

Compositional Controller Synthesis for Stochastic Games

Nicolas Basset, Marta Kwiatkowska, and Clemens Wiltsche

Department of Computer Science, University of Oxford, United Kingdom

Abstract. Design of autonomous systems is facilitated by automatic synthesis of correct-by-construction controllers from formal models and specifications. We focus on stochastic games, which can model the interaction with an adverse environment, as well as probabilistic behaviour arising from uncertainties. We propose a synchronising parallel composition for stochastic games that enables a compositional approach to controller synthesis. We leverage rules for compositional assume-guarantee verification of probabilistic automata to synthesise controllers for games with multi-objective quantitative winning conditions. By composing winning strategies synthesised for the individual components, we can thus obtain a winning strategy for the composed game, achieving better scalability and efficiency at a cost of restricting the class of controllers.

1 Introduction

With increasing pervasiveness of technology in civilian and industrial applications, it has become paramount to provide formal guarantees of safety and reliability for autonomous systems. We consider the development of correct-by-construction controllers satisfying high-level specifications, based on formal system models. Automated synthesis of controllers has been advocated, for example, for autonomous driving [3] and distributed control systems [15].

Stochastic Games. When designing autonomous systems, often a critical element is the presence of an uncertain and adverse environment, which introduces stochasticity and requires the modelling of the non-cooperative aspect in a game-theoretical setting [7,14]. Hence, we model a system we wish to control as a two-player turn-based stochastic game [18], and consider automated synthesis of strategies that are winning against every environment (Player \square), which we can then interpret as controllers of the system (Player \diamond). In addition to probabilities, one can also annotate the model with rewards to evaluate various quantities, for example, profit or energy usage, by means of expectations.

Compositionality. We model systems as a composition of several smaller components. For controller synthesis for games, a compositional approach requires that we can derive a strategy for the composed system by synthesising only for the individual components. Probabilistic automata (PAs) are naturally suited to modelling multi-component probabilistic systems, where synchronising composition is well-studied [17]; see also [19] for a taxonomic discussion. While PAs

can be viewed as stochastic games with strategies already applied, it is not immediately clear how to compose games in a natural way.

Our Composition. We formulate a composition of stochastic games where component games synchronise on shared actions and interleave otherwise. The composition is inspired by interface automata [9], which are well-suited to compositional assume-guarantee verification of component-based systems, where we preserve the identity of Player \diamond and Player \square in the composition by imposing similar compatibility conditions. Our composition is commutative and associative, and reduces to PA composition when only Player \square is present. The results we prove are independent of specific winning conditions, and thus provide a general framework for the development of compositional synthesis methods.

Compositional Synthesis. We show that any rule for compositional verification of PAs carries over as a synthesis rule for games. First, after applying winning strategies to a game, the resulting PAs still satisfy the winning condition for any environment. Then, compositional rules for PAs can be used, such as the assume-guarantee rules in [13] developed for winning conditions involving multi-objective total expected reward and probabilistic LTL queries. These first two steps, described also as “schedule-and-compose” [8], are applicable when strategies can only be implemented locally in practice, for instance, when wanting to control a set of physically separated electrical generators and loads in a microgrid, where no centralised strategy can be implemented. One key property of our game composition is that strategies applied to individual components can be composed to a strategy for the composed game, while preserving the probability measure over the traces. Hence, we obtain a winning strategy for the composed game, which alleviates the difficulty of having to deal with the product state space by trading off the expressiveness of the generated strategies.

Winning Conditions. Each player plays according to a *winning condition*, specifying the desirable behaviour of the game, for example “the probability of reaching a failure state is less than 0.01.” We are interested in synthesis for zero-sum games for which several kinds of winning conditions are defined in the literature, including ω -regular [4], expected total and average reward [10], and multi-objective versions thereof [7]. Our game composition is independent of such winning conditions, since they are definable on the trace distributions.

Work we Build Upon. In this paper we extend the work of [13] by lifting the compositional verification rules for PAs to compositional synthesis rules for games. Typically, such rules involve multi-objective queries, and we extend the synthesis methods for such queries in [6,7] to compositional strategy synthesis.

Contributions. Several notions of (non-stochastic) game composition have recently been proposed [11,12], but they do not preserve player identity, i.e. which player controls which actions, and hence are not applicable to synthesising strategies for a specific player. In this paper, we make the following contributions.

- We define a composition for stochastic games, which, to our knowledge, is the first composition for competitive probabilistic systems that preserves the control authority of Player \diamond , enabling compositional strategy synthesis.

- We show how to apply strategies synthesised for the individual components to the composition, such that the trace distribution is preserved.
- We lift compositional rules for PAs to the game framework for synthesis.
- We apply our theory to demonstrate how to compositionally synthesise controllers for games with respect to multi-objective total expected reward queries, and demonstrate the benefits on a prototype implementation.

Structure. In Section 2 we introduce stochastic games, their normal form, and their behaviour under strategies. In Section 3 we define our game composition, and show that strategies for the individual components can be applied to the composed game. We demonstrate in Section 4 how to use proof rules for PAs and previously developed synthesis methods to compositionally synthesise strategies.

2 Stochastic games, induced PAs and DTMCs

We introduce notation and main definitions for stochastic games and their behaviour under strategies.

Distributions. A *discrete probability distribution* (or *distribution*) over a (countable) set Q is a function $\mu : Q \rightarrow [0, 1]$ such that $\sum_{q \in Q} \mu(q) = 1$; its *support* $\{q \in Q \mid \mu(q) > 0\}$ is the set of values where μ is nonzero. We denote by $\mathcal{D}(Q)$ the set of all distributions over Q with finite support. A distribution $\mu \in \mathcal{D}(Q)$ is *Dirac* if $\mu(q) = 1$ for some $q \in Q$, and if the context is clear we just write q to denote such a distribution μ . We denote by $\boldsymbol{\mu}$ the *product distribution* of $\mu^i \in \mathcal{D}(Q^i)$ for $1 \leq i \leq n$, defined on $Q^1 \times \dots \times Q^n$ by $\boldsymbol{\mu}(q^1, \dots, q^n) \stackrel{\text{def}}{=} \mu^1(q^1) \cdot \dots \cdot \mu^n(q^n)$.

Stochastic Games. We consider turn-based action-labelled stochastic two-player games (henceforth simply called *games*), which distinguish two types of nondeterminism, each controlled by a separate player. Player \diamond represents the controllable part for which we want to synthesise a strategy, while Player \square represents the uncontrollable environment. Examples of games are shown in Figure 1.

Definition 1. A game is a tuple $\langle S, (S_\diamond, S_\square), \varsigma, \mathcal{A}, \longrightarrow \rangle$, where S is a countable set of states partitioned into Player \diamond states S_\diamond and Player \square states S_\square ; $\varsigma \in \mathcal{D}(S)$ is an initial distribution; \mathcal{A} is a countable set of actions; and $\longrightarrow \subseteq S \times (\mathcal{A} \cup \{\tau\}) \times \mathcal{D}(S)$ is a transition relation, such that, for all s , $\{(s, a, \mu) \in \longrightarrow\}$ is finite.

We adopt the infix notation by writing $s \xrightarrow{a} \mu$ for a *transition* $(s, a, \mu) \in \longrightarrow$, and if $a = \tau$ we speak of a τ -transition. The action labels \mathcal{A} on transitions model observable behaviours, whereas τ can be seen as internal: it cannot be used in winning conditions and is not synchronised in the composition.

We denote the set of *moves* by $S_\circ \stackrel{\text{def}}{=} \{(a, \mu) \mid \exists s \in S. s \xrightarrow{a} \mu\}$. A move (a, μ) is *incoming* to a state s if $\mu(s) > 0$, and is *outgoing* from a state s if $s \xrightarrow{a} \mu$. Note that, as for PAs [17], there could be several moves associated to each action. We define the set of actions *enabled* in a state s by $\text{En}(s) \stackrel{\text{def}}{=} \{a \in \mathcal{A} \mid \exists \mu. s \xrightarrow{a} \mu\}$.

A finite (infinite) *path* $\lambda = s_0(a_0, \mu_0)s_1(a_1, \mu_1)s_2 \dots$ is a finite (infinite) sequence of alternating states and moves, such that $\varsigma(s_0) > 0$, and, for all $i \geq 0$,




Fig. 1: Three games: the game on the right is the composition of the normal forms of the other two games. Dirac distributions are shown as filled circles.

$s_i \xrightarrow{a_i} \mu_i$ and $\mu_i(s_{i+1}) > 0$. A finite path λ ends in a state, and we write $\text{last}(\lambda)$ for the last state of λ . We denote the set of finite (infinite) paths of a game G by Ω_G^+ (Ω_G), and by $\Omega_{G,\#}^+$ the set of paths ending in a Player $\#$ state, for $\# \in \{\diamond, \square\}$. A finite (infinite) *trace* is a finite (infinite) sequence of actions. Given a path, its trace is the sequence of actions along λ , with τ projected out. Formally, $\text{trace}(\lambda) \stackrel{\text{def}}{=} \text{PROJ}_{\{\tau\}}(a_0 a_1 \dots)$, where, for $\alpha \subseteq \mathcal{A} \cup \{\tau\}$, PROJ_α is the morphism defined by $\text{PROJ}_\alpha(a) = a$ if $a \notin \alpha$, and ϵ (the empty trace) otherwise.

Strategies. Nondeterminism for each player is resolved by a strategy. A *strategy* for Player $\#$, for $\# \in \{\diamond, \square\}$, is a function $\sigma_\# : \Omega_{G,\#}^+ \rightarrow \mathcal{D}(S_\square)$ such that $\sigma_\#(\lambda)(a, \mu) > 0$ only if $\text{last}(\lambda) \xrightarrow{a} \mu$. The set of Player $\#$ strategies in game G is denoted by $\Sigma_\#^G$. A strategy is called *memoryless* if, for each path λ , the choice $\sigma_\#(\lambda)$ is uniquely determined by $\text{last}(\lambda)$.

Normal Form of a Game. We can transform every game into its corresponding normal form, which does not affect the winning conditions. Transforming a game to normal form is the first step of our game composition.

Definition 2. *A game is in normal form if the following hold:*

- Every τ -transition $s \xrightarrow{\tau} \mu$ is from a Player \square state s to a Player \diamond state s' with a Dirac distribution $\mu = s'$.
- Every Player \diamond state s can only be reached by an incoming move (τ, s) . In particular, every distribution μ of a non- τ -transition, as well as the initial distribution, assigns probability zero to all Player \diamond states.

Given a game G without τ -transitions, one can construct its normal form $\mathcal{N}(G)$ by splitting every state $s \in S_\diamond$ into a Player \square state \bar{s} and a Player \diamond state \underline{s} , s.t.

- the incoming (resp. outgoing) moves of \bar{s} (resp. \underline{s}) are precisely the incoming (resp. outgoing) moves of s , with every Player \diamond state $t \in S_\diamond$ replaced by \bar{t} ;
- and the only outgoing (resp. incoming) move of \bar{s} (resp. \underline{s}) is (τ, \underline{s}) .

Intuitively, at \bar{s} the game is idle until Player \square allows Player \diamond to choose a move in \underline{s} . Hence, any strategy for G carries over naturally to $\mathcal{N}(G)$, and we can operate w.l.o.g. with normal-form games. Also, τ can be considered as a scheduling

choice. In the transformation to normal form, at most one such scheduling choice is introduced for each Player \square state, but in the composition more choices can be added, so that Player \square resolves nondeterminism arising from concurrency.

Game Unfolding. Strategy application is defined on the unfolded game. The *unfolding* $\mathcal{U}(G) = \langle \Omega_G^+, (\Omega_{G,\square}^+, \Omega_{G,\diamond}^+), \varsigma, \mathcal{A}, \longrightarrow' \rangle$ of the game G is such that $\lambda \xrightarrow{a'} \mu_{\lambda,a}$ if and only if $\text{last}(\lambda) \xrightarrow{a} \mu$ and $\mu_{\lambda,a}(\lambda(a, \mu)s) \stackrel{\text{def}}{=} \mu(s)$ for all $s \in S$.

An unfolded game is a set of trees (the roots are the support of the initial distribution), with potentially infinite depth, but finite branching. The entire history is stored in the states, so memoryless strategies suffice for unfolded games; formally, each strategy $\sigma_\diamond \in \Sigma_\diamond^G$ straightforwardly maps to a memoryless strategy $\mathcal{U}(\sigma_\diamond) \in \Sigma_\diamond^{\mathcal{U}(G)}$ by letting $\mathcal{U}(\sigma_\diamond)(\lambda)(a, \mu_{\lambda,a}) = \sigma_\diamond(\lambda)(a, \mu)$. We denote by $\mathcal{U}(G)_\square$ the set of moves of the unfolded form of a game G and by $\mathcal{U}(G)_{\#\square}$ the set of moves following a Player $\#$ state that is of the form $(a, \mu_{\lambda,a})$ with $\lambda \in \Omega_{G,\#}^+$. We remark that the unfolding of a normal form game is also in normal form.

2.1 Induced PA

When only one type of nondeterminism is present in a game, it is a *probabilistic automaton (PA)*. PAs are well-suited for compositional modelling [17], and can be used for verification, i.e. checking whether *all* behaviours satisfy a specification (when only Player \square is present), as well as strategy synthesis (when only Player \diamond is present) [14]. A PA is a game where $S_\diamond = \emptyset$ and $S_\square = S$, which we write here as $\langle S, \varsigma, \mathcal{A}, \longrightarrow \rangle$. This definition corresponds to modelling nondeterminism as an adverse, uncontrollable, environment, and so, by applying a Player \diamond strategy to a game to resolve the controllable nondeterminism, we are left with a PA where only uncontrollable nondeterminism for Player \square remains.

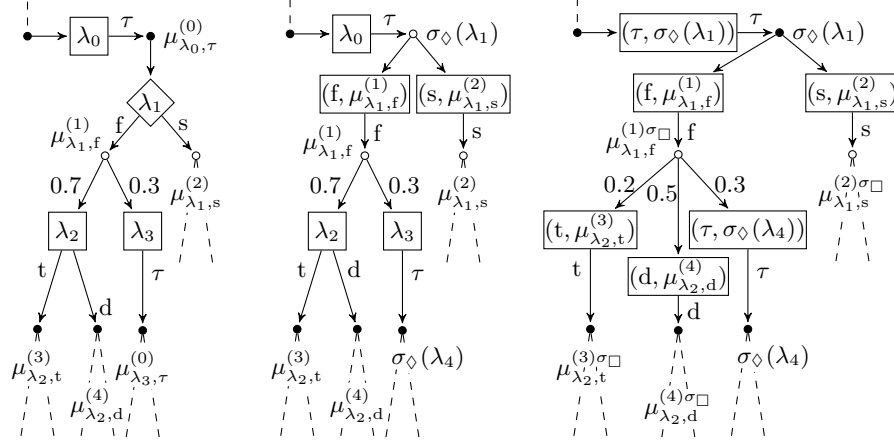
Definition 3. *Given an unfolded game $\mathcal{U}(G) = \langle \Omega_G^+, (\Omega_{G,\square}^+, \Omega_{G,\diamond}^+), \varsigma, \mathcal{A}, \longrightarrow \rangle$ in normal form and a strategy $\sigma_\diamond \in \Sigma_\diamond^G$, the induced PA is $G^{\sigma_\diamond} = \langle S', \varsigma, \mathcal{A}, \longrightarrow' \rangle$, where $S' \subseteq \Omega_{G,\square}^+ \cup \mathcal{U}(G)_{\diamond\square}$ is defined inductively as the reachable states, and*

- (I1) $\lambda \xrightarrow{\tau'} \mathcal{U}(\sigma_\diamond)(\lambda')$ iff $\lambda \xrightarrow{\tau} \lambda'$ (Player \diamond strategy chooses a move);
- (I2) $(a, \mu_{\lambda,a}) \xrightarrow{a'} \mu_{\lambda,a}$ for $(a, \mu_{\lambda,a}) \in \mathcal{U}(G)_{\diamond\square}$ (the chosen move is performed);
- (I3) $\lambda \xrightarrow{a'} \mu_{\lambda,a}$ iff $\lambda \xrightarrow{a} \mu_{\lambda,a}$ and $\lambda \in \Omega_{G,\square}^+$ (external transitions from older Player \square state remain unchanged).

The unfolded form of the game in Figure 1(right) is shown in Figure 2(a), and strategy application is illustrated in Figure 2(b).

2.2 Induced DTMC

A discrete-time Markov chain (DTMC) is a model for systems with probabilistic behaviour only. When applying a Player \square strategy to an induced PA, all nondeterminism is resolved, and a DTMC is obtained. A (labelled) *DTMC* D is a PA such that, for each $s \in S$, there is at most one transition $s \xrightarrow{a} \mu$.



(a) Game in Figure 1 (b) Induced PA via the (c) Induced DTMC via the strat-
 (right) unfolded. strategy σ_{\diamond} assigning egy $\sigma_{\square} = \mathcal{U}(\sigma_{\square}^{\square})$ assigning $\frac{2}{7}$ to t
 uniform probabilities. and $\frac{5}{7}$ to d .

Fig. 2: Game unfoldings. The dashed lines represent recursive application of unfolding and strategy application. Denote $\lambda_0 = \bar{s}u$, $\lambda_1 = \lambda_0(\tau, \mu_{\lambda_0, \tau}^{(0)})\underline{s}u$, $\lambda_2 = \lambda_1(f, \mu_{\lambda_1, f}^{(1)})\bar{s}u$, $\lambda_3 = \lambda_1(f, \mu_{\lambda_1, f}^{(1)})tu$, $\lambda_4 = \lambda_3(\tau, \mu_{\lambda_3, \tau}^{(0)})\underline{s}u$.

Definition 4. Given an unfolded PA $\mathcal{U}(M) = \langle S, \varsigma, \mathcal{A}, \longrightarrow \rangle$ and a strategy $\sigma_{\square} \in \Sigma_{\square}^M$, the induced DTMC $M^{\sigma_{\square}} = \langle S', \varsigma', \mathcal{A}, \longrightarrow' \rangle$ is such that $S' \subseteq \mathcal{U}(M)_{\square}$ is defined inductively as the states reachable via ς' and \longrightarrow' , where

- $(a, \mu_{\lambda, a}) \xrightarrow{a'} \mu_{\lambda, a}^{\sigma_{\square}}$, such that, for all moves $(b, \nu_{\lambda(a, \mu)t, b})$, we let the distribution $\mu_{\lambda, a}^{\sigma_{\square}}(b, \nu_{\lambda(a, \mu)t, b}) \stackrel{\text{def}}{=} \mu(t)\sigma_{\square}(\lambda(a, \mu)t)(b, \nu)$; and
- for all moves $(b, \nu_{t, b})$, $\varsigma'(b, \nu_{t, b}) \stackrel{\text{def}}{=} \varsigma(t)\sigma_{\square}(t)(b, \nu)$.

Note that an induced PA is already unfolded, and does not need to be unfolded again. We illustrate in Figure 2(c) the application of a Player \square strategy.

Probability Measures. We define the probability measure Pr_D of a DTMC D in the usual way. The *cylinder set* of a path $\lambda \in \Omega_D^+$ (resp. trace $w \in \mathcal{A}^*$) is the set of infinite paths (resp. traces) with prefix λ (resp. w). For a finite path $\lambda = s_0(a_0, \mu_0)s_1(a_1, \mu_1) \dots s_n$ we define $\text{Pr}_D(\lambda)$, the measure of its cylinder set, by: $\text{Pr}_D(\lambda) \stackrel{\text{def}}{=} \varsigma(s_0) \prod_{i=0}^{n-1} \mu_i(s_{i+1})$. We write $\text{Pr}_G^{\sigma_{\diamond}, \sigma_{\square}}$ (resp. $\text{Pr}_M^{\sigma_{\square}}$) for the measure $\text{Pr}_G^{\sigma_{\diamond}, \sigma_{\square}}$ (resp. $\text{Pr}_M^{\sigma_{\square}}$). The measures uniquely extend to infinite paths due to Carathéodory's extension theorem.

Given a finite trace w , $\text{paths}(w)$ denotes the set of minimal finite paths with trace w , i.e. $\lambda \in \text{paths}(w)$ if $\text{trace}(\lambda) = w$ and there is no path $\lambda' \neq \lambda$ with $\text{trace}(\lambda') = w$ and λ' being a prefix of λ . The measure of the cylinder set of w is $\text{Pr}_D(w) \stackrel{\text{def}}{=} \sum_{\lambda \in \text{paths}(w)} \text{Pr}_D(\lambda)$, and we call Pr_D the *trace distribution* of D .

Winning Conditions. Providing strategies for both players resolves all non-determinism in a game, resulting in a distribution over paths. A *specification* φ is a predicate on trace distributions, and for a DTMC D we write $D \models \varphi$ if $\varphi(P_D)$ holds. A specification φ is *defined on traces* if $\varphi(P_D) = \varphi(P_{D'})$ for all DTMCs D, D' such that $P_D(w) = P_{D'}(w)$ for all traces w .

3 Composing stochastic games

We now introduce composition operators for games and Player \diamond strategies, leading towards a framework for compositional synthesis in Section 4. Games can be composed of several component games, and our composition is inspired by interface automata [9], which have a natural interpretation as (concurrent) games.

3.1 Game Composition

We provide a synchronising composition of games so that controllability is preserved for Player \diamond , that is, actions controlled by Player \diamond in the components are controlled by Player \diamond in the composition. We endow each component game with an *alphabet* of actions \mathcal{A} , where synchronisation on *shared actions* in $\mathcal{A}^1 \cap \mathcal{A}^2$ is viewed as a (blocking) communication over ports, as in interface automata, though for simplicity we do not distinguish inputs and outputs. Synchronisation is multi-way and we do not impose input-enabledness of IO automata [8].

Given n games G^i in normal form with respective state spaces S^i , for $i \in I$ (let $I = \{1, \dots, n\}$ be the *index set* of the composed game), the state space of the composition is a subset of the Cartesian product $S^1 \times \dots \times S^n$, whose states contain at most one Player \diamond component thanks to the normal form. We denote by s^i the i th component of $\mathbf{s} \in S^1 \times \dots \times S^n$. Furthermore, every probability distribution in the composition is a product distribution. We say that a transition $\mathbf{s} \xrightarrow{a} \boldsymbol{\mu}$ *involves* the i th component if $s^i \xrightarrow{a} \mu^i$.

Definition 5. Given n games in normal form $G^i = \langle S^i, (S_\diamond^i, S_\square^i), \zeta^i, \mathcal{A}^i, \longrightarrow^i \rangle$, $i \in I$, their composition is the game $\parallel_{i \in I} G^i \stackrel{\text{def}}{=} \langle S, (S_\diamond, S_\square), \boldsymbol{\zeta}, \mathcal{A}, \longrightarrow \rangle$, where

- $S \subseteq S_\diamond \cup S_\square$, with $S_\square \subseteq S_\square^1 \times \dots \times S_\square^n$, and $S_\diamond \subseteq \{\mathbf{s} \in S^1 \times \dots \times S^n \mid \exists! i. s^i \in S_\diamond^i\}$ inductively defined to contain the states reachable from the initial distribution through the transition relation;
- $\boldsymbol{\zeta} = \zeta^1 \times \dots \times \zeta^n$; (note that, due to the normal form, $\boldsymbol{\zeta}(\mathbf{s}) > 0$ only if $\mathbf{s} \in S_\square$)
- $\mathcal{A} = \bigcup_{i=1}^n \mathcal{A}^i$;
- The transition relation \longrightarrow is defined such that
 - $\mathbf{s} \xrightarrow{a} \boldsymbol{\mu}$ for $a \neq \tau$ if
 - (C1) at least one component is involved;
 - (C2) the components involved in the transition are exactly those having a in their action alphabet;
 - (C3) for the uninvolved components j (equivalently, that do not have a in their action alphabet), the state remains the same ($\mu^j = s^j$);

- (C4) if \mathbf{s} is a $\text{Player } \diamond$ state then its only $\text{Player } \diamond$ component G^i is involved in the transition; and
- $\mathbf{s} \xrightarrow{\tau} \mathbf{t}$ if only one component G^i is involved, $\mathbf{s} \in S_{\square}$, and s.t. $\text{En}(\mathbf{t}) \neq \emptyset$.

We take the view that identity of the players must be preserved through composition to facilitate synthesis, and thus $\text{Player } \diamond$ actions of the individual components are controlled by a single $\text{Player } \diamond$ in the composition. $\text{Player } \square$ in the composition acts as a scheduler, controlling which component advances and, in $\text{Player } \square$ states, selecting among available actions, whether synchronised or not. Synchronisation in $\text{Player } \diamond$ states means that $\text{Player } \diamond$ in one component may indirectly control some $\text{Player } \square$ actions in another component. In particular, we can impose assume-guarantee contracts at the component level in the following sense. $\text{Player } \diamond$ of different components can cooperate to achieve a common goal: in one component $\text{Player } \diamond$ satisfies the goal B under an assumption A on its environment behaviour (i.e. $A \rightarrow B$), while $\text{Player } \diamond$ in the other component ensures that the assumption is satisfied, against all $\text{Player } \square$ strategies.

Our game composition is both associative and commutative, facilitating a modular model development, and is closely related to PA composition [17], with the condition (C4) added. As PAs are just games without $\text{Player } \diamond$ states, the game composition restricted to PAs is the same as classical PA composition. The condition $\text{En}(\mathbf{t}) \neq \emptyset$ for τ -transitions ensures that a $\text{Player } \diamond$ state is never entered if it were to result in deadlock introduced by the normal form transformation. Deadlocks that were present before the transformation are still present in the normal form. In the composition of normal form games, τ -transitions are only enabled in $\text{Player } \square$ states, and $\text{Player } \diamond$ states are only reached by such transitions; hence, composing normal form games yields a game in normal form.

In Figure 1, the game on the right is the composition of the normal forms of the two games on the left. The actions “f” and “s” are synchronised and controlled by $\text{Player } \diamond$ in \underline{su} . The “d” action is synchronised and controlled by $\text{Player } \square$ in both components, and so it is controlled by $\text{Player } \square$ in the composition, in tu . The action “t” is not synchronised, and thus available in $t\bar{v}$ and tu ; it is, however, not available in $t\underline{v}$, as it is a $\text{Player } \diamond$ state controlled by \underline{v} . The action “c” is also not synchronised, and is available in $t\underline{v}$. The “r” action is synchronised; it is available both in t and in \underline{v} , and hence also in $t\underline{v}$.

Constructing the composition of n components of size $|S|$ clearly requires time $\mathcal{O}(|S|^n)$. In strategy synthesis, the limiting factor is that applying the method on a large product game may be computationally infeasible. For example, the synthesis methods for multi-objective queries of [6] that we build upon are exponential in the number of objectives and polynomial in the size of the state space, and the theoretical bounds can be impractical even for small systems (see [7]). To alleviate such difficulties we focus on compositional strategy synthesis.

3.2 Strategy Composition

For compositional synthesis, we assume the following compatibility condition on component games: we require that actions controlled by $\text{Player } \diamond$ in one game are enabled and fully controlled by $\text{Player } \diamond$ in the composition.

Definition 6. Games G^1, \dots, G^n are compatible if, for every Player \diamond state $\mathbf{s} \in S_\diamond$ in the composition with $s^i \in S_\diamond^i$, if $s^i \xrightarrow{a} \mu^i$ then there is exactly one distribution ν , denoted by $\langle \mu^i \rangle_{\mathbf{s}, a}$, such that $\mathbf{s} \xrightarrow{a} \nu$ and $\nu^i = \mu^i$. (That is, for $i \neq i'$ such that $a \in \mathcal{A}^i$, there exists exactly one a -transition enabled in s^i .)

Our compatibility condition is analogous to that for single-threaded interface automata [9]. It remains to be seen if this condition can be relaxed without affecting preservation properties of the winning conditions.

Composing Strategies. Given a path λ of a composed game $\mathcal{G} = \parallel_{i \in I} G^i$, for each individual component G^i one can retrieve the corresponding path $[\lambda]^i$ that contains precisely the transitions G^i is involved in. The *projection* $[\cdot]^i : \Omega_{\mathcal{G}}^+ \rightarrow \Omega_{G^i}^+$ is defined inductively so that, for all states $\mathbf{t} \in S$ and paths $\lambda(a, \boldsymbol{\mu})\mathbf{t} \in \Omega_{\mathcal{G}}^+$ (with possibly $a = \tau$), we have $[\mathbf{s}]^i \stackrel{\text{def}}{=} s^i$; and inductively that $[\lambda(a, \boldsymbol{\mu})\mathbf{t}]^i \stackrel{\text{def}}{=} [\lambda]^i(a, \boldsymbol{\mu}^i)\mathbf{t}^i$ if $\text{last}(\lambda)^i \xrightarrow{a} \mu^i$, and $[\lambda]^i$ otherwise.

Recall that a Player \diamond state \mathbf{s} of the composed game has exactly one component s^i that is a Player \diamond state in G^i ; we say that the i th Player \diamond controls \mathbf{s} . Given a Player \diamond strategy for each component, the strategy for the composed game plays the strategy of the Player \diamond controlling the respective states.

Definition 7. Let σ_\diamond^i , $i \in I$, be Player \diamond strategies for compatible games. Their composition, $\sigma_\diamond = \parallel_{i \in I} \sigma_\diamond^i$, is defined such that $\sigma_\diamond(\lambda)(a, \langle \mu^i \rangle_{\mathbf{s}, a}) \stackrel{\text{def}}{=} \sigma_\diamond^i([\lambda]^i)(a, \mu^i)$ for all $\lambda \in \Omega_{\mathcal{G}}^+$ with $\mathbf{s} = \text{last}(\lambda) \in S_\diamond$.

From this definition, strategy composition is clearly associative. Note that, for each choice, the composed strategy takes into account the history of only one component, which is less general than using the history of the composed game.

3.3 Properties of the Composition

We now show that synthesising strategies for compatible individual components is sufficient to obtain a composed strategy for the composed game.

Functional Simulations. We introduce functional simulations, which are a special case of classical PA simulations [17], and show that they preserve winning conditions over traces. Intuitively, a PA M' functionally simulates a PA M if all behaviours of M are present in M' , and if strategies translate from M to M' .

Given a distribution μ , and a partial function $f : S \rightarrow S'$ defined on the support of μ , we write $\bar{f}(\mu)$ for the distribution defined by $\bar{f}(\mu)(s') \stackrel{\text{def}}{=} \sum_{f(s)=s'} \mu(s)$. A *functional simulation* from a PA M to a PA M' is a partial function $f : S \rightarrow S'$ such that $\bar{f}(\zeta) = \zeta'$, and if $s \xrightarrow{a} \mu$ in M then $f(s) \xrightarrow{a} \bar{f}(\mu)$ in M' .

Lemma 1. Given a functional simulation from a PA M to a PA M' and a specification φ defined on traces, for every strategy $\sigma_\square \in \Sigma_\square^M$ there is a strategy $\sigma'_\square \in \Sigma_\square^{M'}$ such that $(M')^{\sigma'_\square} \models \varphi \Leftrightarrow M^{\sigma_\square} \models \varphi$.

Key Lemma. The PA $\parallel_{i \in I} (G^i)^{\sigma_\diamond^i}$ is constructed by first unfolding each component, applying the Player \diamond strategies, and then composing the resulting PAs,

while the PA $(\|_{i \in I} G^i)^{\|_{i \in I} \sigma_{\diamond}^i}$ is constructed by first composing the individual components, then unfolding, and applying the composed **Player** \diamond strategy. The following lemma justifies, via the existence of a functional simulation, that composing **Player** \diamond strategies preserves the trace distribution between such PAs, and hence yields a fully compositional approach.

Lemma 2. *Given compatible games G^i , $i \in I$, and respective **Player** \diamond strategies σ_{\diamond}^i , there is a functional simulation from $(\|_{i \in I} G^i)^{\|_{i \in I} \sigma_{\diamond}^i}$ to $\|_{i \in I} (G^i)^{\sigma_{\diamond}^i}$.*

In general, there is no simulation in the other direction, as in the PA composition one can no longer distinguish states in the induced PA that were originally **Player** \diamond states, and so condition (C4) of the composition is never invoked.

4 Compositional Synthesis

Applying strategies synthesised for games to obtain induced PAs allows us to reuse compositional rules for PAs. Using Lemma 2, we can then lift the result back into the game domain. This process is justified in Theorem 1 below.

4.1 Composition Rules

We suppose the designer is supplying a game $\mathcal{G} = \|_{i \in I} G^i$ composed of *atomic games* G_i , together with specifications defined on traces φ_i , $i \in I$, and show how, using our framework, strategies σ_{\diamond}^i synthesised for G^i and φ_i can be composed to a strategy $\sigma_{\diamond} = \|_{i \in I} \sigma_{\diamond}^i$ for \mathcal{G} , satisfying a specification φ defined on traces.

Theorem 1. *Given a rule \mathfrak{P} for PAs \mathcal{M}_i and specifications φ_j^i and φ defined on traces, then the rule \mathfrak{G} holds for all **Player** \diamond strategies σ_{\diamond}^i of compatible games G^i with the same action alphabets as the corresponding PAs, where*

$$\mathfrak{P} \equiv \frac{\mathcal{M}_i \models \varphi_j^i \quad 1 \leq j \leq m \quad i \in I}{\|_{i \in I} \mathcal{M}_i \models \varphi} \quad \text{and} \quad \mathfrak{G} \equiv \frac{(G^i)^{\sigma_{\diamond}^i} \models \bigwedge_{j=1}^m \varphi_j^i \quad i \in I}{(\|_{i \in I} G^i)^{\|_{i \in I} \sigma_{\diamond}^i} \models \varphi}.$$

Theorem 1 enables the compositional synthesis of strategies in an automated way. First, synthesis is performed for atomic components G^i , $i \in I$, by obtaining for each i a **Player** \diamond strategy σ_{\diamond}^i for $G^i \models \bigwedge_{j=1}^m \varphi_j^i$. We apply \mathfrak{P} with $\mathcal{M}_i \stackrel{\text{def}}{=} (G^i)^{\sigma_{\diamond}^i}$ to deduce that φ holds in $\|_{i=1}^n (G^i)^{\sigma_{\diamond}^i}$ and, using Lemma 1 and 2, that $\|_{i \in I} \sigma_{\diamond}^i$ is a winning strategy for **Player** \diamond in $\|_{i=1}^n G^i$. The rules can be applied recursively, making use of associativity of the game and strategy composition.

4.2 Multi-Objective Queries

In this section we leverage previous work on compositional verification for PAs in order to compositionally synthesise strategies for games.

Reward and LTL Objectives. The *expected value* of a function $\rho : \mathcal{A}^* \rightarrow \mathbb{R}_{\pm\infty}$ over traces in a DTMC D is $\mathbb{E}_D[\rho] \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} \sum_{w \in \mathcal{A}^n} \Pr_D(w) \rho(w)$, if

the limit exists in $\mathbb{R}_{\pm\infty}$. We denote by $\mathbb{E}_G^{\sigma_\diamond, \sigma_\square}$ (resp. $\mathbb{E}_M^{\sigma_\square}$) the expected value in a game G (resp. PA M) under the respective strategies. A *reward structure* of a game with actions \mathcal{A} is a function $r : \mathcal{A}_r \rightarrow \mathbb{Q}$, where $\mathcal{A}_r \subseteq \mathcal{A}$. Given a reward structure r such that either $r(a) \leq 0$ or $r(a) \geq 0$ for all actions a occurring infinitely often on a path, the *total reward* for a trace $w = a_0 a_1 \dots$ is $\text{rew}(r)(w) \stackrel{\text{def}}{=} \lim_{t \rightarrow \infty} \sum_{i=0}^t r(a_i)$, which is measurable thanks to the restrictions imposed on r . Given reward structures r_j , $j \in J$, for all $a \in \bigcup_{j \in J} \mathcal{A}_{r_j}$ we let $(\sum_{j \in J} r_j)(a)$ be the sum of the $r_j(a)$ that are defined for a .

To express LTL properties over traces, we use the standard LTL operators (cf. [16]); in particular, the operators **F** and **G** stand for *eventually* and *always*, respectively. For a DTMC D , and an LTL formula Ξ over actions \mathcal{A}_Ξ , define $\text{Pr}_D(\Xi) \stackrel{\text{def}}{=} \text{Pr}_D(\{\lambda \in \Omega_D \mid \text{PROJ}_{\mathcal{A} \setminus \mathcal{A}_\Xi}(\text{trace}(\lambda)) \models \Xi\})$, that is, the measure of infinite paths with traces satisfying Ξ , where actions not in \mathcal{A}_Ξ are disregarded.

A *reward* (resp. *LTL*) *objective* is of the form $r \blacktriangleright v$ (resp. $\Xi \blacktriangleright v$), where r is a reward structure, Ξ is an LTL formula, $v \in \mathbb{Q}$ is a bound, and $\blacktriangleright \in \{\geq, >\}$. A reward objective $r \blacktriangleright v$ (resp. LTL objective $\Xi \blacktriangleright v$) is true in a game G under a pair of strategies $(\sigma_\diamond, \sigma_\square)$ if and only if $\mathbb{E}_G^{\sigma_\diamond, \sigma_\square}[\text{rew}(r)] \blacktriangleright v$ (resp. $\text{Pr}_G^{\sigma_\diamond, \sigma_\square}(\Xi) \blacktriangleright v$), and similarly for PAs and DTMCs. Minimisation of rewards can be expressed by reverting signs.

Multi-Objective Queries. A *multi-objective query* (MQ) φ is a Boolean combination of reward and LTL objectives, and its truth value is defined inductively on its syntax. An MQ φ is a *conjunctive query* (CQ) if it is a conjunction of objectives. Given an MQ with bounds v_1, v_2, \dots , call $\mathbf{v} = (v_1, v_2, \dots)$ the *target*. Denote by $\varphi[\mathbf{x}]$ the MQ φ , where, for all i , $r_i \blacktriangleright v_i$ is replaced by $r_i \blacktriangleright x_i$, and $\Xi_i \blacktriangleright v_i$ is replaced by $\Xi_i \blacktriangleright x_i$. Given a game G (resp. PA M) we write $G^{\sigma_\diamond, \sigma_\square} \models \varphi$ (resp. $M^{\sigma_\square} \models \varphi$), if the query φ evaluates to true under the respective strategies. We write $M \models \varphi$ if $M^{\sigma_\square} \models \varphi$ for all $\sigma_\square \in \Sigma_\square^M$. We say that an MQ φ is *achievable* in a game G if there is a **Player** \diamond strategy σ_\diamond such that $G^{\sigma_\diamond} \models \varphi$, that is, σ_\diamond is winning for φ against all possible **Player** \square strategies. We require that expected total rewards are bounded, that is, we ask for any reward structure r in an MQ φ that $G^{\sigma_\diamond, \sigma_\square} \models r < \infty \wedge r > -\infty$ for all σ_\diamond and σ_\square .

Fairness. Since PA rules as used in Theorem 1 often include fairness conditions, we recall here the concept of unconditional process fairness based on [1]. Given a composed PA $\mathcal{M} = \parallel_{i \in I} M^i$, a strategy σ_\square is *unconditionally fair* if $\mathcal{M}^{\sigma_\square} \models \mathbf{u}$, where $\mathbf{u} \stackrel{\text{def}}{=} \bigwedge_{i \in I} \text{GF } \mathcal{A}_i \geq 1$, that is, each component makes progress infinitely often with probability 1. We write $\mathcal{M} \models^{\mathbf{u}} \varphi$ if, for all unconditionally fair strategies $\sigma_\square \in \Sigma_\square$, $\mathcal{M}^{\sigma_\square} \models \varphi$; this is equivalent to $\mathcal{M} \models \mathbf{u} \rightarrow \varphi$ (the arrow \rightarrow stands for the standard logical implication), and so MQs can incorporate fairness.

Applying Theorem 1. In particular, for the premises in Theorem 1 we can use the compositional rules for PAs developed in [13], which are stated for MQs. Thus, the specification φ for the composed game can, for example, be a CQ, or a summation of rewards, among others. Unconditional fairness corresponds precisely to the fairness conditions used in the PA rules of [13]. When the PA rules

include fairness assumptions, note that, for a single component, unconditional fairness is equivalent to only requiring deadlock-freedom.

4.3 Compositional Pareto Set Computation

We describe in this section how to pick the targets of the objectives φ_i in compositional rules, such as those in Theorem 1, so that φ is achievable.

Pareto Sets. Given an MQ φ with N objectives, vector $\mathbf{v} \in \mathbb{R}^N$ is a *Pareto vector* if and only if $\varphi[\mathbf{v} - \varepsilon]$ is achievable for all $\varepsilon > 0$, and $\varphi[\mathbf{v} + \varepsilon]$ is not achievable for any $\varepsilon > 0$. The downward closure of the set of all such vectors is called a *Pareto set*, where the downward closure of a set X is defined as $\text{dwc}(X) \stackrel{\text{def}}{=} \{\mathbf{y} \in \mathbb{R}^N \mid \exists \mathbf{x} \in X. \mathbf{x} \geq \mathbf{y}\}$. Given $\varepsilon > 0$, an ε -*approximation of a Pareto set* P is a set of vectors Q satisfying that, for any $\mathbf{w} \in Q$, there is a vector $\mathbf{v} \in P$ such that $\|\mathbf{v} - \mathbf{w}\| \leq \varepsilon$, and for every $\mathbf{v} \in P$ there is a vector $\mathbf{w} \in Q$ such that $\|\mathbf{v} - \mathbf{w}\| \leq \varepsilon$, where $\|\cdot\|$ is the Manhattan norm.

Under-approximating Pareto Sets. We can compositionally compute an under-approximation of the Pareto set for φ , which we illustrate in Figure 3.

Consider N reward structures, r_1, \dots, r_N , and objectives $\varphi^i, i \in I$, over these reward structures for respective games G^i , as well as an objective φ , over the same reward structures, for the composed game $\mathcal{G} = \parallel_{i \in I} G^i$. Note that, for each $1 \leq j \leq N$, the reward structure r_j may be present in several objectives φ^i . Let P^i be the Pareto set for $G^i \models \varphi^i$, for $i \in I$, and so each point $\mathbf{v}^{(i)} \in P^i$ represents a target vector for the MQ $\varphi^i[\mathbf{v}^{(i)}]$ achievable in the game G^i .

For a Pareto set P^i , define the *lifting* $[P^i]$ to all N reward structures by $[P^i] \stackrel{\text{def}}{=} \{\mathbf{v} \in \mathbb{R}_{\pm\infty}^N \mid \text{the coordinates of } \mathbf{v} \text{ appearing in } \varphi^i \text{ are in } P^i\}$. The set $P' \stackrel{\text{def}}{=} \bigcap_{i \in I} [P^i]$ is the set of target vectors for all M reward structures, which are consistent with achievability of all objectives φ^i in the respective games. The projection¹ P'' of P' onto the space of reward structures appearing in φ then yields an under-approximation of the Pareto set P for φ in the composed game \mathcal{G} , that is, $P'' \subseteq P$. A point $\mathbf{v} \in P''$ can be achieved by instantiating the objectives φ^i with any targets $\mathbf{v}^{(i)}$ in P' that match \mathbf{v} .

4.4 Compositional MQ Synthesis

We now describe our compositional strategy synthesis method.

MQ Synthesis for Component Games. A game is *stopping* if, under any strategy pair, with probability 1 a part of the game is reached where the properties no longer change (see the technical report [2] for details.) A strategy is ε -*optimal* for an MQ φ with target \mathbf{v} if it achieves $\varphi[\mathbf{v} - \varepsilon]$ for all $\varepsilon > 0$. From [6] we have that, for atomic stopping games with MQs, it is decidable whether an ε -optimal strategy exists (optimal strategies may not exist), and, ε -optimal strategies can be represented finitely using stochastic memory update [7].

¹ More generally, if φ contains items such as $r_i + r_j \blacktriangleright v_i + v_j$, as in the (SUM-REWARD) rule of [13], a new dimension is introduced combining the rewards as required.

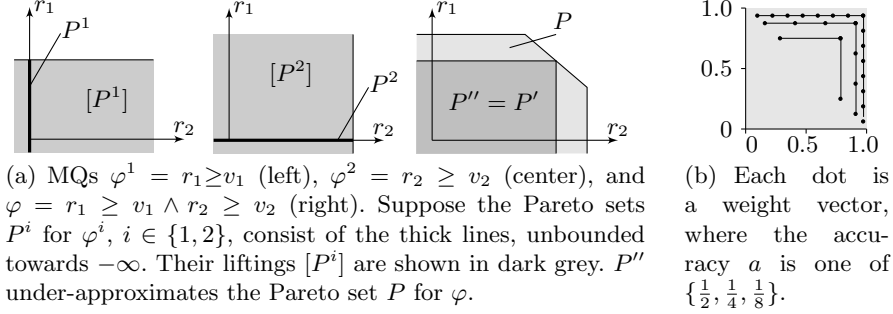


Fig. 3: Compositional Pareto set computation (a); weight vector selection (b).

We compute a strategy for an MQ in CNF $\bigwedge_{i=1}^n \bigvee_{j=1}^m r_{i,j} \geq v_{i,j}$ by implementing value iteration based on [6]. First, set an initial accuracy, $a \leftarrow \frac{1}{2}$. For each $0 \leq i < n$, select a corresponding $0 \leq j_i < m$. Then uniformly iterate over *weights* $\mathbf{x}_i \in [0, 1 - a/2]^m$ by gridding with accuracy a , keeping the j_i th dimension constant at $1 - a/2$. The pattern of selected vectors is shown for $m = 2$ dimensions in Figure 3(b). At each selection of \mathbf{x}_i , check, using the CQ algorithm of [7], if $\bigwedge_{i=1}^n (\sum_{j=1}^m \mathbf{x}_i^j \cdot r_{i,j} \geq \sum_{j=1}^m \mathbf{x}_i^j \cdot v_{i,j})$ is realisable, and, if so, return the winning strategy. Otherwise, if all options for selecting j_i are exhausted, refine the accuracy to $a \leftarrow \frac{a}{2}$ and repeat.

Every point $\mathbf{y} \in \mathbb{R}^n$ in a CQ Pareto set with weights $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}_{\geq 0}^m$ corresponds to intersection of half-spaces $\mathbf{x}_i \cdot \mathbf{z} \geq y^i$; the union over all choices of weight vectors is the ε -approximation of the corresponding MQ Pareto set.

MQ Synthesis for Composed Games. Our method for compositional strategy synthesis, based on synthesis for atomic games, is summarised as follows:

- (S1) **User Input:** A composed game $\mathcal{G} = \parallel_{i \in I} G^i$, MQs φ^i , φ , and matching PA rules for use in Theorem 1.
- (S2) **First Stage:** Obtain ε -approximate Pareto sets P^i for $G^i \models \varphi^i$, and compute P'' as in Section 4.3.
- (S3) **User Feedback:** Pick targets \mathbf{v} for φ from P'' and matching targets $\mathbf{v}^{(i)}$ for φ^i from P^i .
- (S4) **Second Stage:** Synthesise strategies σ_{\diamond}^i , for $G^i \models \varphi^i[\mathbf{v}^{(i)}]$, and compose them using Definition 7 (see the technical report [2] for composing strategies in the stochastic memory update representation.)
- (S5) **Output:** A composed strategy $\parallel_{i \in I} \sigma_{\diamond}^i$, winning for $\mathcal{G} \models \varphi[\mathbf{v}]$ by Theorem 1.

Steps (S1), (S4) and (S5) are sufficient if the targets are known, while (S2) and (S3) are an additional feature enabled by the Pareto set computation.

4.5 Case Study

We illustrate our approach with an example, briefly outlined here; see the technical report [2] for more detail. We model an Internet-enabled fridge that au-

Table 1: Run time comparison between compositional and monolithic strategy synthesis. For the CQ value iteration (cf. [7]) we use 60, 400 and 200 iterations for the fridge, trader and composed model, respectively. Computations were done on a 2.8 GHz Intel[®] Xeon[®] CPU, with 32 GB RAM, under Fedora 14.

Traders (n)	State Space Size		Running Time [s]					
			Composition		Compositional		Monolithic	
	F	T_i	C	Pareto Set	Strategies	Pareto Set	Strategy	
1	11	7	17	0.006	31.0	0.2	10.9	0.26
2	23	7	119	0.1	400.0	0.37	6570.0	161.0
3	39	7	698	2.0	407.0	2.4	> 3h	–
4	59	7	3705	75.0	4870.0	1.4	> 5h	–

tonomously selects between different digital agents selling milk whenever restocking is needed. We compute the Pareto sets and strategies in a prototype implementation as an extension of PRISM-games [5].

The fridge repeatedly invites offers from several traders, and decides whether to accept or decline the offers, based on the quality of each offer. The objective is for the traders to maximise the unit price, and for the fridge to maximise the amount of milk it purchases. For n traders T_i , $1 \leq i \leq n$, and a fridge F , denote the composition $C \stackrel{\text{def}}{=} (\|_{i=1}^n T_i) \parallel F$. We use the following reward objectives $O_i \equiv$ “offers made by T_i ” $\geq v_{o_i}$, $A_i \equiv$ “offers of T_i accepted” $\geq v_{a_i}$, $Q_i \equiv$ “quality of offers made by T_i ” $\geq v_{q_i}$, $\$i \equiv$ “unit price of T_i ” $\geq v_{\$i}$, and $\# \equiv$ “amount of milk obtained by F ” $\geq v_{\#}$, and synthesise strategies as explained in Section 4.1 according to the rule:

$$\frac{F \models \bigwedge_{j=1}^n (O_j \rightarrow A_j) \wedge (\bigwedge_{j=1}^n Q_j \rightarrow \#) \quad T_i \models A_i \rightarrow (Q_i \wedge \$i) \quad 1 \leq i \leq n}{C \models \bigwedge_{j=1}^n (O_j \rightarrow \$j) \wedge (\bigwedge_{j=1}^n O_j \rightarrow \#)}.$$

The main advantages of compositional synthesis are a dramatic improvement in efficiency and the compactness of strategies, as indicated in Table 1. In general, the strategies are randomised and history dependent. For the case of two traders, with the target that we selected, we generate a strategy where the traders make an expensive offer in the first round with probability 0.91, but from then on consistently make less expensive bulk offers.

5 Conclusion

We have defined a synchronising composition for stochastic games, and formulated a compositional approach to controller synthesis by leveraging techniques for compositional verification of PAs [13] and multi-objective strategy synthesis of [6,7]. We have extended the implementation of [7] to synthesise ε -optimal strategies for two-player stochastic games for total expected reward objectives in conjunctive normal form. We intend to investigate relaxing the compatibility

condition and consider notions of fairness weaker than unconditional fairness to broaden the applicability of our methods.

Acknowledgements. The authors thank Dave Parker and Vojtěch Forejt for helpful discussions. The authors are partially supported by ERC Advanced Grant VERIWARE.

References

1. C. Baier, M. Größer, and F. Ciesinski. Quantitative analysis under fairness constraints. In *ATVA'09*, volume 5799 of *LNCS*, pages 135–150. Springer.
2. N. Basset, M. Kwiatkowska, and C. Wiltsche. Compositional controller synthesis for stochastic games. University of Oxford, 2014. Technical Report CS-RR-14-05.
3. M. Campbell, M. Egerstedt, J. P. How, and R. M. Murray. Autonomous driving in urban environments: approaches, lessons and challenges. *Phil. Trans. R. Soc. A*, 368(1928):4649–4672, 2010.
4. K. Chatterjee, M. Jurdziński, and T. Henzinger. Simple stochastic parity games. In *CSL'03*, volume 2803 of *LNCS*, pages 100–113. Springer.
5. T. Chen, V. Forejt, M. Kwiatkowska, D. Parker, and A. Simaitis. PRISM-games: A model checker for stochastic multi-player games. In *TACAS'13*, volume 7795 of *LNCS*, pages 185–191. Springer.
6. T. Chen, V. Forejt, M. Kwiatkowska, A. Simaitis, and C. Wiltsche. On stochastic games with multiple objectives. In *MFCS'13*, volume 8087 of *LNCS*, pages 266–277. Springer.
7. T. Chen, M. Kwiatkowska, A. Simaitis, and C. Wiltsche. Synthesis for multi-objective stochastic games: An application to autonomous urban driving. In *QEST'13*, volume 8054 of *LNCS*, pages 322–337. Springer.
8. L. Cheung, N. Lynch, R. Segala, and F. Vaandrager. Switched PIOA: Parallel composition via distributed scheduling. *TCS*, 365(1–2):83–108, 2006.
9. L. de Alfaro and T.A. Henzinger. Interface automata. *SIGSOFT Softw. Eng. Notes*, 26(5):109–120, 2001.
10. J. Filar and K. Vrieze. *Competitive Markov decision processes*. Springer, 1996.
11. M. Gelderie. Strategy composition in compositional games. In *ICALP'13*, volume 7966 of *LNCS*, pages 263–274. Springer.
12. S. Ghosh, R. Ramanujam, and S. Simon. Playing extensive form games in parallel. In *CLIMA'10*, volume 6245 of *LNCS*, pages 153–170. Springer.
13. M. Kwiatkowska, G. Norman, D. Parker, and H. Qu. Compositional probabilistic verification through multi-objective model checking. *Information and Computation*, 232:38–65, 2013.
14. M. Kwiatkowska and D. Parker. Automated verification and strategy synthesis for probabilistic systems. In *ATVA'13*, volume 8172 of *LNCS*, pages 5–22. Springer.
15. N. Ozay, U. Topcu, R.M. Murray, and T. Wongpiromsarn. Distributed synthesis of control protocols for smart camera networks. In *ICCP'11*, pages 45–54. IEEE.
16. A. Pnueli. The temporal logic of programs. In *FOCS'77*, pages 46–57. IEEE.
17. R. Segala. *Modelling and Verification of Randomized Distributed Real Time Systems*. PhD thesis, Massachusetts Institute of Technology, 1995.
18. Lloyd S Shapley. Stochastic games. *Proc. Natl. Acad. Sci. USA*, 39(10):1095, 1953.
19. A. Sokolova and E.P. de Vink. Probabilistic automata: system types, parallel composition and comparison. In *VOSS'04*, volume 2925 of *LNCS*, pages 1–43. Springer.