ensuing play is finite with probability one. Thus, by Proposition 3, either the σ -value of the visited vertices is consistently below η , or Eve wins with probability one. We can use reset strategies to expose several links between the different notions of “winning regions”.

Theorem 5. *In any finite turn-based stochastic tail game, Eve has an almost-sure winning strategy in the region with value one, and Adam has an almost-sure strategy in the region with value zero.*

Sketch of proof. Let us prove the result for Eve. As the arena is finite, we can choose η and ν such that $0 < \eta < \nu < 1$, and any vertex whose value is less than one is also less than η . Now, if σ has value ν from the vertices with value one, $\sigma_{\downarrow \eta}$ is almost-sure for Eve from these vertices. \square

Theorem 5 states that the *limit* and *almost-sure* winning criteria are equivalent in finite turn-based stochastic tail games. It follows directly that the *positive* and *bounded* winning criteria are also equivalent. Using the standard reduction to parity games, these results can be extended to finite simple stochastic ω -regular games. However, Theorem 5 does not hold for games with context-free conditions, infinite arenas, or concurrent moves: in each of the three games of Figure 1, the value of the initial vertex is one, yet Eve has no almost-sure strategy.

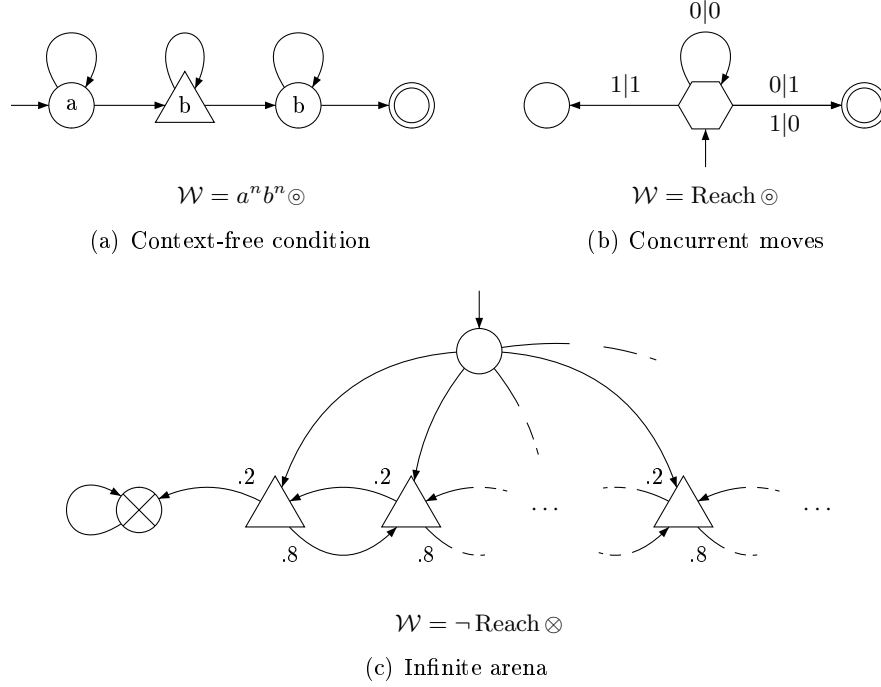


Fig. 1. Limit-sure is not almost-sure

We can also use reset strategies directly to derive a positive-almost property for finite turn-based stochastic tail games, extending Chatterjee’s bounded-limit property for finite concurrent tail games [Cha07a]¹:

Theorem 6 (Positive-almost property). *In any finite turn-based stochastic tail game, if Eve has a positive strategy from every vertex, she has an almost-sure strategy from every vertex. If she has a positive strategy from at least one vertex, she has an almost-sure strategy from at least one vertex. The same holds for Adam.*

Sketch of proof. In the proof of the “universal” part of Theorem 6 for Eve, the trick is to choose η and ν between zero and the lowest value for a vertex in the game (by Theorem 5, all vertices have positive value). Once again, if σ has value ν on the vertices with value one, $\sigma_{\downarrow\eta}$ is an almost-sure strategy for Eve. The other statements follow by duality. \square

¹ It is called a positive-limit property in the paper, but relies on the existence of a vertex with positive value: a “bounded” vertex, according to [dAH00]’s taxonomy.

Although the result is out of the scope of this work, a large part of this proof still holds in the more general case of finite concurrent games whose winning condition is suffix-closed. However, Theorem 6 itself does not: Figure 1 again provides counter-examples. But we could derive an existential positive-limit property for Eve and a universal bounded-almost property for Adam in these games:

Claim 7. *In any finite concurrent game with a suffix-closed winning condition, if Eve has a positive strategy from at least one vertex, then there is at least one vertex with value one. If no vertex has value one, Adam has an almost-sure strategy from every vertex.*

Last, but not least of our triptych is Theorem 8, which extends quantitative determinacy in prefix-independent games:

Theorem 8 (Qualitative determinacy). *In any finite turn-based stochastic tail game, from any vertex, either Eve has an almost-sure strategy or Adam has a positive strategy, and vice versa.*

Proof. Theorem 8 follows directly from Theorem 5 and the quantitative determinacy of Blackwell games [Mar98]. \square

By contrast with Theorem 5, we did not find any counter-example for natural extensions of Theorem 8. In particular, the three games of Figure 1 are qualitatively determined.

4 Values and optimal strategies

The algorithms computing the values of simple stochastic tail games are often adaptations of algorithms for reachability games which use qualitative algorithms as oracles. For example, one can guess a solution to (1) and use a qualitative algorithm to check necessary and sufficient conditions on the value regions: see [CdAH05] for Rabin and Streett games, [Cha07b] for Muller games, and [CHH08] for finitary games. It is also possible to adapt the strategy improvement algorithm of [HK66] when one of the players has positional strategies: see [CJH04] for parity, and [CH06] for Rabin games. Finally, in one-player stochastic tail games (Markov Decision Processes), one can compute first the almost-sure region, and then the values of the reachability game to this region [Cha07a].

In this section, we show how our permutation algorithms from [GH08] can be automatically modified to solve any finite turn-based stochastic

tail game. In fact, the resulting algorithms uses very similar concepts (and proofs). The main idea is that if Adam does not make obvious mistakes, Eve can only hope to win by reaching her almost-sure region. This can only be done through random vertices: there is no vertex of Eve leading to it (it would belong to the almost-sure region), and she cannot hope that Adam will enter it voluntarily (that would be an obvious mistake). We can thus consider the winning condition as a tool for Eve to ensure that the token reaches the best possible random vertex: if Adam refuses to comply, he loses. The behaviour of both players is then determined by their preferences over the random vertices. Furthermore, it is sufficient to consider the cases where Eve and Adam share the same estimation over the respective quality of random vertices, *i.e.*, when their preferences are opposed. These preferences are represented by permutations over the random vertices. In the remainder of the paper, a *permutation* π designates a permutation over the k random vertices, such that $\{\pi_1, \dots, \pi_k\} = \mathcal{Q}_R$. We often consider the sink and target vertices as random vertices in permutation-based concepts, with the implicit assumption that they are respectively the lowest and greatest vertices: $\pi_0 = \otimes$ and $\pi_{k+1} = \odot$.

For simplicity (*and* efficiency), we first *normalise* the games we consider: we compute the almost-sure regions of both players, and merge each of them into a single sink vertex (\otimes for Adam, \odot for Eve). The winning condition is modified accordingly: a play that reaches \otimes is winning for Adam and a play that reaches \odot is winning for Eve.

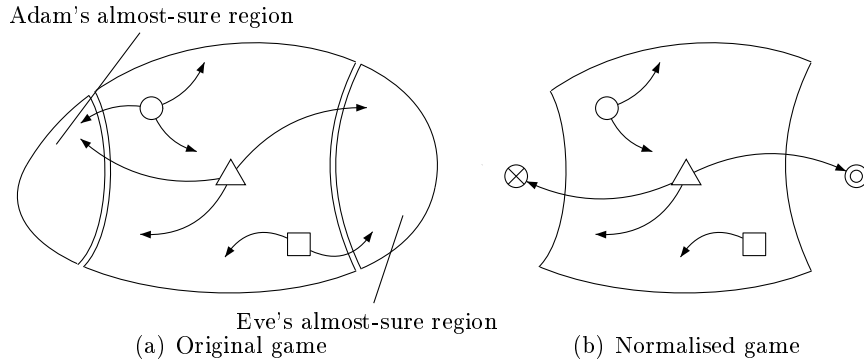


Fig. 2. Game normalisation

In each of our algorithms, the atomic loop considers a permutation π , and decides whether it is “correct”. The first question is to determine,

for each vertex, the best (with respect to π) random vertex that Eve can ensure to reach. We do so with the help of a qualitative oracle, which computes embedded almost-sure regions for Eve:

Definition 9. Let $\mathcal{G} = (\mathcal{A}, \Phi)$ be a normalised finite turn-based stochastic tail game, and π be a permutation over the k random vertices of \mathcal{A} . The π -regions of \mathcal{G} are defined as follows:

- $W_\pi[k+1] = \{\odot\}$;
- for any $1 \leq i \leq k$, $W_\pi[i]$ is the almost-sure region of Eve in \mathcal{A} with respect to the objective $\Phi \cup \text{Reach}(\cup_{j \geq i} \{\pi_j\})$, minus $\cup_{j > i} W_\pi[j]$;
- $W_\pi[0] = \{\otimes\}$.

Notice that a random vertex π_i may belong to a region $W_\pi[j]$ with $i < j$ (but not $i > j$). In this case, the region $W_\pi[i]$ is empty. Once the π -regions have been computed, we derive from them a Markov Chain \mathcal{G}^π , with $k+2$ vertices numbered $0 \dots k+1$: for any $i, j \in [0, k+1]$, the probability of going from π_i to π_j is equal to the probability of going from π_i to $W_\pi[j]$ in \mathcal{G} . We denote by $v_\pi[i]$ the value of π_i in \mathcal{G}^π .

We use two logical relations to decide about the pertinence of a permutation. *Self-consistency* is a most natural condition, as it simply expresses the adequation between *a priori* preferences, and resulting values:

Definition 10. Let \mathcal{G} be a finite turn-based stochastic tail game with k random vertices. A permutation π over \mathcal{Q}_R is self-consistent if for any $1 \leq i \leq j \leq k$, $v_\pi[i] \leq v_\pi[j]$.

It is easy to derive a solution to (1) from a self-consistent permutation. However, in general, there is more than one solution to this system, so we need another property, that we dub *liveness*:

Definition 11. Let \mathcal{G} be a finite turn-based stochastic tail game with k random vertices. A permutation π over \mathcal{Q}_R is live if for any $1 \leq i \leq k$, $\delta(\pi_i)(\cup_{j > i} W_\pi[j]) > 0$.

It may seem that this notion is already captured by self-consistency, as it is a “bad idea” for Eve to send the token to a random vertex that does not verify the internal property. However, the choice of the permutation also effects Adam’s behaviour [Mur07], and there may be spurious non-live self-consistent permutations.

Liveness and self-consistency are used in a straightforward way. In any finite turn-based stochastic tail game, there is a live and self-consistent permutation. Moreover, if a permutation π is live and self-consistent, then

for any vertex q in \mathcal{Q} , $q \in W_\pi[i] \Rightarrow \mathbf{v}(q) = v_\pi[i]$. It is then easy to derive algorithms computing the values from a qualitative algorithm:

Theorem 12. *Let \mathfrak{C} be a class of finite turn-based stochastic tail games. If the almost-sure region of Eve in a \mathfrak{C} -game \mathcal{G} can be computed in time $t(|\mathcal{G}|)$, then the values of any \mathfrak{C} -game \mathcal{G} can be computed in time $|\mathcal{Q}_R + 1|! \cdot t(|\mathcal{G}|)$.*

Sketch of proof. For any given permutation π , the π -regions can be computed in time $|\mathcal{Q}_R| \cdot t(|\mathcal{G}|)$. We can then decide whether π is live and self-consistent in linear time. We may need to do so for each of the $|\mathcal{Q}_R|!$ permutations, leading to a worst-time complexity of $|\mathcal{Q}_R + 1|! \cdot t(|\mathcal{G}|)$. \square

Theorem 13. *Let \mathfrak{C} be a class of finite turn-based stochastic tail games. If the problem of computing the almost-sure regions of \mathfrak{C} -games belongs to the complexity class \mathcal{K} , then the quantitative problems of \mathfrak{C} -games belongs to the classes $\text{NP}^\mathcal{K}$ and $\text{co-NP}^\mathcal{K}$.*

Sketch of proof. Instead of searching exhaustively for a live and self-consistent permutation, we can guess it non-deterministically, and check that it is correct in linear time with $|\mathcal{Q}_R|$ calls to a \mathcal{K} -oracle. \square

An interesting by-product of the proof is that the π -strategies derived from a live and self-consistent permutation are optimal:

Theorem 14. *In any finite turn-based stochastic tail game, both players have optimal strategies.*

As the parity acceptance condition, which can be used to represent any ω -regular language, is a tail condition, Theorem 14 also yields an alternative proof of the existence of optimal strategies in finite simple stochastic ω -regular games [dAH00].

It can also be noted that Eve's strategy is defined as a spatial composition of residually almost-sure strategies, and does not use more memory than its components:

Theorem 15. *Let \mathfrak{C} be a class of finite turn-based stochastic tail games. If Eve has almost-sure strategies with memory Υ in \mathfrak{C} -games, then she also has optimal strategies with memory Υ in \mathfrak{C} -games.*

Note that Theorem 15 does not hold when the winning condition is not a tail condition: see the weak parity game of Figure 3.

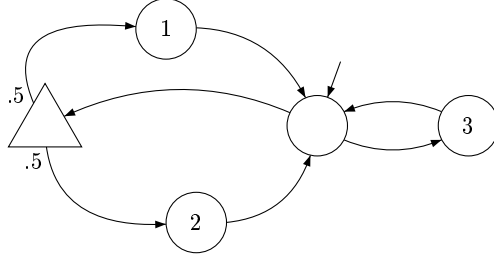


Fig. 3. Optimal strategies require memory in weak parity games

In this game, the value of the initial vertex is $\frac{1}{2}$. Indeed, if Eve sends the token once to the left and then always to the right, the lowest occurring colour has equal chances to be 1 or 2. However, this value cannot be achieved by means of a positional strategy:

- if there is a positive probability to send the token to the left, the lowest occurring colour is almost surely 1;
- otherwise, the lowest occurring colour is surely 3.

Both players have positional almost-sure strategies in weak parity games [GZ05]. Optimal strategies for weak parity games with d colours may require up to $d - 1$ memory states.

5 Conclusion

We proved the existence of optimal strategies for both players in all finite turn-based stochastic tail games. This also yields an alternative proof for the existence of optimal strategies in finite simple stochastic ω -regular games [dAH00]. Furthermore, we presented a single procedure to compute the values of finite turn-based stochastic tail games, provided that we already have a qualitative algorithm. The cost of this procedure is either a $|Q_R|!$ factor, or a non-deterministic guess, generalising several results on the complexity of quantitative problems. Once again, these results can be used to compute the values of ω -regular games, although it is necessary to first reduce them to equivalent parity games.

The existence of optimal strategies is very sensitive to each of our hypotheses, as demonstrated by Figure 1. However, the “qualitative determinacy” may hold in more general settings. It would also be interesting to look for a procedure to compute the values of simple stochastic games with arbitrary winning conditions, and/or infinite arenas.

References

- [BL69] J. R. Büchi and L. H. Landweber. Solving Sequential Conditions by Finite-State Strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
- [CdAH05] K. Chatterjee, L. de Alfaro, and T. A. Henzinger. The Complexity of Stochastic Rabin and Streett Games. In *Proceedings of ICALP'05*, volume 3580 of *Lecture Notes in Computer Science*, pages 878–890. Springer, 2005.
- [CH06] K. Chatterjee and T. A. Henzinger. Strategy Improvement for Stochastic Rabin and Streett Games. In *Proceedings of CONCUR'06*, volume 4137 of *Lecture Notes in Computer Science*, pages 375–389. Springer, 2006.
- [Cha07a] K. Chatterjee. Concurrent Games with Tail Objectives. *Theoretical Computer Science*, 388(1–2):181–198, 2007.
- [Cha07b] K. Chatterjee. Stochastic Müller Games are PSPACE-Complete. In *Proceedings of FSTTCS'07*, volume 4855 of *LNCS*, pages 436–448. Springer, 2007.
- [CHH08] K. Chatterjee, T. A. Henzinger, and F. Horn. Stochastic Finitary Games. Technical Report 2008–002, LIAFA, CNRS UMR 7089, 2008.
- [CJH04] K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Quantitative Stochastic Parity Games. In *Proceedings of SODA'04*, pages 121–130. Society for Industrial and Applied Mathematics, 2004.
- [dA97] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
- [dAH00] L. de Alfaro and Thomas A. Henzinger. Concurrent ω -regular Games. In *Proceedings of LICS'00*, pages 141–154. IEEE Computer Society, 2000.
- [GH08] H. Gimbert and F. Horn. Simple Stochastic Games with Few Random Vertices are Easy to Solve. In *Proceedings of FoSSaCS'08*, volume 4962 of *LNCS*, pages 5 – 19. Springer, 2008.
- [GTW02] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *LNCS*. Springer, 2002.
- [GZ05] H. Gimbert and W. Zielonka. Games Where You Can Play Optimally Without Any Memory. In *Proceedings of CONCUR'05*, volume 3653 of *LNCS*, pages 428–442. Springer, 2005.
- [HK66] A. J. Hoffman and R. M. Karp. On Nonterminating Stochastic Games. *Management Science*, 12(5):359–370, 1966.
- [Hor08] F. Horn. *Random Games*. PhD thesis, Université Paris 7 and RWTH Aachen, 2008.
- [Mar98] D. A. Martin. The Determinacy of Blackwell Games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- [Mur07] R. Muroe. Effect an Effect. *XKCD 326*, 2007.
- [PR89] A. Pnueli and R. Rosner. On the Synthesis of a Reactive Module. In *Proceedings of POPL'89*, pages 179–190, 1989.