

1 Probabilistic Model Checking for 2 Strategic Equilibria-based Decision Making: 3 Advances and Challenges

4 **Marta Kwiatkowska** ✉ 

5 University of Oxford, Oxford, UK

6 **Gethin Norman** ✉ 

7 University of Glasgow, Glasgow, UK

8 University of Oxford, Oxford, UK

9 **David Parker** ✉ 

10 University of Birmingham, Birmingham, UK

11 **Gabriel Santos** ✉ 

12 University of Oxford, Oxford, UK

13 **Rui Yan** ✉ 

14 University of Oxford, Oxford, UK

15 — Abstract —

16 Game-theoretic concepts have been extensively studied in economics to provide insight into competit-
17 18 19 20 21 22 23 24 25 26 27 28
ive behaviour and strategic decision making. As computing systems increasingly involve concurrently
acting autonomous agents, game-theoretic approaches are becoming widespread in computer science
as a faithful modelling abstraction. These techniques can be used to reason about the competitive
or collaborative behaviour of multiple rational agents with distinct goals or objectives. This paper
provides an overview of recent advances in developing a modelling, verification and strategy syn-
thesis framework for concurrent stochastic games implemented in the probabilistic model checker
PRISM-games. This is based on a temporal logic that supports finite- and infinite-horizon temporal
properties in both a zero-sum and nonzero-sum setting, the latter using Nash and correlated equilibria
with respect to two optimality criteria, social welfare and social fairness. We summarise the key
concepts, logics and algorithms and the currently available tool support. Future challenges and
recent progress in adapting the framework and algorithmic solutions to continuous environments
and neural networks are also outlined.

29 **2012 ACM Subject Classification** Theory of computation

30 **Keywords and phrases** Probabilistic model checking, stochastic games, equilibria

31 **Digital Object Identifier** 10.4230/LIPIcs.MFCS.2022.4

32 **Category** Invited Talk

33 **Funding** This project was funded by the ERC under the European Union’s Horizon 2020 research
34 and innovation programme (FUN2MODEL, grant agreement No. 834115).

35 **1** Introduction

36 Game-theoretic techniques have long been a source of fundamental insights into strategic
37 38 39 40 41 42
decision making for multi-agent systems. They have been widely studied in areas such as
economics [27], control [43] and robotics [36]. Concurrent stochastic multi-player games
(CSGs), in particular, provide a natural framework for modelling a set of interactive, rational
agents operating concurrently within an uncertain or stochastic environment. They can be
viewed as a collection of players (agents) with strategies for determining their actions based
on the execution so far, and where the resulting evolution of the system is probabilistic.



© Marta Kwiatkowska and Gethin Norman and David Parker and Gabriel Santos and Rui Yan;
licensed under Creative Commons License CC-BY 4.0

47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022).

Editors: Stefan Szeider, Robert Ganian, and Alexandra Silva; Article No. 4; pp. 4:1–4:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

43 Game-theoretic analysis is versatile, in that it can support both zero-sum and nonzero-
44 sum (equilibria) analysis. Zero-sum properties focus on scenarios in which one player (or
45 a coalition of players) aims to optimise some objective, while the remaining players form
46 a coalition with the directly opposing goal. On the other hand, nonzero-sum (equilibria)
47 properties correspond to situations where two or more players (or coalitions of players) in a
48 CSG have distinct objectives to be maximised or minimised. In nonzero-sum properties the
49 goals of the players (or coalitions) are not necessarily directly opposing, and therefore it may
50 be beneficial for players to collaborate. Competitive scenarios occur in many applications,
51 e.g., attackers and defenders in the context of computer security. Similarly, collaborative
52 behaviour can be essential, e.g., to effectively control a multi-robot system, or for users to
53 send data efficiently through a shared medium in a communication protocol.

54 Probabilistic model checking is a powerful approach to the formal analysis of systems
55 with stochastic behaviour. It relies on the construction and analysis of a probabilistic model,
56 guided by a formal specification of its desired behaviour in temporal logic. It is of particular
57 benefit in the context of models, such as stochastic games, which combine nondeterministic
58 and probabilistic behaviour. This is because the interplay between these aspects of the model
59 can be subtle and lead to unexpected results if not carefully modelled and analysed. This is
60 exacerbated when the system comprises multiple agents with differing objectives.

61 Until recently, practical applications of probabilistic model checking based on stochastic
62 games had focused primarily on *turn-based* models [15], in which simultaneous decision
63 making by agents is forbidden. Alternatively, model checking of *non-stochastic* games has
64 been extensively studied, and tool support developed [5, 39]. CSGs provide a more powerful
65 and realistic modelling formalism, but also bring considerable challenges, in terms of the
66 higher computational complexity or undecidability for some key problems.

67 There has nonetheless been significant amounts of work on tackling verification problems
68 for CSGs. A number of algorithms have been proposed for solving CSGs against formally
69 specified zero-sum properties, e.g. [17, 18, 11]. In the case of nonzero-sum properties,
70 [14, 25] study the existence of and the complexity of finding equilibria for stochastic games.
71 Complexity results for finding equilibria are also considered in [9] and [23] for quantitative
72 reachability properties and temporal logic properties, respectively. Other work concerns
73 finding equilibria for discounted properties; we mention [49], which formulates a learning-
74 based algorithm, and [40], which presents iterative algorithms. However these advances are
75 mostly lacking in implementations, tool support or case studies. Tools exist for solving turn-
76 based stochastic games [13, 16] and non-stochastic concurrent games [16, 8, 10, 55, 24, 45],
77 with the latter class including support for computing equilibria.

78 At the same time, there is an increasing trend to incorporate data-driven decision making,
79 which necessitates the incorporation on of machine learning components within autonomous
80 systems, which are built largely using conventional, symbolic methods. Examples of such
81 *neuro-symbolic* systems are self-driving cars whose vision function is provided via a neural
82 network image classifier, or an aircraft controller whose collision avoidance system uses a
83 neural network for decision support. Design automation support for such systems is lacking,
84 yet automatic computation of equilibria aids in ensuring stable solutions.

85 This paper provides an overview of recent advances in developing a modelling, verification
86 and strategy synthesis framework for concurrent stochastic games, as implemented in the
87 PRISM-games probabilistic model checker [33]. The framework uses a temporal logic that
88 supports a wide range of finite- and infinite-horizon properties, relating to the probability
89 of an event's occurrence or the expected amount of reward or cost accumulated. The logic
90 allows specification of both *zero-sum* and *nonzero-sum* properties, with the latter expressed

91 using either *Nash equilibria* or *correlated equilibria*. For both types of equilibria, strategies
 92 are synthesised in which it is not beneficial for any player to unilaterally alter their chosen
 93 actions, but correlated equilibria also allow players to coordinate through *public signals*.
 94 Since several, varied such equilibria may exist, we also support distinct optimality criteria
 95 to select between them; we consider *social welfare*, which maximises the sum of the players
 96 utilities, and *social fairness*, which minimises the difference between the utilities.

97 We summarise the key concepts, logics and algorithms that underlie this framework and
 98 discuss the tool support provided by PRISM-games, including an illustrative case study
 99 of formally modelling and analysing a multi-agent communication protocols using CSGs.
 100 Future challenges and recent progress in extending the framework and algorithmic solutions
 101 to modelling of neuro-symbolic CSGs are also outlined. In contrast to the majority of prior
 102 research, the focus of this strand of work is on software tool development, applications and
 103 case studies.

104 2 Normal form games

105 We introduce the main concepts used in this paper by means of simple one-shot games known
 106 as *normal form games* (NFGs), where players make their choices at the same time. We
 107 consider both zero-sum NFGs and nonzero-sum NFGs, then define equilibria concepts for
 108 these games and summarise existing algorithms for equilibria computation.

109 We first require the following notation. Let $Dist(X)$ denote the set of probability
 110 distributions over set X . For any vector $v \in \mathbb{R}^n$, we use $v(i)$ to refer to the i th entry of the
 111 vector. For any tuple $x = (x_1, \dots, x_n) \in X^n$, element $x' \in X$ and $i \leq n$, we define the tuples
 112 $x_{-i} \stackrel{\text{def}}{=} (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and $x_{-i}[x'] \stackrel{\text{def}}{=} (x_1, \dots, x_{i-1}, x', x_{i+1}, \dots, x_n)$.

113 ► **Definition 1** (Normal form game). A (finite, n -player) normal form game (NFG) is a tuple
 114 $\mathbf{N} = (N, A, u)$ where:

- 115 ■ $N = \{1, \dots, n\}$ is a finite set of players;
- 116 ■ $A = A_1 \times \dots \times A_n$ and A_i is a finite set of actions available to player $i \in N$;
- 117 ■ $u = (u_1, \dots, u_n)$ and $u_i: A \rightarrow \mathbb{R}$ is a utility function for player $i \in N$.

118 In a normal form game \mathbf{N} , the players choose actions simultaneously, with player $i \in N$
 119 choosing an action from the set A_i and, assuming that each player $i \in N$ selects action
 120 a_i , player j receives the utility $u_j(a_1, \dots, a_n)$. The *objective* of each player is to maximise
 121 their utility and their choices are governed by *strategies*, which we now define. We will also
 122 distinguish *strategy profiles*, which comprise a strategy for each player, and *correlated profiles*,
 123 which correspond to choices of the players when they are allowed to coordinate through a
 124 (probabilistic) *public signal*.

125 ► **Definition 2** (Strategies, profiles and correlated profiles). For an NFG \mathbf{N} :

- 126 ■ a strategy σ_i for player i in an NFG \mathbf{N} is a probability distribution over the set of actions
 127 A_i and we let $\Sigma_{\mathbf{N}}^i$ denote the set of all strategies for player i ;
- 128 ■ a strategy profile (or profile) $\sigma = (\sigma_1, \dots, \sigma_n)$ is a tuple of strategies for each player;
- 129 ■ a correlated profile is a tuple (τ, ς) comprising $\tau \in Dist(D_1 \times \dots \times D_n)$, where D_i is a
 130 finite set of signals for player i , and $\varsigma = (\varsigma_1, \dots, \varsigma_n)$, where $\varsigma_i: D_i \rightarrow A_i$ is a function
 131 from the signals of player i to the actions of player i .

132 For a correlated profile (τ, ς) of \mathbf{N} , the public signal τ is a joint distribution over signals D_i
 133 for each player i such that, if player i receives the signal $d_i \in D_i$, then it chooses action

134 $\varsigma_i(d_i)$. We can consider any correlated profile (τ, ς) as a *joint strategy*, i.e., a distribution
 135 over $A_1 \times \dots \times A_n$ where:

$$136 \quad (\tau, \varsigma)(a_1, \dots, a_n) = \sum \{ \tau(d_1, \dots, d_n) \mid d_i \in D_i \wedge \varsigma(d_i) = a_i \text{ for all } i \in N \}.$$

137 Conversely, any joint strategy $\tau \in \text{Dist}(A_1 \times \dots \times A_n)$ of \mathbf{N} can be considered as a correlated
 138 profile (τ, ς) , where $D_i = A_i$ and ς_i is the identity function for $i \in N$. Any profile σ of an
 139 NFG \mathbf{N} can be mapped to an equivalent correlated profile (in which τ is the joint distribution
 140 $\sigma_1 \times \dots \times \sigma_n$ and ς_i is the identity function). On the other hand, there are correlated profiles
 141 with no equivalent strategy profile.

142 Under profile σ or correlated profile (τ, ς) the expected utilities of player i are:

$$143 \quad \begin{aligned} u_i(\sigma) &\stackrel{\text{def}}{=} \sum_{(a_1, \dots, a_n) \in A} u_i(a_1, \dots, a_n) \cdot \left(\prod_{j=1}^n \sigma_j(a_j) \right) \\ u_i(\tau, \varsigma) &\stackrel{\text{def}}{=} \sum_{(d_1, \dots, d_n) \in D} \tau(d_1, \dots, d_n) \cdot u_i(\varsigma_1(d_1), \dots, \varsigma_n(d_n)). \end{aligned}$$

144

145 ► **Example 3.** Consider the two-player NFG with available action sets $A_i = \{\text{heads}_i, \text{tails}_i\}$
 146 for $1 \leq i \leq 2$ and a correlated profile corresponding to the joint distribution $\tau \in \text{Dist}(A_1 \times A_2)$,
 147 where $\tau(\text{heads}_1, \text{heads}_2) = \tau(\text{tails}_1, \text{tails}_2) = 0.5$. Under this correlated profile, the players
 148 share a fair coin and choose their action based on the outcome of the coin toss. There is no
 149 equivalent strategy profile.

150 2.1 Zero-sum NFGs

151 A *zero-sum NFG* is a two-player NFG \mathbf{N} such that $u_1(\alpha) + u_2(\alpha) = 0$ for all $\alpha \in A$, meaning
 152 that the objectives of the players are directly opposing. Such an NFG is often called a
 153 *matrix game*, as it can be represented by a single matrix $Z \in \mathbb{Q}^{l \times m}$, where $A_1 = \{a_1, \dots, a_l\}$,
 154 $A_2 = \{b_1, \dots, b_m\}$ and $z_{ij} = u_1(a_i, b_j) = -u_2(a_i, b_j)$.

155 We next introduce the notion of the *value* of a zero-sum NFG and recall classical results
 156 about the existence of optimal strategies.

157 ► **Theorem 4** (Minimax theorem [56, 57]). *For any zero-sum NFG $\mathbf{N} = (N, A, u)$ and
 158 corresponding matrix game Z , there exists $v^* \in \mathbb{Q}$, called the value of the game and denoted
 159 $\text{val}(Z)$, such that:*

- 160 ■ *there is a strategy σ_1^* for player 1, called an optimal strategy of player 1, such that under
 161 this strategy the player's expected utility is at least v^* regardless of the strategy of player
 162 2, i.e., $\inf_{\sigma_2 \in \Sigma_N^2} u_1(\sigma_1^*, \sigma_2) \geq v^*$;*
- 163 ■ *there is a strategy σ_2^* for player 2, called an optimal strategy of player 2, such that under
 164 this strategy the player's expected utility is at least $-v^*$ regardless of the strategy of player
 165 1, i.e., $\inf_{\sigma_1 \in \Sigma_N^1} u_2(\sigma_1, \sigma_2^*) \geq -v^*$.*

166 The value of a matrix game $Z \in \mathbb{Q}^{l \times m}$ can be found by solving a linear programming (LP)
 167 problem [56, 57].

168 ► **Example 5.** Table 1 shows a classic example of a two-player zero-sum game known as
 169 *matching pennies*. Columns α and u_i represent the collective choice (profile) and player i 's
 170 utility, respectively. In this example, each player has a coin for which they may choose the
 171 value to be heads or tails, i.e., $A_i = \{\text{heads}_i, \text{tails}_i\}$. If the coins match, player 1 wins the
 172 round, which is indicated by being awarded a utility of 1, while player 2 receives utility -1 .
 173 If the coins do not match, then the players' utilities are negated.

α	$u_1(\alpha)$	$u_2(\alpha)$	α	$u_1(\alpha)$	$u_2(\alpha)$
$(heads_1, heads_2)$	1	-1	$(tails_1, heads_2)$	-1	1
$(heads_1, tails_2)$	-1	1	$(tails_1, tails_2)$	1	-1

■ **Table 1** Matching pennies game in normal form.

174 The value for the corresponding matrix game is the solution to the following LP problem:
 175 Maximise v subject to:

176
$$x_1 - x_2 \geq v, \quad x_2 - x_1 \geq v, \quad x_1 + x_2 = 1$$

177 which yields the value $v^* = 0$ with optimal strategy $\sigma_1^* = (\frac{1}{2}, \frac{1}{2})$ for player 1 (the optimal
 178 strategy for player 2 is the same).

179 2.2 Nonzero-sum NFGs

180 The requirement for players to have directly opposing objectives is often too limiting, and it is
 181 necessary to allow distinct objectives, which cannot be modelled in a zero-sum fashion. These
 182 scenarios can be captured using the notion of *equilibria*, defined by a separate, independent
 183 objective for each agent. We now define the concepts of *Nash equilibrium* [57] and *correlated*
 184 *equilibrium* [6] for NFGs, which ensure stability against deviations by individual agents,
 185 improving the overall game outcomes. Since many equilibria may exist, we also introduce
 186 optimality criteria for these equilibria: *social welfare*, which is standard [46], and *social*
 187 *fairness*, which was first defined in [35].

188 Before giving the formal definitions, we first extend our notation as follows: for any
 189 profile σ and strategy σ_i^* , the strategy tuple σ_{-i} corresponds to σ with the strategy of player
 190 i removed and $\sigma_{-i}[\sigma_i^*]$ to the profile σ after replacing player i 's strategy with σ_i^* .

191 ► **Definition 6** (Best response). *For any nonzero-sum NFG N and profile σ or correlated*
 192 *profile (τ, ς) of N , the best response moves for player i to σ_{-i} and (τ, ς_{-i}) are, respectively:*


- 193 ■ *a strategy σ_i^* for player i such that $u_i(\sigma_{-i}[\sigma_i^*]) \geq u_i(\sigma_{-i}[\sigma_i])$ for all $\sigma_i \in \Sigma_N^i$;*
- 194 ■ *a function $\varsigma_i^*: D_i \rightarrow A_i$ for player i such that $u_i(\tau, \varsigma_{-i}[\varsigma_i^*]) \geq u_i(\tau, \varsigma_{-i}[\varsigma_i])$ for all functions*
 195 *$\varsigma_i: D_i \rightarrow A_i$.*

196 ► **Definition 7** (NE and CE). *For any nonzero-sum NFG N , a strategy profile σ^* is a Nash*
 197 *equilibrium (NE) and a correlated profile (τ, ς^*) of N is a correlated equilibrium (CE) if:*

- 198 ■ *σ_i^* is a best response to σ_{-i}^* for all $i \in N$;*
- 199 ■ *ς_i^* is a best response to (τ, ς_{-i}^*) for all $i \in N$;*
 200 *respectively.*

201 Any NE of N is also a CE, while there exist CEs that cannot be represented by a strategy
 202 profile, and therefore are not NEs. For each class of equilibria, NE and CE, we introduce
 203 two optimality criteria, the first maximising *social welfare* (SW), defined as the *sum* of the
 204 utilities, and the second maximising *social fairness* (SF), which minimises the *difference*
 205 between the players' utilities. Other variants of fairness have been considered for NEs, such
 206 as in [38], where the authors seek to maximise the lowest utility among the players.

207 ► **Definition 8** (SW and SF). *An equilibrium σ^* is a social welfare (SW) equilibrium if the*
 208 *sum of the utilities of the players under σ^* is maximal over all equilibria, while σ^* is a social*
 209 *fair (SF) equilibrium if the difference between the player's utilities under σ^* is minimised*
 210 *over all equilibria.*



α	$u_1(\alpha)$	$u_2(\alpha)$	$u_3(\alpha)$
(pro_1, pro_2, pro_3)	-1000	-1000	-100
(pro_1, pro_2, yld_3)	-1000	-100	-5
(pro_1, yld_2, pro_3)	5	-5	5
(pro_1, yld_2, yld_3)	5	-5	-5
(yld_1, pro_2, pro_3)	-5	-1000	-100
(yld_1, pro_2, yld_3)	-5	5	-5
(yld_1, yld_2, pro_3)	-5	-5	5
(yld_1, yld_2, yld_3)	-10	-10	-10

■ **Figure 1** Example from [35]: Cars at an intersection and the corresponding NFG.

211 We can also define the dual concept of *social cost* (SC) equilibria [34], where players try
 212 to minimise, rather than maximise, their expected utilities by considering equilibria of the
 213 game $N^- = (N, A, -u)$ in which the utilities of N are negated. We remark that SC equilibria
 214 strategies are not a subset of classically defined NE or CE strategies of N .

215 ► **Example 9.** Consider the scenario from [35], based on an example from [50], where three
 216 cars meet at an intersection and want to proceed as indicated by the arrows in Figure 1.
 217 Each car can either *proceed* or *yield*. If two cars with intersecting paths proceed, then there
 218 is an accident. If an accident occurs, the car having the right of way, i.e., the other car is
 219 to its left, has a utility of -100 and the car that should yield has a utility of -1000 . If a
 220 car proceeds without causing an accident, then its utility is 5 and the cars that yield have
 221 a utility of -5 . If all cars yield, then, since this delays all cars, all have utility -10 . The
 222 3-player NFG is given in Figure 1. The different optimal equilibria of the NFG are:

- 223 ■ the SWNE and SWCE are the same: for c_2 to yield and c_1 and c_3 to proceed, with the
 224 expected utilities of the players $(5, -5, 5)$;
- 225 ■ the SFNE is for c_1 to yield with probability 1 , c_2 to yield with probability 0.863636
 226 and c_3 to yield with probability 0.985148 , with the expected utilities of the players
 227 $(-9.254050, -9.925742, -9.318182)$;
- 228 ■ the SFCE gives a joint distribution where the probability of c_2 yielding and of c_1 and c_3
 229 yielding are both 0.5 with the expected utilities of the players $(0, 0, 0)$.

230 Modifying u_2 such that $u_2(pro_1, pro_2, pro_3) = -4.5$ to, e.g., represent a reckless driver, the
 231 SWNE becomes for c_1 and c_3 to yield and c_2 to proceed with the expected utilities of the
 232 players $(-5, 5, -5)$, while the SWCE is still for c_2 to yield and c_1 and c_3 to proceed. The
 233 SFNE and SFCE also do not change.

234 **Algorithms for computing equilibria in NFGs.** Finding NEs in two-player NFGs is in
 235 the class of *linear complementarity* problems (LCPs). Established algorithms include the
 236 Lemke-Howson algorithm [37], which is based on the method of labelled polytopes [46],
 237 support enumeration [48] and regret minimisation [52]. In [34] a method for NE computation
 238 is developed, which reduces the problem to SMT via labelled polytopes [46] by considering
 239 the regions of the strategy profile space. This method iteratively reduces the search space
 240 of profiles as positive probability assignments are found and added as constraints on the
 241 profiles. This approach can also be used for finding both an SWNE and SFNE by computing
 242 all NEs and then selecting an optimal one.

243 In the case of NFGs with more than two players, the computation of NEs is more
 244 complex since, for a given support (i.e., a sub-region of the strategy profile space which fixes
 245 the set of actions chosen with nonzero probability by each player), finding NEs cannot be

reduced to an LP problem. A method for such NFGs is presented in [32], based on support enumeration [48], which exhaustively examines all supports one at a time, checking whether that sub-region contains NEs. For each support, finding an SWNE can be reduced to a *nonlinear programming problem* [32]. This nonlinear programming problem can be modified to find an SFNE in each support [35].

In the case of CEs, the approach introduced in [35] is to first find a joint strategy for the players, i.e., a distribution over the action tuples, which can then be mapped to a correlated profile. For SWCEs, [35] reduces the computation to solving a LP problem which has $|A|$ variables, one for each action tuple, and $\sum_{i \in N} (|A_i|^2 - |A_i|) + |A| + 1$ constraints. For SFCEs, on the other hand, the method of [35] involves solving an optimisation problem with an additional $|N| + 2$ variables and $3 \cdot |N|$ constraints compared to the LP problem for finding SWCEs.

3 Concurrent Stochastic Games

This section introduces *concurrent stochastic games* (CSGs) [54], in which players repeatedly make simultaneous choices over actions and the action choices cause a probabilistic update of the game state. CSGs thus provide a natural framework for modelling a set of interactive, rational agents operating concurrently within an uncertain or probabilistic environment. Compared to normal form games, they are classified as *multi-stage*, which is more convenient for specifying repeated or sequential interactions among agents. The introduction of stochasticity facilitates modelling of a wide range of important phenomena, for example uncertain behaviour due to noisy sensors or unreliable hardware in a multi-robot system, or the use of randomisation for coordination in a distributed security or networking protocol.

► **Definition 10** (Concurrent stochastic game). A concurrent stochastic multi-player game (CSG) is a tuple $G = (N, S, \bar{s}, A, \Delta, \delta)$ where:

- $N = \{1, \dots, n\}$ is a finite set of players;
- S is a finite set of states and $\bar{s} \in S$ is an initial state;
- $A = (A_1 \cup \{\perp\}) \times \dots \times (A_n \cup \{\perp\})$ where A_i is a finite set of actions available to player $i \in N$ and \perp is an idle action disjoint from the set $\cup_{i=1}^n A_i$;
- $\Delta: S \rightarrow 2^{\cup_{i=1}^n A_i}$ is an action assignment function;
- $\delta: S \times A \rightarrow \text{Dist}(S)$ is a probabilistic transition function.

Given a CSG G , the set of actions available to player $i \in N$ in state $s \in S$ is given by $A_i(s) \stackrel{\text{def}}{=} \Delta(s) \cap A_i$. The CSG G starts in the initial state \bar{s} and, if G is in the game state s , then each player $i \in N$ selects an action from its available actions in state s if this set is non-empty, and from $\{\perp\}$ otherwise. Next, supposing each player $i \in N$ chooses action a_i , the game state is updated according to the distribution $\delta(s, (a_1, \dots, a_n))$. We allow sets of players $C \subseteq N$ to form *coalitions*, and will consider the induced CSG, called the *coalition game*, with coalitions as players.

A *path* π of a CSG G is a sequence $\pi = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots$, where $s_i \in S$, $\alpha_i \in A$ and $\delta(s_i, \alpha_i)(s_{i+1}) > 0$ for all $i \geq 0$. We denote by $FPaths_{G,s}$ and $IPaths_{G,s}$ the sets of finite and infinite paths starting in state s of G , respectively, and drop the subscript s when considering all finite and infinite paths of G . As for NFGs, we can define *strategies* of G that resolve the choices of the players. Here, a strategy for player i is a function $\sigma_i: FPaths_G \rightarrow \text{Dist}(A_i \cup \{\perp\})$ mapping finite paths to distributions over available actions, such that, if $\sigma_i(\pi)(a_i) > 0$, then $a_i \in A_i(\text{last}(\pi))$ where $\text{last}(\pi)$ is the final state of π . Furthermore, we can define strategy profiles, correlated profiles and joint strategies analogously to Section 2.

291 A *labelled* CSG is a tuple (G, AP, L) , where G is a CSG, as in Definition 10, AP is a
 292 set of atomic propositions and $L: S \rightarrow 2^{AP}$ is a labelling function, specifying which atomic
 293 propositions are true in each state. We also associate CSGs with *reward structures*, which
 294 annotate states and transitions with real values. More precisely a reward structure is a pair
 295 $r=(r_A, r_S)$ consisting of an action reward function $r_A: S \times A \rightarrow \mathbb{R}$ and state reward function
 296 $r_S: S \rightarrow \mathbb{R}$. We use atomic propositions and rewards as the building blocks to specify players'
 297 utilities in a CSG, which will be described in Section 4.

298 Formally, the utility function or *objective* of player i in a CSG is given by a random
 299 variable $X_i: IPaths_G \rightarrow \mathbb{R}$ over infinite paths. For a profile σ and state s , using standard
 300 techniques [31], we can construct a probability measure $Prob_{G,s}^\sigma$ over the paths that start
 301 in state s corresponding to σ , denoted $IPaths_{G,s}^\sigma$, and define the expected value $\mathbb{E}_{G,s}^\sigma(X_i)$ of
 302 player i 's utility from s under σ . Similarly, we can also define such a probability measure
 303 and expected value given a correlated profile or joint strategy of G .

304 3.1 Zero-sum CSGs

305 Similarly to NFGs (see Section 2.1), *zero-sum* CSGs are two-player games that have a single
 306 utility function X for player 1, with the utility function of player 2 given by $-X$, and both
 307 players aiming to maximise the expected value of their utility. Equivalently, we can suppose
 308 that player 1 tries to maximise the expected value of X , while player 2 tries to minimise
 309 it. As for NFGs (see Theorem 4), a CSG has a *value* with respect to X if it is determined,
 310 i.e., if the maximum value that player 1 can ensure equals the minimum value that player
 311 2 can ensure when starting from any state of the CSG. Since the CSGs we discuss in this
 312 paper are finite-state and finitely-branching, it follows that they are determined for all of the
 313 objectives that we consider [44].

314 Given a multi-player CSG and objective X , we can divide the players into two coalitions,
 315 $C \subseteq N$ and $N \setminus C$, and then construct a *two-player zero-sum* coalition game, in which each
 316 coalition acts as a single player, with one coalition trying to maximise the value of X and
 317 the other trying to minimise that value.

318 3.2 Nonzero-sum CSGs

319 We define *nonzero-sum* CSGs similarly to NFGs: we assume that there is a distinct and
 320 independent objective X_i for each player i (or coalitions of players). We can then define NE
 321 and CE for CSGs (see Definition 7), as well as the restricted classes of SW and SF equilibria,
 322 similarly to those for NFGs (see Definition 8). Following [34, 32], we focus on *subgame-perfect*
 323 equilibria [47], which are equilibria in *every state* of G . Furthermore, because we include
 324 infinite-horizon objectives, where the existence of NE is an open problem [7], we will in some
 325 cases use ε -NE, which do exist for any $\varepsilon > 0$ for all the infinite-horizon objectives we consider.
 326

327 ► **Definition 11** (Subgame-perfect ε -NE). *For CSG G and $\varepsilon > 0$, a strategy profile σ^* is*
 328 *a subgame-perfect ε -Nash equilibrium for objectives $\langle X_i \rangle_{i \in N}$ if and only if $\mathbb{E}_{G,s}^{\sigma^*}(X_i) \geq$*
 329 *$\sup_{\sigma_i \in \Sigma_i} \mathbb{E}_{G,s}^{\sigma_i, \sigma_i^*}(X_i) - \varepsilon$ for all $i \in N$ and $s \in S$.*

330 ► **Example 12.** As an example scenario that can be modelled as a CSG, consider a number
 331 of users trying to send packets using the slotted ALOHA protocol studied in [32, 34, 35]. If
 332 there is a collision or if sending a packet fails, a user waits for some number of slots before
 333 resending, with the wait set according to an exponential backoff scheme.

334 If we model this scenario as a CSG then, when a player has a packet to send, the actions
 335 available to the player correspond to either sending their packet or waiting to send the packet
 336 at some future time step. In the case when one coalition of players has an objective related
 337 to sending their packets efficiently, e.g., minimising the expected time to send their packets,
 338 and the remaining players form a second coalition and have the dual objective, we can model
 339 this scenario as a zero-sum CSG. In such a zero-sum CSG, the optimal strategy for the first
 340 coalition is to try and choose times to send that avoid collisions, while the second coalition
 341 will do the opposite and instead try and cause collisions. On the other hand, when there are
 342 more coalitions and each coalition's goal corresponds to sending their own packets efficiently,
 343 we can model this as a nonzero-sum CSG. Here we would be looking for equilibria, i.e.,
 344 profiles such that no coalition could improve its objective by changing its strategy, which are
 345 also optimal, e.g., the sum of the expected times is minimal or the difference between the
 346 expected time to send for each coalition is minimal.

347 4 Property specifications and model checking for CSGs

348 Probabilistic model checking is a technique for systematically constructing a stochastic model
 349 and analysing it against a quantitative property formally specified in temporal logic. This
 350 approach can be used either to *verify* that a specification is always satisfied or to perform
 351 *strategy synthesis*, i.e., to construct a witness to the satisfaction of a property. In the context
 352 of CSGs, the latter means synthesising strategies for one or more players (or coalitions) such
 353 that the resulting behaviour of the game satisfies the specification.

354 To specify properties of labelled CSGs, we use the property specification language of the
 355 PRISM-games model checker [33], which is based on the logic PCTL (probabilistic compu-
 356 tation tree logic) [26], extended with operators to specify expected reward properties [21]
 357 and the *coalition* operator $\langle\langle C \rangle\rangle$ from alternating temporal logic (ATL) [4]. The variant
 358 of this logic that just considers *zero-sum* formulae is referred to as rPATL (probabilistic
 359 alternating-time temporal logic with rewards) in [15], but here we use a further extended
 360 version that also supports *nonzero-sum* properties, using the notion of equilibria [34, 35].

361 ► **Definition 13** (PRISM-games logic [34, 35]). *The syntax of the PRISM-games logic is given*
 362 *by the grammar:*

$$\begin{aligned}
 363 \quad \phi &::= \mathbf{true} \mid \mathbf{a} \mid \neg\phi \mid \phi \wedge \phi \mid \langle\langle C \rangle\rangle P_{\sim q}[\psi] \mid \langle\langle C \rangle\rangle R_{\sim x}[\rho] \mid \langle\langle C \rangle\rangle (\star_1, \star_2)_{\text{opt} \sim x}(\theta) \\
 364 \quad \psi &::= \mathbf{X} \phi \mid \phi \mathbf{U}^{\leq k} \phi \mid \phi \mathbf{U} \phi \\
 365 \quad \rho &::= \mathbf{I}^{\leq k} \mid \mathbf{C}^{\leq k} \mid \mathbf{F} \phi \\
 366 \quad \theta &::= \mathbf{P}[\psi] + \dots + \mathbf{P}[\psi] \mid \mathbf{R}^r[\rho] + \dots + \mathbf{R}^r[\rho]
 \end{aligned}$$

367 where \mathbf{a} is an atomic proposition, $\mathbb{C} = C_1 : \dots : C_m$, C and C_1, \dots, C_m are coalitions of players
 368 such that $C' = N \setminus C$, $C_i \cap C_j = \emptyset$ for all $1 \leq i \neq j \leq m$, $(\star_1, \star_2) \in \{\text{NE}, \text{CE}\} \times \{\text{SW}, \text{SF}\}$,
 369 $\text{opt} \in \{\text{min}, \text{max}\}$, $\sim \in \{\text{<}, \text{≤}, \text{≥}, \text{>}\}$, $q \in \mathbb{Q} \cap [0, 1]$, $x \in \mathbb{Q}$, r is a reward structure and $k \in \mathbb{N}$.

370 The syntax distinguishes between state (ϕ), path (ψ) and reward (ρ) formulae. State formulae
 371 are evaluated over states of a CSG, while path and reward formulae are both evaluated over
 372 paths. Sums of formulae (θ) are used to specify multiple objectives for equilibria.

373 We omit the formal semantics, which can be found in [34, 35]. Path and reward formulae
 374 are used to express the utilities of the players, i.e., random variables over paths. For path
 375 formulae, we allow *next* ($\mathbf{X} \phi$), *bounded until* ($\phi \mathbf{U}^{\leq k} \phi$) and *unbounded until* ($\phi \mathbf{U} \phi$). We
 376 also allow the usual equivalences such as $\mathbf{F} \phi \equiv \mathbf{true} \mathbf{U} \phi$ (i.e., *probabilistic reachability*)
 377 and $\mathbf{F}^{\leq k} \phi \equiv \mathbf{true} \mathbf{U}^{\leq k} \phi$ (i.e., *bounded probabilistic reachability*). The random variable

378 corresponding to the path formula ψ returns 1 for paths that satisfy ψ and zero otherwise.
 379 For reward formulae, we allow instantaneous (state) reward at the k th step (*instantaneous*
 380 *reward* \mathbf{I}^k), reward accumulated over k steps (*bounded cumulative reward* $\mathbf{C}^{\leq k}$), and reward
 381 accumulated until a formula ϕ is satisfied (*expected reachability* $\mathbf{F}\phi$). The random variable
 382 corresponding to the reward formula ρ returns for a path the reward corresponding to ρ .

383 4.1 Zero-sum formulae

384 A state satisfies a formula $\langle\langle C \rangle\rangle \mathbf{P}_{\sim q}[\psi]$ if the coalition of players $C \subseteq N$ can ensure that
 385 the probability of the path formula ψ being satisfied is $\sim q$, regardless of the actions of the
 386 other players ($N \setminus C$) in the game. A state satisfies a formula $\langle\langle C \rangle\rangle \mathbf{R}_{\sim x}^r[\rho]$ if the players in
 387 C can ensure that the expected value of the reward formula ρ for reward structure r is $\sim x$,
 388 whatever the other players do.

389 The model checking algorithms presented in [34] involve graph-based analysis followed
 390 by backward induction [53, 57] for exact computation of *finite-horizon properties* and value
 391 iteration [51, 12] for approximate computation of *infinite-horizon properties*. During both
 392 backward induction and value iteration for each state, at each iteration, an LP problem of
 393 size $|A|$ must be solved (corresponding to finding the value of a zero-sum one-shot game),
 394 which has complexity PTIME [30].

395 Strategy synthesis for the formulae $\langle\langle C \rangle\rangle \mathbf{P}_{\sim q}[\psi]$ and $\langle\langle C \rangle\rangle \mathbf{R}_{\sim x}^r[\rho]$ corresponds to finding
 396 optimal strategies for the players in coalition C when their objective, respectively, is max-
 397 imising the probability of satisfying the formula ψ and maximising the expected value of
 398 the reward formula with respect to the reward structure r . All strategies synthesised are
 399 randomised and can be found during model checking by extracting not just the value of the
 400 zero-sum one-shot game solved in each state, but also an optimal (randomised) strategy. For
 401 infinite-horizon objectives, the synthesised strategies are memoryless, while for finite-horizon
 402 objectives, the synthesised strategies are finite-memory, with a separate distribution required
 403 for each state and each time step.

404 4.2 Nonzero-sum formulae

405 Nonzero-sum formulae allow us to reason about equilibria, for either of the types (NE
 406 or CE) and optimality criteria (SW or SF) considered here. A probabilistic formula
 407 $\langle\langle C_1 \cdots C_m \rangle\rangle (\star_1, \star_2)_{\max \sim x} (\mathbf{P}[\psi_1] + \cdots + \mathbf{P}[\psi_m])$ is true in a state if, when the players form
 408 the coalitions C_1, \dots, C_m , there is a subgame-perfect equilibrium of type \star_1 meeting the
 409 optimality criterion \star_2 for which the *sum* of the values of the objectives $\mathbf{P}[\psi_1], \dots, \mathbf{P}[\psi_m]$
 410 for the coalitions C_1, \dots, C_m satisfies $\sim x$. The objective of coalition C_i is to maximise the
 411 probability of satisfying a path formula ψ_i .

412 For a reward formula $\langle\langle C_1 \cdots C_m \rangle\rangle (\star_1, \star_2)_{\max \sim x} (\mathbf{R}^{r_1}[\rho_1] + \cdots + \mathbf{R}^{r_m}[\rho_m])$ the meaning is
 413 similar; however, here the objective of coalition C_i refers to a reward formula ρ_i with respect
 414 to reward structure r_i . Formulae of the form $\langle\langle C_1 \cdots C_m \rangle\rangle (\star_1, \star_2)_{\min \sim x}(\theta)$ correspond to the
 415 dual notion of cost equilibria, which are also supported. We also allow *numerical* queries of the
 416 form $\langle\langle C_1 \cdots C_m \rangle\rangle (\star_1, \star_2)_{\text{opt}=?}(\theta)$, which return the sum of the subgame-perfect equilibrium's
 417 values of of type \star_1 meeting the optimality criterion \star_2 .

418 Model checking algorithms, presented in [34, 32, 35], involve solving an m -player *coalition*
 419 *game* $\mathbf{G}^{\mathcal{C}}$, where $\mathcal{C} = \{C_1, \dots, C_m\}$ and the choices of each player i in $\mathbf{G}^{\mathcal{C}}$ correspond to
 420 the choices of the players in coalition C_i in \mathbf{G} . If all the objectives in θ are finite-horizon,
 421 then *backward induction* [53, 57] can be applied to compute (precise) optimal equilibria
 422 values. On the other hand, if all the objectives are infinite-horizon, *value iteration* [12] can

423 be used to approximate optimal equilibria values. When there is a combination of finite- and
 424 infinite-horizon objectives, the game under study is modified in a standard manner to make
 425 all objectives infinite-horizon.

426 Both backward induction and value iteration over the CSG G^C work by iteratively
 427 computing new values for each state s of G^C . The values for each state, in each iteration,
 428 are found by computing optimal equilibria values, with respect to the criterion \star_2 and
 429 equilibrium type \star_1 , of an NFG N whose utility function is derived from the outgoing
 430 transition probabilities from s in the CSG and the values computed for successor states of s
 431 in the previous iteration.

432 We can synthesise a strategy profile representing the appropriate type of equilibrium for
 433 the CSG by combining the optimal strategies for the equilibria generated in each individual
 434 state during solution. As for zero-sum formulae, randomisation is required and memory
 435 is needed both to keep track of both the step bound of finite-horizon objectives and the
 436 satisfaction of each player's objective.

437 ► **Example 14.** We now return to the scenario from Example 12, where a number users
 438 are attempting to send packets using the slotted ALOHA protocol. The zero-sum formulae
 439 $\langle\langle usr_1, \dots, usr_k \rangle\rangle_{\min=?} \mathbf{R}^{time}[\mathbf{F} \text{ sent}_{1\dots k}]$ and $\langle\langle usr_1, \dots, usr_k \rangle\rangle_{\max=?} \mathbf{P}[\mathbf{F} \text{ sent}_{1\dots k} \wedge t \leq D]$ represent
 440 the case where the first k users form a coalition and try to minimise the expected time
 441 to send their packets or maximise the probability they send their packets within a deadline
 442 $D \in \mathbb{N}$, respectively, while the remaining users form a second coalition and try and achieve
 443 the opposite objective, i.e., maximise the expected time or minimise the probability.

444 On the other hand, in the nonzero-sum case, if we suppose there are m users and the object-
 445 ive of each user is to minimise the expected time to send their packet, this can be expressed by
 446 the nonzero-sum formula $\langle\langle usr_1 : \dots : usr_m \rangle\rangle_{(\star_1, \star_2)_{\min=?}} (\mathbf{R}^{time}[\mathbf{F} \text{ sent}_1] + \dots + \mathbf{R}^{time}[\mathbf{F} \text{ sent}_m])$.

447 5 Tool support and case studies

448 Tool support for the modelling and automated verification of CSGs has been implemented in
 449 PRISM-games [33], which is available from [61]. A variety of case studies have been modelled
 450 and analysed as CSGs with the tool, using both zero-sum and nonzero-sum properties. These
 451 include: a robot coordination problem [34]; futures market investors [34, 35]; medium access
 452 control [32, 34]; power control [34, 35]; a public good game [32, 35] and secret sharing [32].
 453 The results for these case studies demonstrate: the advantages of using CSGs for modelling
 454 (for example, with respect to simpler turn-based games); that using nonzero-sum properties
 455 can yield gains for the players (or coalitions); and that the use of correlated equilibria and
 456 social fairness results may be advantageous compared to Nash equilibria and social welfare.
 457 We give a brief description of the functionality and implementation of PRISM-games and
 458 then present a representative case study: the slotted ALOHA protocol.

459 5.1 PRISM-games

460 PRISM-games [33] is an extension of the PRISM model checker, which provides support
 461 for a variety of stochastic game models, including turn-based and concurrent multi-player
 462 stochastic games, and (turn-based) timed probabilistic games.

463 These are all described in the PRISM-games modelling language, a stochastic extension
 464 of the Reactive Modules formalism [3]. The language facilitates the specification of systems
 465 comprising multiple components, referred to as modules, that operate in parallel, both
 466 asynchronously and synchronously through action labels. Each module has a number of

467 finite-valued variables and a state of the system specifies the values of the variables of all
 468 modules. The behaviour of each module is defined by probabilistic guarded commands,
 469 where the guard is a predicate over the variables of the modules and the command specifies
 470 a probabilistic update of the module’s variables. In a CSG model, each player constitutes a
 471 set of modules, and these therefore execute concurrently.

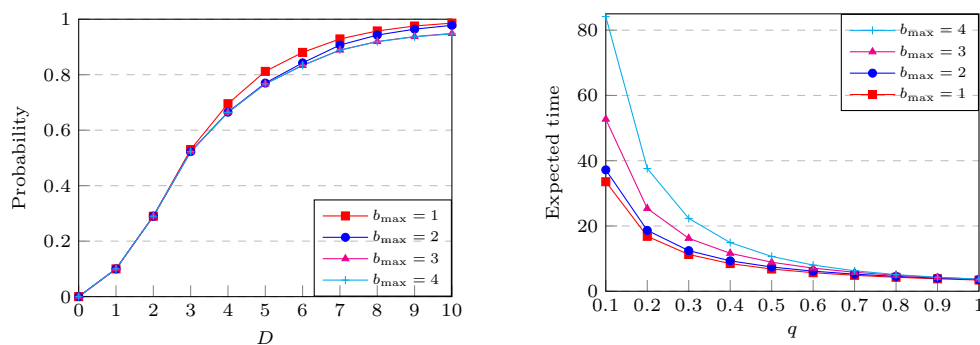
472 PRISM-games provides a graphical user interface for designing and simulating stochastic
 473 games models, but its core functionality is to exhaustively construct a game and perform
 474 verification and strategy synthesis against a logical specification. For CSGs, the PRISM-
 475 games logic described in Definition 13 is supported, and the resulting strategies can be
 476 exported, simulated or further verified.

477 The implementation of CSG model checking is built within PRISM’s ‘explicit’ engine,
 478 which is based on sparse matrices and implemented in Java. Computing values (and optimal
 479 strategies) for zero-sum NFGs, needed for zero-sum formulae, is performed using the LPSolve
 480 library [41] via linear programming. The computation of SWNE or SFNE for nonzero-
 481 sum NFG, required for nonzero-sum formulae, depends on the number of players. For
 482 two players [34], labelled polytopes are used to characterise and find NE values through a
 483 reduction to SMT in both Z3 [19] and Yices [20]. If there are more than two players, the
 484 implementation [32] is based on support enumeration and uses a combination of the SMT
 485 solver Z3 [19] and the nonlinear optimisation suite IPOPT [58]. In the case of SWCE for
 486 nonzero-sum NFGs, as the problem reduces to an LP problem [35], either Gurobi [22] or the
 487 SMT solver Z3 [19] is used. Finally, for SFCE, since the problem does not reduce directly to
 488 an LP problem, only Z3 can be used.

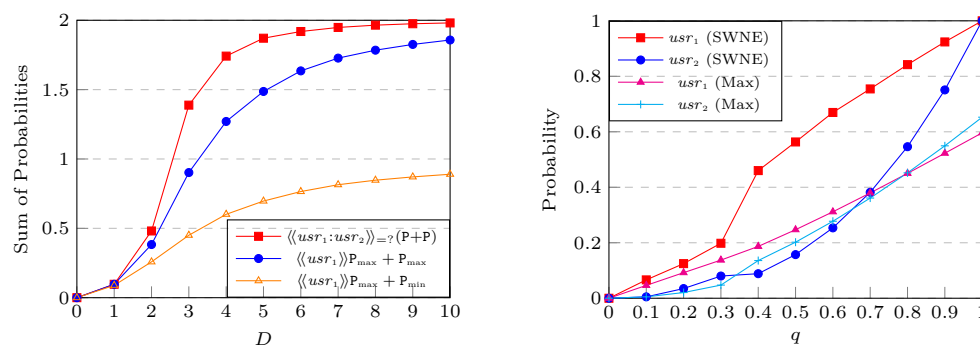
489 5.2 The ALOHA case study

490 We now return to the slotted ALOHA protocol discussed in Examples 12 and 14 to illustrate
 491 the benefits of game-theoretic analysis with CSGs. For further details of this case study, as
 492 well as several others, see [61]. Recall that, in the slotted ALOHA protocol, a number of
 493 users are attempting to send packets on a shared medium. We assume that, in any time
 494 slot, if a single user tries to send a packet then there is a probability (q) that the packet is
 495 sent and, as more users try and send, then the probability of success decreases. If sending a
 496 packet fails, the user waits for a number of slots before resending, defined according to an
 497 exponential backoff scheme. More precisely, each user maintains a backoff counter, which
 498 it increases each time there is a failure (up to b_{\max}) and, if the counter equals k , randomly
 499 chooses the slots to wait from $\{0, 1, \dots, 2^k - 1\}$.

500 **Zero-sum properties.** We first consider the zero-sum properties $\langle\langle usr_1 \rangle\rangle_{P_{\max=?}}[\mathbf{F}^{\leq D} \text{sent}_1]$
 501 and $\langle\langle usr_1 \rangle\rangle_{R_{\min=?}^{time}}[\mathbf{F} \text{sent}_1]$ from Example 14, which correspond to the first user trying to
 502 maximise the probability that their packet is sent before a deadline and trying to minimise
 503 the expected time to send their packet, respectively. The results for the first property when
 504 $q = 0.9$ as the deadline D varies, and for the second property as the probability q varies,
 505 are presented in Figure 2 for different values of b_{\max} . We see that the probability decreases
 506 and the expected time decreases as b_{\max} increases; this is because, as b_{\max} increases, the
 507 additional time the first user can spend in backoff outweighs the gains in reducing the chance
 508 of avoiding further collisions. By performing strategy synthesis we see that it is optimal
 509 for the first user to initially randomly decide as to when to send their packet in order to
 510 avoid collisions with the coalition of the second and third user. However, this changes to a
 511 deterministic strategy of just sending its packet when the other users have sent their packets
 512 or the deadline is getting close, and therefore waiting will mean the deadline is missed.



■ **Figure 2** Results from a CSG model of the ALOHA protocol: one user maximising the probability of sending their packet before a deadline D (left); and minimising the expected time to send the packet, assuming a message transmission failure probability q (right).




■ **Figure 3** Results from CSG equilibria synthesis on the ALOHA protocol, maximising the probabilities of sending packets by deadline D for two coalitions (user 1, and users 2 and 3): probability sums (left) and individual probabilities (right).

513 **Benefits of equilibria.** We next highlight the analysis from [34], which demonstrates the
 514 advantages of cooperation through nonzero-sum properties when using Nash equilibria (NE)
 515 and the social welfare (SW) optimality criterion, as opposed to adopting a strategy that
 516 assumes antagonistic behaviour. The first non-zero sum property we consider corresponds
 517 to the case when each user is trying to maximise the probability of sending their packet
 518 before a deadline D , with users 2 and 3 forming a coalition, represented by the formula
 519 $\langle\langle usr_1 : usr_2, usr_3 \rangle\rangle_{(NE, SW)_{\max=?}} (P[\mathbf{F}(\text{sent}_1 \wedge t \leq D)] + P[\mathbf{F}(\text{sent}_2 \wedge \text{sent}_3 \wedge t \leq D)])$.

520 Figure 3 presents total values (the sum of the probabilities for user 1 and the coalition
 521 of user 2 and 3) as D varies (left) and individual values as q varies (right). By performing
 522 strategy synthesis, the analysis found that the collaboration is dependent on both D and q .
 523 In particular, if the users have more time there is a greater chance for the users to collaborate
 524 by sending in different slots, whereas, when q is large, it is unlikely users need to repeatedly
 525 send, so again can send in different slots. As Figure 3 (right) demonstrates, since the coalition
 526 has more packets to send, their probabilities are lower.

527 **Equilibria types and optimality criteria.** Finally, we report on the experiments of [35], which
 528 investigate the benefits of using different types of equilibria, i.e., *correlated* (CE) over *Nash*
 529 equilibria, and optimality criteria, i.e., *social fairness* (SF) over *social welfare* (SW). The
 530 experiments varied the number of users and considered the case when the objective of each
 531 individual user is to minimise the expected time to send their packet, which is represented by
 532 the nonzero-sum formula $\langle\langle usr_1 : \dots : usr_m \rangle\rangle_{(\star_1, \star_2)_{\min=?}} (R^{time}[\mathbf{F} \text{ sent}_1] + \dots + R^{time}[\mathbf{F} \text{ sent}_m])$.



■ **Figure 4** Results from different types of equilibria (correlated vs. Nash) and optimality criteria (social fairness vs. social welfare) for minimising the expected times for users to send packets in the ALOHA protocol., for varying numbers of users.

533 Synthesising optimal strategies for this specification, it was found that the cases for
 534 SWNE and SWCE coincide (although SWCE returns a joint strategy for the users, this joint
 535 strategy can be separated to form a strategy profile). This profile required one user to try
 536 and send first, and then for the remaining users to take turns to try and send afterwards. If
 537 a user fails to send, then they enter backoff and allow all remaining users to try and send
 538 before trying to send again. The reason for this is that there is no gain in a user trying to
 539 send at the same time as another user, as this will increase the probability of a collision and
 540 thus their packets not being sent, and therefore the users having to spend time in backoff.

541 For SFNE, which has only been implemented for the two-player case, the two users
 542 followed identical strategies, which involve randomly deciding whether to wait or transmit,
 543 unless they are the only user that has not transmitted, and then they always try to send
 544 when not in backoff. In the case of SFCE, users employed a shared probabilistic signal to
 545 coordinate which user sends next. Initially, this was a uniform choice over the users, but as
 546 time progresses the signal favoured the users with lower backoff counters as these users had
 547 fewer opportunities to send their packet previously.

548 Figure 4 plots the optimal values for the users, where SW_{*i*} corresponds to the optimal
 549 values (expected times to send their packets) for user *i* for both SWNE and SWCE for the
 550 cases of two, three and four users. We see that the optimal values for the different users
 551 under SFNE and SFCE coincide, while under SWNE and SWCE they are different for each
 552 user (with the user sending first having the lowest and the user sending last the highest).
 553 Comparing the sum of the SWNE (and SWCE) values and that of the SFCE values, we see
 554 a small decrease in the sum of less than 2% of the total, whereas for SFNE there is a greater
 555 difference as the users cannot coordinate, and hence try and send at the same time.

556 6 Recent Developments: Neuro-symbolic CSGs

557 The recent encouraging advances of AI, and particularly deep learning, have resulted in
 558 computing architectures that integrate components that are synthesized from data (e.g.,
 559 implemented as neural networks) with conventional, symbolic modules (e.g., controllers).
 560 Design automation support for such *neuro-symbolic* systems is, however, lacking. To this
 561 end, we have developed the model of *neuro-symbolic concurrent stochastic games* (NS-
 562 CSGs) [60, 59], which is targeted at AI-based autonomous systems, e.g., autonomous driving
 563 or aircraft controllers. NS-CSGs are a variant of (continuous-space) CSGs, in which each
 564 player is a neuro-symbolic *agent* and the agents act concurrently in a shared, continuous-state
 565 environment. As for the players of CSGs, each agent has a finite set of available actions and
 566 agents choose their actions simultaneously; however, in NS-CSGs the action choices cause

567 the agents' local states to be updated probabilistically and the agents are endowed with a
 568 perception mechanism implemented as a neural network, through which they can observe the
 569 local states of the other agents and that of the environment and encode these observations as
 570 locally stored *percepts*. The global states of NS-CSGs comprise the state of the environment
 571 together with the local state and percept of each agent, and are therefore infinite-state, in
 572 contrast to the CSGs discussed in the rest of this paper.

573 ► **Definition 15** (Neuro-symbolic concurrent stochastic game [60, 59]). *A neuro-symbolic*
 574 *concurrent stochastic multi-player game (NS-CSG) NSC comprises players $(\mathbf{Ag}_i)_{i \in N}$, for*
 575 *$N = \{1, \dots, n\}$, and an environment E where:*

$$576 \quad \mathbf{Ag}_i = (S_i, A_i, \Delta_i, obs_i, \delta_i) \text{ for } i \in N, \quad E = (S_E, \delta_E)$$


577 *and we have:*

- 578 ■ $S_i = Loc_i \times Per_i$ is a set of states for \mathbf{Ag}_i , where $Loc_i \subseteq \mathbb{R}^{b_i}$ and $Per_i \subseteq \mathbb{R}^{d_i}$ are finite
 579 sets of local states and percepts, respectively;
- 580 ■ $S_E \subseteq \mathbb{R}^e$ is a closed infinite set of environment states;
- 581 ■ A_i is a nonempty finite set of actions for \mathbf{Ag}_i and $A := (A_1 \cup \{\perp\}) \times \dots \times (A_n \cup \{\perp\})$ is
 582 the set of joint actions, where \perp is an idle action disjoint from $\cup_{i=1}^n A_i$;
- 583 ■ $\Delta_i : S_i \rightarrow 2^{A_i}$ is an available action function, defining the actions \mathbf{Ag}_i can take in each
 584 state;
- 585 ■ $obs_i : (Loc_1 \times \dots \times Loc_n \times S_E) \rightarrow Per_i$ is a perception function for \mathbf{Ag}_i , mapping the
 586 local states of all agents and the environment to a percept of the agent, implemented via
 587 a neural network (NN) classifier;
- 588 ■ $\delta_i : S_i \times A \rightarrow \mathbb{P}(Loc_i)$ is a partial probabilistic transition function for \mathbf{Ag}_i determining
 589 the distribution over the agent's local states given its current state and joint action;
- 590 ■ $\delta_E : S_E \times A \rightarrow S_E$ is a partial deterministic environment transition function determining
 591 the environment's next state given its current state and joint action.

592 A (global) state for an NS-CSG NSC comprises a state $s_i = (loc_i, per_i)$ for each agent
 593 \mathbf{Ag}_i (a pair of a local state and percept) and an environment state s_E . If an NS-CSG is
 594 in a state $s = (s_1, \dots, s_n, s_E)$, then each \mathbf{Ag}_i simultaneously chooses one of the actions
 595 available in its state s_i (if no action is available, i.e., $\Delta_i(s_i) = \emptyset$, it picks the idle action \perp)
 596 yielding a joint action $\alpha = (a_1, \dots, a_n) \in A$. Next, each \mathbf{Ag}_i updates its local state to
 597 $loc'_i \in Loc_i$, according to its probabilistic local transition function δ_i , applied to its current
 598 state (loc_i, per_i) and the joint action α . The environment updates its state to $s'_E \in S_E$
 599 according to its local deterministic transition function δ_E , based on its state s_E and on α .
 600 Finally, each agent, based on its new local state loc'_i , observes the new local states of the
 601 other agents and environment through its perception function obs_i to generate a new percept
 602 $per'_i = obs_i(loc'_1, \dots, loc'_n, s'_E)$. Thus, the game reaches the state $s' = (s'_1, \dots, s'_n, s'_E)$, where
 603 $s'_i = (loc'_i, per'_i)$ for $i \in N$. We assume that each perception function obs_i is implemented
 604 via an NN $f_i : \mathbb{R}^{b+e} \rightarrow \mathbb{P}(Per_i)$ yielding a normalised score over different percept values,
 605 where $b = \sum_{i=1}^n b_i$; however, it can be any function including other types of machine learning
 606 models. A rule is then applied that selects the percept value with the maximum score.

607 Formally, the semantics of an NS-CSG NSC is given by an infinite-state CSG $\llbracket \text{NSC} \rrbracket$ over
 608 the product of the states of the agents and environment, which assumes a particular structure
 609 of the transition function that distinguishes between agent and environment states and uses
 610 the NN-based perception function to define which states have the same characteristics.

611 ► **Definition 16** (Semantics of an NS-CSG). *Given an NS-CSG NSC consisting of n players*
 612 *and an environment, the semantics of NSC is $\llbracket \text{NSC} \rrbracket = (N, S, (A_i)_{i \in N}, \Delta, \delta)$ where:*



■ **Figure 5** Geometry with trust levels and advisories for the agents of the VCAS[2] case study.

- 613 ■ $S = S_1 \times \dots \times S_n \times S_E$ is the set of (global) states, which contain both discrete and
 614 continuous elements;
 615 ■ $\Delta(s_1, \dots, s_n, s_E) = \cup_{i=1}^n \Delta_i(s_i)$;
 616 ■ $\delta : (S \times ((A_1 \cup \{\perp\}) \times \dots \times (A_n \cup \{\perp\}))) \rightarrow \mathbb{P}(S)$ is the partial probabilistic transition
 617 function, where for states $s = (s_1, \dots, s_n, s_E)$, $s' = (s'_1, \dots, s'_n, s'_E) \in S$ and joint action
 618 $\alpha = (a_1, \dots, a_n) \in A$, if $a_i \in \Delta_i(s_i)$ when $\Delta_i(s_i) \neq \emptyset$ and $a_i = \perp$ otherwise, then
 619 $\delta(s, \alpha)$ is defined and if $s'_i = (loc'_i, per'_i)$, $per'_i = obs_i(loc'_1, \dots, loc'_n, s'_E)$ for all $i \in N$ and
 620 $s'_E = \delta_E(s_E, \alpha)$, then

$$621 \quad \delta(s, \alpha)(s') = (\prod_{i=1}^n \delta_i(s_i, \alpha)(loc'_i))$$

622 and otherwise $\delta(s, \alpha)(s') = 0$.

623 To illustrate NS-CSGs, we present the VerticalCAS Collision Avoidance Scenario (VCAS[2])
 624 [28, 29], modelled in [59], which is a variant of the one studied in [2], where the agents' trust
 625 level is modelled probabilistically to account for possible uncertainty.

626 ► **Example 17.** The geometry of the VCAS[2] case study is shown in Figure 5. There are
 627 two aircraft (ownship and intruder), constituting the agents of the NS-CSG, both equipped
 628 with an NN-controlled collision avoidance system called VCAS. At each time unit (i.e., every
 629 second), VCAS issues an advisory ($ad \in \{1, \dots, 9\}$) from which, together with the current
 630 trust level ($tr \in \{1, \dots, 4\}$) from the previous advisory, the pilot needs to make a decision
 631 about the rate of acceleration, aimed at avoiding a near mid-air collision (NMAC) [1].

632 The input to the VCAS system is the tuple $(h, \dot{h}_{own}, \dot{h}_{int}, t) \in \mathbb{R}^4$ including the relative
 633 altitude h of the aircraft, the climb rate \dot{h}_{own} of ownship, the climb rate \dot{h}_{int} of intruder, and
 634 the time t until loss of horizontal separation between the aircraft. VCAS is implemented
 635 via nine feed-forward NNs $F = \{f_i : \mathbb{R}^4 \rightarrow \mathbb{R}^9 \mid 1 \leq i \leq 9\}$, each of which corresponds to an
 636 advisory and outputs the scores of the nine possible advisories. Each advisory will provide a
 637 set of accelerations for the agent to select from and the trust level increases probabilistically
 638 if the current advisory is compliant with the executed accelerations, and decreases otherwise.
 639 We formulate the NS-CSG with agents \mathbf{Ag}_i for $i \in \{own, int\}$ as follows:

- 640 ■ the set of states for \mathbf{Ag}_i is given by $S_i = \{1, \dots, 4\} \times \{1, \dots, 9\}$, where the agent state
 641 $s_i = (tr_i, ad_i) \in S_i$ has local state (trust level) tr_i and percept (advisory) ad_i ;
 642 ■ the set of environment states is given by $S_E = \mathbb{R}^4$, where $s_E = (h, \dot{h}_{own}, \dot{h}_{int}, t) \in S_E$
 643 represents the relative altitude, the climb rate of the ownship, the climb rate of the
 644 intruder, and the time until loss of their horizontal separation;
 645 ■ the set of actions of \mathbf{Ag}_i is given by $A_i = \{0, \pm 3.0, \pm 7.33, \pm 9.33, \pm 9.7, \pm 11.7\}$ representing
 646 the acceleration options of \mathbf{Ag}_i ;
 647 ■ the available action function $\Delta_i : Per_i \rightarrow A_i$ of \mathbf{Ag}_i is independent of the local state of
 648 the agent and returns the set consisting two non-zero acceleration actions from Table 2
 649 for a given percept and the zero acceleration action;

Label (ad_i)	Advisory	Description	Vertical Range (Min, Max) ft/min	Actions ft/s ²
1	COC	Clear of Conflict	$(-\infty, +\infty)$	-3, +3
2	DNC	Do Not Climb	$(-\infty, 0]$	-9.33, -7.33
3	DND	Do Not Descend	$[0, +\infty)$	7.33, +9.33
4	DES1500	Descend at least 1500 ft/min	$(-\infty, -1500]$	-9.33, -7.33
5	CL1500	Climb at least 1500 ft/min	$[+1500, +\infty)$	+7.33, +9.33
6	SDES1500	Strengthen Descend to at least 1500 ft/min	$(-\infty, -1500]$	-11.7, -9.7
7	SCL1500	Strengthen Climb to at least 1500 ft/min	$[+1500, +\infty)$	+9.7, +11.7
8	SDES2500	Strengthen Descend to at least 2500 ft/min	$(-\infty, -2500]$	-11.7, -9.7
9	SCL2500	Strengthen Climb to at least 2500 ft/min	$[+2500, +\infty)$	+9.7, +11.7

■ **Table 2** Actions available for the agents of VCAS[2] for each advisory [2].

- 650 ■ the perception function $obs_i : Per_i \times S_E \rightarrow \{1, \dots, 9\}$ of Ag_i is independent of the local
651 state of the agent and is given by the feed forward NNs F of VCAS;
652 ■ the local transition function δ_i of Ag_i updates the agent's trust level probabilistically
653 according to its current trust level, its current advisory and its executed action;
654 ■ the environment transition function δ_E is given by $\delta_E((h, \dot{h}_{own}, \dot{h}_{int}, t), (\ddot{h}_{own}, \ddot{h}_{int})) =$
655 $(h', \dot{h}'_{own}, \dot{h}'_{int}, t')$ where for time step $\Delta t = 1$:
656 ■ $h' = h - \Delta t(\dot{h}_{own} - \dot{h}_{int}) - 0.5\Delta t^2(\ddot{h}_{own} - \ddot{h}_{int})$;
657 ■ $\dot{h}'_{own} = \dot{h}_{own} + \ddot{h}_{own}\Delta t$;
658 ■ $\dot{h}'_{int} = \dot{h}_{int} + \ddot{h}_{int}\Delta t$;
659 ■ $t' = t - \Delta t$.


6.1 Zero-sum NS-CSGs

660 In view of the uncountable state spaces, [60] presents an approach for zero-sum *discounted*
661 *infinite-horizon* cumulative rewards under the assumption of full state observability for
662 NS-CSGs, which exploits Borel state space decomposition and identifies model restrictions to
663 ensure determinacy, and therefore existence of a value that corresponds to a unique fixed point.
664 Value iteration and policy iteration algorithms to compute values and synthesise optimal
665 strategies are also derived based on formulating piecewise linear or constant representations
666 of the value functions and strategies for NS-CSGs.

6.2 Nonzero-sum NS-CSGs

668 In the case of nonzero-sum NS-CSGs, [59] studies the *undiscounted, finite-horizon* equilibria
669 synthesis problem. The use of finite-horizon objectives simplifies the analysis (note that
670 the existence of infinite-horizon NE for CSGs is an open problem [7], and the verification
671 of non-probabilistic infinite-horizon reachability properties for neuro-symbolic games is
672 undecidable [2]). Both NE and CE using the SW optimality criteria are considered. The
673 algorithms, based on backward induction and non-linear programming, compute *globally*
674 *optimal* equilibria which, from a fixed initial state, are optimal over the chosen time horizon,
675 in contrast to the local optimality of equilibria for finite-state CSGs [34, 35].

677 ► **Example 18.** The NS-CSG model of the VCAS[2] system described in Example 17 is
678 studied in [59], comparing equilibria strategies to the zero-sum strategies analysed in [2].
679 Figure 6 plots the relative altitude h of the two aircraft for equilibria and zero-sum strategies
680 when maximising this value for a given time instant k , plotted for several different initial



■ **Figure 6** Relative altitude h of the two aircraft at time instants k for equilibria and zero-sum strategies for the VCAS[2] case study.

681 values of h . It can be seen that, with respect to the safety criterion established in [28, 2], i.e.,
 682 avoiding an NMAC when two aircraft are separated by less than 100 ft vertically (dotted
 683 line) and 500 ft horizontally, equilibria strategies allow the two aircraft to reach a safe
 684 configuration within a shorter horizon, which would be missed through a zero-sum analysis.

685 The analysis of [59] also considers a reward structure that incorporates both the trust
 686 level and fuel consumption. Figure 7 shows the resulting equilibria strategies when both
 687 safety and trust are prioritised, using two different time horizons. When $t = 3$ initially
 688 (left), it is optimal to follow the advisories and the trust levels tr_{own} and tr_{int} of the two
 689 aircraft never decrease from their initial values of 4. However, when $t = 3$ initially (right),
 690 the optimal strategy shows a deviation from the advisory, denoted by action $a_{own} = 0$, in
 691 state s_2 , resulting in tr_{own} dropping to 3 in s_3 with probability 0.9.


692 The experiments from [59], summarised above, and from [60] are developed using prototype
 693 tools that build upon parts of PRISM-games, but full tool support for NS-CSGs is not yet
 694 provided. Efforts in this direction are continuing.

695 7 Conclusions and Future Challenges

696 This paper has provided an overview of modelling, verification and strategy synthesis
 697 techniques that have been developed and implemented for concurrent stochastic games in
 698 the PRISM-games model checker. Through case studies, we have demonstrated the uses
 699 and advantages of zero-sum and equilibria-based reasoning in strategic decision making,
 700 highlighting Nash and correlated equilibria in conjunction with two optimality criteria, social
 701 welfare and social fairness. We have also discussed recent trends in autonomous systems
 702 towards neural-symbolic architectures, and summarised the first steps towards developing a
 703 modelling framework to support the development of such AI-based systems. Despite some
 704 progress, many problems remain open in this area, in particular, the development of (efficient)
 705 approximate algorithms for (undiscounted infinite-horizon) temporal logic specifications
 706 where the underlying problem is undecidable even in the finite-state case [42].

707 — References —

- 708 1 M. Akintunde, E. Botoeva, P. Kouvaros, and A. Lomuscio. Formal verification of neural agents
 709 in non-deterministic environments. In *Proc. AAMAS'20*, pages 25–33. Springer, 2020.



■ **Figure 7** Optimal strategies for the VCAS[2] case study over two different time horizons, using initial t values of 3 (left) and 4 (right).

- 710 2 M. Akintunde, E. Botoeva, P. Kouvaros, and A. Lomuscio. Verifying Strategic Abilities of
 711 Neural-symbolic Multi-agent Systems. In *Proc. KR'20*, pages 22–32. IJCAI Organization, 9
 712 2020.
- 713 3 R. Alur and T. Henzinger. Reactive modules. *Formal Methods in System Design*, 15(1):7–48,
 714 1999.
- 715 4 R. Alur, T. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the*
 716 *ACM*, 49(5):672–713, 2002.
- 717 5 R. Alur, T. Henzinger, F. Mang, S. Qadeer, S. Rajamani, and S. Tasiran. MOCHA: Modularity
 718 in model checking. In *Proc. 10th Int. Conf. Computer Aided Verification (CAV'98)*, volume
 719 1427 of *LNCS*, pages 521–525, Vancouver, 1998. Springer.
- 720 6 R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical*
 721 *Economics*, 1(1):67–96, 1974.
- 722 7 P. Bouyer, N. Markey, and D. Stan. Mixed Nash equilibria in concurrent games. In *Proc.*
 723 *FSTTCS'14*, volume 29 of *LIPICs*, pages 351–363, 2014.
- 724 8 R. Brenguier. PRALINE: A tool for computing Nash equilibria in concurrent games. In *Proc.*
 725 *CAV'13*, volume 8044 of *LNCS*, pages 890–895. Springer, 2013. lsv.fr/Software/praline/.
- 726 9 T. Brihaye, V. Bruyère, A. Goeminne, J.-F. Raskin, and M. van den Bogaard. The complexity
 727 of subgame perfect equilibria in quantitative reachability games. In *Proc. CONCUR'19*, volume
 728 140 of *LIPICs*, pages 13:1–13:16. Leibniz-Zentrum für Informatik, 2019.
- 729 10 P. Čermák, A. Lomuscio, F. Mogavero, and A. Murano. MCMAS-SLK: A model checker for
 730 the verification of strategy logic specifications. In *Proc. CAV'14*, volume 8559 of *LNCS*, pages
 731 525–532. Springer, 2014.
- 732 11 K. Chatterjee, L. de Alfaro, and T. Henzinger. Strategy improvement for concurrent reachability
 733 and turn-based stochastic safety games. *Journal of Computer and System Sciences*, 79(5):640–
 734 657, 2013.

- 735 12 K. Chatterjee and T. Henzinger. Value iteration. In *25 Years of Model Checking*, volume 5000
736 of *LNCS*, pages 107–138. Springer, 2008.
- 737 13 K. Chatterjee, T. Henzinger, B. Jobstmann, and A. Radhakrishna. GIST: A solver for
738 probabilistic games. In *Proc. CAV'10*, volume 6174 of *LNCS*, pages 665–669. Springer, 2010.
739 pub.ist.ac.at/gist/.
- 740 14 K. Chatterjee, R. Majumdar, and M. Jurdziński. On Nash equilibria in stochastic games. In
741 *Proc. CSL'04*, volume 3210 of *LNCS*, pages 26–40. Springer, 2004.
- 742 15 T. Chen, V. Forejt, M. Kwiatkowska, D. Parker, and A. Simaitis. Automatic verification of
743 competitive stochastic systems. *Formal Methods in System Design*, 43(1):61–92, 2013.
- 744 16 C. Cheng, A. Knoll, M. Luttenberger, and C. Buckl. GAVS+: An open platform for the
745 research of algorithmic game solving. In *Proc. TACAS'11*, volume 6605 of *LNCS*, pages
746 258–261. Springer, 2011. sourceforge.net/projects/gavsplus/.
- 747 17 L. de Alfaro, T. Henzinger, and O. Kupferman. Concurrent reachability games. *Theoretical
748 Computer Science*, 386(3):188–217, 2007.
- 749 18 L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *Journal of
750 Computer and System Sciences*, 68(2):374–397, 2004.
- 751 19 L. De Moura and N. Bjørner. Z3: An efficient SMT solver. In *Proc. TACAS'08*, volume 4963
752 of *LNCS*, pages 337–340. Springer, 2008. github.com/Z3Prover/z3.
- 753 20 B. Dutertre. Yices 2.2. In *Proc. CAV'14*, volume 8559 of *LNCS*, pages 737–744. Springer,
754 2014. yices.csl.sri.com.
- 755 21 V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker. Automated verification techniques for
756 probabilistic systems. In *SFM'11*, volume 6659 of *LNCS*, pages 53–113. Springer, 2011.
- 757 22 Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2021. gurobi.com.
- 758 23 J. Gutierrez, M. Najib, P. Giuseppe, and M. Wooldridge. Equilibrium design for concurrent
759 games. In *Proc. CONCUR'19*, volume 140 of *LIPICs*, pages 22:1–22:16. Leibniz-Zentrum für
760 Informatik, 2019.
- 761 24 J. Gutierrez, M. Najib, G. Perelli, and M. Wooldridge. EVE: A tool for temporal equilib-
762 rium analysis. In *Proc. ATVA'18*, volume 11138 of *LNCS*, pages 551–557. Springer, 2018.
763 github.com/eve-mas/eve-parity.
- 764 25 Julian Gutierrez, Lewis Hammond, Anthony W. Lin, Muhammad Najib, and Michael J.
765 Wooldridge. Rational verification for probabilistic systems. In *Proc. 18th International
766 Conference on Principles of Knowledge Representation and Reasoning (KR'21)*, pages 312–322,
767 2021.
- 768 26 H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *FAC*, 6(5):512–535,
769 1994.
- 770 27 L. Hurwicz and S. Reiter. *Designing Economic Mechanisms*. Cambridge University Press,
771 2006.
- 772 28 K. Julian and M. Kochenderfer. A reachability method for verifying dynamical systems with
773 deep neural network controllers. *CoRR*, abs/1903.00520, 2019.
- 774 29 K. Julian, S. Sharma, J.-B. Jeannin, and M. Kochenderfer. Verifying aircraft collision avoidance
775 neural networks through linear approximations of safe regions. *CoRR*, abs/1903.00762, 2019.
- 776 30 N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*,
777 4(4):373–395, 1984.
- 778 31 J. Kemeny, J. Snell, and A. Knapp. *Denumerable Markov Chains*. Springer, 1976.
- 779 32 M. Kwiatkowska, G. Norman, D. Parker, and G. Santos. Multi-player equilibria verification
780 for concurrent stochastic games. In *Proc. QEST'20*, volume 12289 of *LNCS*, pages 74–95.
781 Springer, 2020.
- 782 33 M. Kwiatkowska, G. Norman, D. Parker, and G. Santos. PRISM-games 3.0: Stochastic game
783 verification with concurrency, equilibria and time. In *Proc. CAV'20*, volume 12225 of *LNCS*,
784 pages 475–487. Springer, 2020.
- 785 34 M. Kwiatkowska, G. Norman, D. Parker, and G. Santos. Automatic verification of concurrent
786 stochastic systems. *Formal Methods in System Design*, pages 1–63, 2021.

- 787 35 M. Kwiatkowska, G. Norman, D. Parker, and G. Santos. Correlated equilibria and fairness
788 in concurrent stochastic games. In *Proc. TACAS'22*, volume 13244 of *LNCS*, pages 60–78.
789 Springer, 2022.
- 790 36 S. LaValle. Robot motion planning: A game-theoretic foundation. *Algorithmica*, 26:430–465,
791 2000.
- 792 37 C. Lemke and Jr J. Howson. Equilibrium points of bimatrix games. *Journal of the Society for*
793 *Industrial and Applied Mathematics*, 12(2):413–423, 1964.
- 794 38 M. Littman, N. Ravi, A. Talwar, and M. Zinkevich. An efficient optimal-equilibrium algorithm
795 for two-player game trees. In *Proc. UAI'06*, pages 298–305. AUAI Press, 2006.
- 796 39 A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: A model checker for the verification of
797 multi-agent systems. In *Proc. 21st Int. Conf. Computer Aided Verification (CAV'09)*, volume
798 5643 of *LNCS*, pages 682–688. Springer, 2009.
- 799 40 D. Lozovanu and S. Pickl. Determining Nash equilibria for stochastic positional games with
800 discounted payoffs. In *Proc. ADT'17*, volume 10576 of *LNAI*, pages 339–343. Springer, 2017.
- 801 41 LPSolve (version 5.5). lpsolve.sourceforge.net/5.5/.
- 802 42 O. Madani, S. Hanks, and A. Condon. On the undecidability of probabilistic planning and
803 related stochastic optimization problems. *Artificial Intelligence*, 147(1–2):5–34, 2003.
- 804 43 J. Marden and J. Shamma. Game theory and control. *Annual Review of Control, Robotics,*
805 *and Autonomous Systems*, 1(1):105–134, 2018.
- 806 44 D. Martin. The determinacy of Blackwell games. *J. Symbolic Logic*, 63(4):1565–1581, 1998.
- 807 45 R. McKelvey, A. McLennan, and T. Turocy. Gambit: Software tools for game theory. [gambit-](http://gambit-project.org)
808 [project.org](http://gambit-project.org).
- 809 46 N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge
810 University Press, 2007.
- 811 47 M. Osborne and A. Rubinstein. *An Introduction to Game Theory*. Oxford University Press,
812 2004.
- 813 48 R. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding a Nash equilibrium.
814 In *Proc. AAAI'04*, pages 664–669. AAAI Press, 2004.
- 815 49 H. Prasad, L. Prashanth, and S. Bhatnagar. Two-timescale algorithms for learning Nash
816 equilibria in general-sum stochastic games. In *Proc. AAMAS'15*, pages 1371–1379. IFAAMAS,
817 2015.
- 818 50 E. Prisner. *Game Theory Through Examples*. Mathematical Association of America, 1 edition,
819 2014.
- 820 51 T. Raghavan and J. Filar. Algorithms for stochastic games — a survey. *Zeitschrift für*
821 *Operations Research*, 35(6):437–472, Nov 1991.
- 822 52 T. Sandholm, A. Gilpin, and V. Conitzer. Mixed-integer programming methods for finding
823 Nash equilibria. In *Proc. AAAI'05*, pages 495–501. AAAI Press, 2005.
- 824 53 U. Schwalbe and P. Walker. Zermelo and the early history of game theory. *Games and*
825 *Economic Behavior*, 34(1):123–137, 2001.
- 826 54 L. Shapley. Stochastic games. *PNAS*, 39:1095–1100, 1953.
- 827 55 A. Toumi, J. Gutierrez, and M. Wooldridge. A tool for the automated verification of Nash
828 equilibria in concurrent games. In *Proc. ICTAC'15*, volume 9399 of *LNCS*, pages 583–594.
829 Springer, 2015.
- 830 56 J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320,
831 1928.
- 832 57 J. von Neumann, O. Morgenstern, H. Kuhn, and A. Rubinstein. *Theory of Games and*
833 *Economic Behavior*. Princeton University Press, 1944.
- 834 58 A. Wächter. Short tutorial: Getting started with IPOPT in 90 minutes. In *Combinatorial*
835 *Scientific Computing*, number 09061 in Dagstuhl Seminar Proceedings. Leibniz-Zentrum für
836 Informatik, 2009. github.com/coin-or/Ipopt.
- 837 59 R. Yan, G. Santos, X. Duan, D. Parker, and M. Kwiatkowska. Finite-horizon equilibria for
838 neuro-symbolic concurrent stochastic games. In *Proc. UAI'22*, 2022.

4:22 Probabilistic Model Checking for Strategic Equilibria-based Decision Making

- 839 **60** R. Yan, G. Santos, G. Norman, D. Parker, and M. Kwiatkowska. Strategy synthesis for
840 zero-sum neuro-symbolic concurrent stochastic games. arXiv.2202.06255, 2022.
- 841 **61** PRISM-games web site. prismmodelchecker.org/games/.