Stochastic Games: A Tutorial

Marcin Jurdziński
Department of Computer Science
University of Warwick, UK

Game theory [1] is a formalism for the study of competitive interaction in the rich spectrum of relationships ranging between conflict and cooperation. Originally conceived as a mathematical foundation of economics, it proved its robustness by providing new techniques and insights in logic and set theory [15, 13], evolutionary and population biology [22], auction design and implementation, the design and study of the internet [19], analysis of cold war strategies, etc.

Dynamic games are used to model competitive processes evolving over time. *Stochastic* transitions are used for abstraction in modelling or to formalize inherent uncertainty, leading to quantitative statistical analysis. *Stochastic games* are dynamic games with stochastic transitions. They enjoy rich and mature mathematical theory [11, 23] and a wide range of applications including economics, cell, population and evolutionary biology [26], queueing theory and performance evaluation, and quantitative temporal model checking [17, 8].

In this tutorial we explore the spectrum of competitive dynamic behavioural models ranging from Markov chains, to transition systems, to Markov decision processes [20], to 2-player and multi-player perfect-information and general stochastic games [11]. We also consider the spectrum of *qualitative* and *quantitative* objectives for the players contained in the class of *Borel-measurable functions* [16], and we survey most popular concrete choices: *discounted* and *limiting-average* objectives [20, 11] used in economics and performance evaluation, and *omega-regular* objectives [13] encompassing safety, liveness and fairness objectives used in *temporal logic model checking*.

We present the general *determinacy* result for *zero-sum* stochastic games with Borel-measurable objectives due to Martin [16], as well as its many refinements for stochastic games with discounted, limiting-average, and omega-regular objectives [11, 6, 9, 17, 4, 28]. Then we survey algorithmic techniques for solving various classes of stochastic games and Markov decision processes, including *value iteration* and *strategy improvement* [20, 7], *convex optimization* [7, 27, 2] and graph-theoretic algorithms [3]. Special attention is given to algorithmic solutions of *(perfect-information) omega-regular stochastic games* [14, 3, 4] laying the foundation for automated *quantitative temporal model checking*.

We also consider *non-zero sum* stochastic games. After defining the fundamental notion of a *Nash equilibrium* we survey recent results [24, 21, 5] on its existence in various classes of stochastic games as well as a few challenging open problems in the area.

Then we present the *quantitative mu-calculus* [10, 17, 8] which is a very expressive, robust, and yet algorithmically feasible logical formalism for specification and automated quantitative analysis. We argue that the quantitative mu-calculus can be as fundamental to the quantitative temporal analysis as the modal mu-calculus is for the qualitative temporal model checking by showing how it subsumes various quantitative temporal logics proposed in literature. The unification of semantic and algorithmic techniques for quantitative specification and verification that the quantitative mu-calculus offers is an added benefit.

Finally, we present a game theoretic semantics for the quantitative mu-calculus [17] by reducing its satisfaction problem to solving stochastic omega-regular games. This allows to apply the rich theory of stochastic games both to the semantic study and to the algorithmic applications of the the quantitative mu-calculus in automated quantitative model checking.

Recent surveys of related material include invited presentations at CSL'02 [18], CONCUR'03 [8], STACS'04 [12] and LICS'04 [25].

References

[1] R. J. Aumann and S. Hart, editors. *Handbook of Game Theory with Economic Applications*, volume 3. Elsevier Science, 2002.

[2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[3] K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Simple stochastic parity games. In *CSL'03*, volume 2803 of *LNCS*, pages 100–113. Springer-Verlag, 2003.

[4] K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Quantitative stochastic parity games. In *SODA 2004*, pages 114–123. ACM/SIAM, 2004.

[5] K. Chatterjee, M. Jurdziński, and R. Majumdar. On Nash equilibria in stochastic games. Manuscript, 2003.

[6] A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.

[7] A. Condon. On algorithms for simple stochastic games. In Jin-Yi Cai, editor, *Advances in Computational Complexity Theory*, volume 13 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 51–73. American Mathematical Society, 1993.

[8] L. de Alfaro. Quantitative verification and control via the mu-calculus. In *CONCUR'03*, volume 2761 of *LNCS*, pages 102–126. Springer-Verlag, 2003.

[9] L. de Alfaro and T. A. Henzinger. Concurrent ω-regular games. In *LICS'00*, pages 141–154. IEEE Computer Society Press, 2000.

[10] L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *Journal of Computer and System Sciences*, 68:374–397, 2004.

[11] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.

[12] E. Grädel. Positional determinacy of infinite games. In *STACS'04*, volume 2996 of *LNCS*, pages 4–18. Springer-Verlag, 2004.

[13] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games*, volume 2500 of *LNCS*. Springer-Verlag, 2002.

[14] M. Jurdziński, O. Kupferman, and T. A. Henzinger. Trading probability for fairness. In *CSL'02*, volume 2471 of *LNCS*, pages 292–305. Springer-Verlag, 2002.

[15] D. A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.

[16] D. A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.

[17] A. K. McIver and C. C. Morgan. Games, probability and the quantitative mu-calculus $qM\mu$. In *LPAR'02*, volume 2514 of *LNCS*, pages 292–310. Springer-Verlag, 2002.

[18] D. Niwiński. μ-calculus via games. In *CSL'02*, volume 2471 of *LNCS*, pages 27–43. Springer-Verlag, 2002.

[19] C. Papadimitriou. Algorithms, games, and the internet. In *STOC'01*, pages 749–753. ACM Press, 2001.

[20] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.

[21] P. Secchi and W. D. Sudderth. Stay-in-a-set games. *International Journal of Game Theory*, 30(4):479–490, 2002.

[22] J. M. Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.

[23] S. Sorin. *A First Course on Zero-Sum Repeated Games*, volume 37 of *Mathématiques & Applications*. Springer-Verlag, 2002.

[24] N. Vieille. Stochastic games: Recent results. In *Handbook of Game Theory*, pages 1833–1850. Elsevier Science, 2002. Chapter 48 of [1].

[25] I. Walukiewicz. A landscape with games in the background. In *LICS'04*. IEEE Computer Society Press, 2004.

[26] D. M. Wolf and A. P. Arkin. Motifs, modules and games in bacteria. *Current Opinion in Microbiology*, 6:125–134, 2003.

[27] Y. Ye. A new complexity result on solving the Markov decision problem. Manuscript, April 2003.

[28] W. Zielonka. Perfect-information stochastic parity games. In *FOSSACS'04*, volume 2987 of *LNCS*, pages 499–513. Springer-Verlag, 2004.