



Concurrent Stochastic Lossy Channel Games

Daniel Stan ✉ 🏠 

EPITA, Le Kremlin-Bicêtre, France

Muhammad Najib ✉ 🏠 

Heriot-Watt University, Edinburgh, UK

Anthony Widjaja Lin ✉ 🏠 

University of Kaiserslautern-Landau, Germany

Max-Planck Institute for Software Systems, Kaiserslautern, Germany

Parosh Aziz Abdulla ✉ 🏠 

Uppsala University, Sweden

Abstract

Concurrent stochastic games are an important formalism for the *rational verification* of probabilistic multi-agent systems, which involves verifying whether a temporal logic property is satisfied in some or all game-theoretic equilibria of such systems. In this work, we study the rational verification of probabilistic multi-agent systems where agents can cooperate by communicating over *unbounded lossy channels*. To model such systems, we present *concurrent stochastic lossy channel games* (CSLCG) and employ an equilibrium concept from cooperative game theory known as the *core*, which is the most fundamental and widely studied cooperative equilibrium concept. Our main contribution is twofold. First, we show that the rational verification problem is undecidable for systems whose agents have almost-sure LTL objectives. Second, we provide a decidable fragment of such a class of objectives that subsumes almost-sure reachability and safety. Our techniques involve reductions to solving infinite-state zero-sum games with conjunctions of qualitative objectives. To the best of our knowledge, our result represents the first decidability result on the rational verification of stochastic multi-agent systems on infinite arenas.

2012 ACM Subject Classification Theory of computation → Formal languages and automata theory; Theory of computation → Verification by model checking; Theory of computation → Concurrency; Theory of computation → Solution concepts in game theory

Keywords and phrases concurrent, games, stochastic, lossy channels, wqo, finite attractor property, cooperative, core, Nash equilibrium

Digital Object Identifier 10.4230/LIPIcs.CSL.2024.46

Related Version *Full Version:* <http://arxiv.org/abs/2311.17037>

Funding *Anthony Widjaja Lin:* supported by European Research Council under European Union's Horizon research and innovation programme (grant agreement no 501100000781).

Acknowledgements We wish to thank Richard Mayr and all anonymous reviewers for their useful feedback.

1 Introduction

Rational verification concerns the problem of checking which temporal logic properties will be satisfied in game-theoretic equilibria of a multi-agent system, that is, the stable collective behaviours that arise assuming that agents choose strategies/policies rationally in order to achieve their goals [27, 1]. The usual approach to rational verification is to model multi-agent systems as concurrent games [27, 31]. This involves converting a multi-agent system into a game where agents are represented by a collection of independent, self-interested players in a finite-state environment. The game is played over an infinite number of rounds, with each



© Daniel Stan, Muhammad Najib, Anthony Widjaja Lin, and Parosh Aziz Abdulla;
licensed under Creative Commons License CC-BY 4.0

32nd EACSL Annual Conference on Computer Science Logic (CSL 2024).

Editors: Aniello Murano and Alexandra Silva; Article No. 46; pp. 46:1–46:19

Leibniz International Proceedings in Informatics



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


player/agent (we use these terms interchangeably throughout the paper) choosing an action to perform in each round. Each player’s goal is typically given by a temporal logic formula, which the player aims to satisfy. The temporal logic formula may or may not be satisfied by the infinite plays generated from the game, assuming that the players act rationally to achieve their goals.

In this paper, unlike much of previous work in rational verification, we consider systems that give rise to games with *probabilistic transitions* and *infinitely many states*. In particular, we focus on systems that can naturally be modelled by *stochastic lossy channel games* [8] in the *multi-player* and *concurrent* setting, where there are $n \geq 2$ players who can make concurrent moves. Our setting generalises the two-player turn-based framework presented in [8]. We call this model *Concurrent Stochastic Lossy Channel Games* (CSLCG). This model can be used to analyse a wide class of systems that communicate through potentially unreliable FIFO channels, such as communication networks, timed protocols, distributed systems, and memory systems [2, 9, 10, 4]. Stochasticity can be used to represent uncertainty in both the environment (e.g., branching and message losses) and the behaviour of agents. Incorporating such uncertainty is desirable from a practical standpoint, as real-world systems are expected to operate correctly even when communication is not perfect and agents’ behaviour is not deterministic. In the context of memory systems, stochasticity is used as a fairness condition that prevents unrealistic scenarios where the shared memory is never updated by the processes [3, 34].

Given the possibility of communication, albeit imperfect, among agents, it seems quite natural to assume that some form of cooperation may arise in games. Thus, a relevant and fundamental question within the rational verification framework is: “*What temporal logic property is satisfied by the rational cooperation that emerges in such a setting?*” To address this question, we consider an equilibrium concept from cooperative game theory called the *core* [14, 41, 30], which is the most fundamental and widely-studied cooperative equilibrium concept. With this concept, the standard assumption is that there exists some mechanism¹ that the players in a game can use to make *binding agreements*. These binding agreements enable players to cooperate and work in teams/coalitions, providing a way to eliminate undesirable equilibria that may arise in non-cooperative settings [30, 29]. We illustrate that this is also true in our setting in Example 14. Despite using a cooperative equilibrium concept, we emphasize that players are still self-interested, meaning they rationally pursue their individual goals. As such, the games we consider in this work are general-sum games instead of strictly positive-sum games, which are purely cooperative.

Contributions. We study the rational verification problem in CSLCG with the core as the key equilibrium concept. It is shown in [30] that the core of a game with *qualitative objectives* can never be empty, which also applies to the setting considered in this paper. Thus, two relevant decision problems pertaining to rational verification in the present work are E-CORE and A-CORE. E-CORE asks whether there *exists* a strategy profile in the core satisfying a given property Γ , whereas A-CORE asks whether *all* profiles in the core satisfy Γ . We first show that these problems are undecidable for games in which the players’ objectives and property Γ are almost-sure LTL formulae, i.e., of the form $AS(\varphi)$, where φ is a LTL formula. We consider LTL *with regular valuations* [24] where the set of *states/configurations* satisfying an atomic proposition is represented by a regular language. Then, for our main

¹ Such mechanism is assumed to be exogenous (e.g., via contracts) and beyond the scope of the present paper.



■ **Figure 1** A simple adversarial transmission system from Example 1 (left) and its graphical representation as a 2-player arena (right). The special action $*$ is a shorthand for any possible action. The resulting operation f on the channel is written after $|$ and omitted if $f = \text{nop}$.

contribution, we show the following: when players' goals are given as *almost-sure reachability* or *almost-sure safety* objectives, and the property Γ is given as *almost-sure reachability*, *almost-sure safety*, or *almost-sure Büchi*, the problems of E-CORE and A-CORE become decidable. Our decidability proof is obtained via a reduction to *concurrent 2.5-player*² lossy channel games with conjunction of objectives. This approach differs from previous work in two ways. First, our reduction considers *concurrent* plays, in contrast to *turn-based 2.5-player* games considered by [8]. Second, we do not assume finite-memory strategies, as opposed to finite-memory assumption in [7, 16]. This is because finite-memory strategies do not ensure determinacy [35]: it is possible that none of the players has a winning strategy in a given concurrent 2.5-player game, even in the finite state case with simple objectives [23]. Therefore, general strategies (which may require infinite memory) are required in equilibrium concepts such as the core, where players (or coalitions) may try to satisfy their objectives while simultaneously preventing other players from achieving theirs. However, the main challenge in using these strategies in the infinite state case is the issue of representation. To address this, we provide a novel encoding of strategies in our proof of decidability. To our knowledge, this is the first decidability result on the rational verification of stochastic multi-player games with infinite-state arenas.

► **Example 1.** To illustrate the model, we consider a simple transmission system depicted in Figure 1 throughout the article. In this example, *Sender* (player 1) tries to emit some message of either type a or b . *Attacker* (player 2) is trying to scramble the communication by concurrently choosing the same message type (action a or b). Moreover, Attacker cannot scramble the communication two times in a row and has to wait (action w) otherwise. The CSLCG arena is depicted on the right with the corresponding transitions and an extra location, reachable by a unilateral decision of player 1 by reading a c -letter from the channel, which is possible only in case of a successful scrambling. Note that although the game structure is deterministic in this example, some stochastic behaviour will still appear both from message losses and from players' strategies, which are played concurrently. As a more concrete example of Attacker's objective, one could specify the condition "reaching l_2 almost-surely, while not having more than 3 queued messages with positive probability". Note that this is a conjunction of reachability and safety conditions over locations and regular sets of channel configurations.

² Henceforth, we use the usual terms 1.5-player and 2.5-player games for, respectively, one-player and two-player stochastic games.

Related Work. As already mentioned above, the most relevant work w.r.t. verification of CSLCG is [8], which shows decidability of *two-player turn-based* stochastic lossy channel games with almost-sure reachability or almost-sure Büchi objectives. This work was extended to parity conditions in [7], where decidability can be shown *assuming finite-memory strategies*; otherwise, it is already undecidable for 1.5-player games over lossy channel systems with almost-sure co-Büchi objectives [16]. We are not aware of any work on rational verification of concurrent stochastic games over infinite arenas. [26] studies verification of the core in a probabilistic setting, while [11] presents *Probabilistic Strategy Logic*, which can be used to characterize the core. However, both of these works are in a finite state setting only. Without probability, we mention the work [37, 22] on concurrent (deterministic) pushdown games with multiple players. In particular, ATL* model checking is decidable in such games, which allows one to reason about the core. Additionally, there has been work on pushdown *module-checking*, which provides some element of non-determinism through an (external) environment. [12] examines the imperfect information setting, while [21] studies multi-agent systems with ATL* specifications. Note that lossy channel systems, which are the focus of our work, are inherently different from the models considered in these studies.

Organization. Section 2 introduces preliminary definitions and notations. Section 3 describes concurrent lossy channel games and the special case of 2.5-player zero-sum games. Section 4 presents a characterization of the core, the problems E-CORE and A-CORE, a procedure to solve them, and an undecidability result of E-CORE and A-CORE. Section 5 addresses the computability of winning regions for concurrent 2.5-player zero-sum games and provides algorithms to compute such regions. Section 6 studies the conjunction of objectives, while Section 7 presents our main result on the decidability of E-CORE and A-CORE. Finally, Section 8 concludes with a discussion and future work.

2 Preliminaries

For a finite alphabet Σ , the set of finite sequences, called *words*, is written Σ^* . Given two words $u, v \in \Sigma^*$, we write $u \cdot v$ for their concatenation and extend this notation to sets of words. Given $L \subseteq \Sigma^*$, L^+ denotes the smallest set containing L and closed under concatenation and $L^* = \{\epsilon\} \cup L^+$ with ϵ the empty word. The class of *regular languages* is the smallest class containing Σ , closed under difference, union, Kleene star and concatenation. We refer to [33] for further references about regular expressions and their link to automata theory.

Let S denote a countable set, for example $S = \Sigma^*$. A *well-quasi-ordering* [25] (wqo) \preceq over S is a quasi-ordering (i.e. reflexive and transitive binary relation) such that any infinite sequence $(s_i)_{i \in \mathbb{N}}$ of elements of S contains an increasing pair $i < j$ such that $s_i \preceq s_j$. As an example, Higman's lemma [32] states that the *sub-word ordering* \preceq defined below is a wqo over Σ^* :

► **Definition 2.** For any $w, w' \in \Sigma^*$, $w \sqsubseteq w'$ if w can be written $w = w_0 \cdot w_1 \cdots w_n$ and $w' \in \Sigma^* \cdot \{w_0\} \cdot \Sigma^* \cdots \Sigma^* \cdot \{w_n\} \cdot \Sigma^*$.

A subset $U \subseteq S$ is *upward-closed* (UC) if for every $s \preceq t$ such that $s \in U$, we also have $t \in U$. Any UC set can be uniquely represented by its set of minimal elements, which is finite (wqo property). A *downward-closed* (DC) set is defined in a similar manner. In particular, any UC or DC set w.r.t. the sub-word ordering is a regular language.

A *distribution over S* is an array $\delta \in \mathbb{R}_{\geq 0}^S$ of values $\delta[s] \in \mathbb{R}_{\geq 0}$ for any $s \in S$, such that $\sum_{s \in S} \delta[s] = 1$. If X is finite and non-empty, we write $\mathcal{U}(X)$ for the *uniform distribution over X* , namely $\forall x \in X, \mathcal{U}(X)[x] = 1/|X|$. We use the notation $\text{Dist}(S)$ to denote the set of distributions over S .

Let S^ω denote the set of infinite sequences over S , called *paths*. A set $X \subseteq S^\omega$ is a *cylinder* if it is of the form $s_0 \cdots s_n \cdot S^\omega$ for some $s_0 \cdots s_n \in S^*$. Such a set is written $\text{Cyl}(s_0 \cdots s_n)$. We introduce $\mathcal{F}(S)$ as the smallest family of sets of paths containing all cylinders, closed under complementation, and countable union. Such sets of $\mathcal{F}(S)$ are called *measurable*.

Given an initial state $s_0 \in S$ and a mapping $\eta : S^+ \rightarrow \text{Dist}(S)$ from *histories* to distributions over S , we define a partial function $\mathbb{P} : \mathcal{F}(S) \rightarrow \mathbb{R}_{\geq 0}$ such that $\mathbb{P}(\text{Cyl}(s_0)) = 1$, $\forall s \neq s_0, \mathbb{P}(\text{Cyl}(s)) = 0$, and for all $h \cdot s \in S^+ \cdot S$, $\mathbb{P}(\text{Cyl}(h \cdot s)) = \mathbb{P}(\text{Cyl}(h)) \cdot \eta(h)[s]$. Carathéodory's criterion [40] ensures this definition is well and uniquely defined on all $\mathcal{F}(S)$, and \mathbb{P} is therefore called a *probability measure*.

We describe infinite paths with the logic LTL, whose usual semantics over infinite words in S^ω , leads to measurable sets [17, Remark 10.57]. More precisely, we consider LTL with regular valuations [24], where atomic propositions are represented by regular languages. We focus on the fragment without the “until” operator:

► **Definition 3** (LTL [39]). *An LTL(\circ, \diamond) formula is any φ in the following grammar, where ν ranges over regular languages over S :*

$$\varphi ::= \nu \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \circ\varphi \mid \diamond\varphi \mid \square\varphi$$

Here, $\square\varphi$, $\diamond\varphi$ and $\square\diamond\varphi$ denote *always* φ , *eventually* φ and *infinitely often* φ , respectively.

Let I be a set of indices and V be a set of values. A *profile* \vec{v} is a mapping from I to V , where v_i is the value assigned to i . In particular if $I = \{1 \dots n\}$, then $\vec{v} = (v_1 \dots v_n)$. We introduce a fresh symbol $\perp \notin V$, and define the notation \vec{v}_{-i} as the profile where v_i has been replaced by \perp . Given any other value $w \in V$, we write (\vec{v}_{-i}, w) for the profile where the value assigned to i has been replaced by w . We extend these notations to a subset $Y \subseteq I$ in the usual way, i.e., \vec{v}_{-Y} denotes \vec{v} where each $v_i, i \in Y$ is replaced by \perp and $(\vec{v}_{-Y}, \vec{v}'_Y)$ where v_i is replaced by v'_i for each $i \in Y$.

3 Lossy Channel Games

In this section, we provide formal definitions for the game model and the equilibrium concept that we use. While the model definitions are direct generalizations of those in [8], the concurrent setting requires extra care. In this setting, all players choose their actions concurrently and independently. The resulting action profile is evaluated on the game graph, which provides a (distribution of) channel operation to apply and a successor control state. Message losses are then processed.

3.1 Lossy Channel

For simplicity, this article focuses on a single channel system. The *channel configuration* is represented by a word $\mu \in M^*$, where M is a finite alphabet of *messages*. This channel is subject to stochastic *message losses*, meaning that every message has a fixed probability $\lambda \in (0, 1)$ of being lost at every round, independently of other messages.

We can derive the following probability values:

► **Example 4** (Message Losses). For any $\mu, \mu' \in M^*$, let us write $P_\lambda(\mu, \mu')$ for the probability of transitioning from channel configuration μ to μ' after random message losses. For example, $P_\lambda(\mu, \mu) = (1 - \lambda)^{|\mu|}$ (no message loss) and whenever $\mu' \not\leq \mu$ (not a sub-word), we have $P_\lambda(\mu, \mu') = 0$. Moreover, for a single message letter $a \in M$, $P_\lambda(a^n, a^m) = \binom{n}{m} \lambda^m (1 - \lambda)^{n-m}$.

As we will see later in Section 5, the exact value of λ is not relevant for the qualitative probabilistic objectives considered in Definition 10.

3.2 Channel Operations

► **Definition 5.** A channel operation f is defined as one of these three type of partial functions:

- If $f = \text{nop}$ then $f(\mu) = \mu$;
- If $f = !m$ then $f(\mu) = m \cdot \mu$;
- If $f = ?m$ and $\mu = w \cdot m$, then $f(\mu) = w$ and $f(\mu) = \perp \notin M^*$ otherwise.

The set of all such partial functions is denoted op_M .

Intuitively, nop , $!m$, and $?m$ denote “no action”, “enqueue the message m ”, and “dequeue the message m ”, respectively. If the channel configuration does not end with the message m , then the effect of $f = ?m$ is the fresh symbol \perp , indicating that the operation is not allowed.

3.3 Lossy Channel Arena

Players choose channel operations through an *arena*, which specifies a control state (*location*) that defines the actions allowed by the players and the resulting effects on the channel:

► **Definition 6.** A n -player concurrent stochastic lossy channel arena (CSLCG arena) is a tuple $\mathcal{A} = (\text{Agt}, L, M, \{\text{Act}_i\}_{i \in \text{Agt}}, \text{Tab}, l_0)$ where:

- $\text{Agt} = \{1 \dots n\}$ is the set of agents;
- L is the set of locations;
- $l_0 \in L$ is the initial location;
- M is the message alphabet;
- For each $i \in \text{Agt}$, Act_i is a finite set of actions available to agent i and require these sets to be pairwise disjoint. We write Act for $\prod_{i \in \text{Agt}} \text{Act}_i$;
- $\text{Tab} : L \times \text{Act} \rightarrow \text{Dist}(L \times \text{op}_M)$.


A configuration, or state of a CSLCG arena \mathcal{A} is a word $s = l \cdot \mu$ composed of a location and a channel configuration. The state space is denoted by $S = L \cdot M^*$ and the initial state is $s_0 = l_0 \cdot \epsilon \in S$.

In the rest of the section we assume \mathcal{A} to be fixed.

3.4 Concurrent Actions and Strategies

Since actions are taken concurrently, players must be prevented from taking certain actions that could result in illegal channel operations. In general, the set of allowed actions and strategy decisions will depend on the state (location and channel) of the arena, as formalized below:

► **Definition 7 (Allowed Actions and Strategies).** For a configuration $s = l \cdot \mu$ and a player $i \in \text{Agt}$, we define $\text{Act}_i(s)$ as the set of allowed actions α such that if $\exists \vec{\beta} \in \text{Act}^n : \text{Tab}(l, (\vec{\beta}_{-i}, \alpha))[l', f] > 0$, then $f(\mu) \neq \perp$. A strategy for player i is a mapping σ_i from sequences of states (histories), to distributions of allowed actions. Namely, for all $h \cdot s \in S^+$, we have $\sigma_i(h \cdot s) \in \text{Dist}(\text{Act}_i(s))$. We write $\sigma_i(\alpha \mid h \cdot s)$ as a shorthand for $\sigma_i(h \cdot s)[\alpha]$. The set of strategies of i is written \mathfrak{S}_i and $\vec{\mathfrak{S}} = \prod_{i \in \text{Agt}} \mathfrak{S}_i$ is the set of strategy profiles.



■ **Figure 2** Fragment of the semantics of the CSLCG arena of Figure 1 as a concurrent arena. Stochastic behaviours such as message losses are represented by the • nodes and probability values are omitted.

► **Remark 8.** A strategy profile $\vec{\sigma}$ can be seen as a strategy of a single player taking any (distribution of) action(s) $\vec{\alpha} \in \text{Act}$ from history $h \cdot s \in S^+$ with probability

$$\vec{\sigma}(\vec{\alpha} \mid h \cdot s) = \prod_{i \in \text{Agt}} \sigma_i(\alpha_i \mid h \cdot s)$$

However the converse does not always hold: for example if $\text{Act}_i(s) = \{a, b\}$ for $i \in \{1, 2\}$, there exist no pair of strategies σ_1, σ_2 such that $\vec{\sigma}(aa \mid s) = \vec{\sigma}(bb \mid s) = 1/2$ and $\vec{\sigma}(ab \mid s) = \vec{\sigma}(ba \mid s) = 0$.

3.5 Semantics

The semantics of a CSLCG arena can be understood as an n -player, infinite-state, concurrent stochastic game arena $(S, \text{Agt}, \{\text{Act}_i\}_{i \in \text{Agt}}, s_0, p)$ where $p : S \times \text{Act} \rightarrow \text{Dist}(S)$. Intuitively, from any state $s = l \cdot \mu$: (1) Every player picks an allowed action from $\text{Act}_i(s)$ based on the probability given by its strategy on the current history; (2) A new location l' and a channel operation are then sampled according to $\text{Tab}(l, \vec{\alpha})$. Since players only play allowed actions, the new channel configuration $f(\mu)$ is guaranteed to be defined; (3) Message losses finally occur from $f(\mu)$ to some channel configuration μ' , resulting in state $s' = l' \cdot \mu'$. This is formally defined below. An example of such a semantics is provided in Figure 2.

► **Definition 9 (Semantics of CSLCG).** For an initial state $s_0 \in L \cdot M^*$ and a strategy profile $\vec{\sigma} = (\sigma_i)_{i \in \text{Agt}}$ in the CSLCG arena $\mathcal{A} = (\text{Agt}, L, M, \{\text{Act}_i\}_{i \in \text{Agt}}, \text{Tab}, l_0)$, we define $\mathbb{P}_{s_0}^{\vec{\sigma}}$ as the probability measure on $\mathcal{F}(S)$ uniquely defined by its values on cylinders:

- $\mathbb{P}_{s_0}^{\vec{\sigma}}(\text{Cyl}(s_0)) = 1$;
- $\forall s \neq s_0, \mathbb{P}_{s_0}^{\vec{\sigma}}(\text{Cyl}(s)) = 0$;
- For every pair of states $s = l \cdot \mu \in S$ and $s' = l' \cdot \mu' \in S$, and any history $h \in S^*$,

$$\mathbb{P}_{s_0}^{\vec{\sigma}}(\text{Cyl}(h \cdot s \cdot s')) = \mathbb{P}_{s_0}^{\vec{\sigma}}(\text{Cyl}(h \cdot s)) \sum_{\vec{\alpha}, f} \underbrace{\vec{\sigma}(\vec{\alpha} \mid h \cdot s)}_{(1)} \times \underbrace{\text{Tab}(l, \vec{\alpha})[(l', f)]}_{(2)} \times \underbrace{P_\lambda(f(\mu), \mu')}_{(3)}.$$

When $s_0 = l_0 \cdot \epsilon$, we simply write $\mathbb{P}^{\vec{\sigma}} = \mathbb{P}_{s_0}^{\vec{\sigma}}$.

We now define qualitative objectives, that will be used to express players' goals and the (global) properties to be checked:

► **Definition 10.** Let φ be a measurable set of paths. For any strategy profile $\vec{\sigma}$ and state s , we define the property $\text{NZ}(\varphi)$ and $\text{AS}(\varphi)$ by:

$$\mathcal{A}, \vec{\sigma}, s \models \text{NZ}(\varphi) \text{ for } \mathbb{P}_s^{\vec{\sigma}}(\varphi) > 0 \quad \text{and} \quad \mathcal{A}, \vec{\sigma}, s \models \text{AS}(\varphi) \text{ for } \mathbb{P}_s^{\vec{\sigma}}(\varphi) = 1.$$

We omit \mathcal{A} or s when they are clear from context.

We extend the definition to any conjunction of objectives: For a conjunction of NZ and AS objectives $\Psi_1 \wedge \dots \wedge \Psi_k$, we have

$$\mathcal{A}, \vec{\sigma}, s \models \Psi_1 \wedge \dots \wedge \Psi_k \text{ if and only if for all } i \mathcal{A}, \vec{\sigma}, s \models \Psi_i.$$

We consider for φ reachability or safety conditions of regular sets of states, namely LTL with regular valuations [24], and identify a formula φ with its semantics $\mathcal{L}(\varphi) \subseteq (L \cdot M^*)^\omega$. For example, to specify the objective “reaching l_2 almost-surely, while not having more than 3 queued messages with positive probability”, one would write:

$$\Gamma = \text{AS}(\diamond l_2 \cdot M^*) \wedge \text{NZ}(\square L \cdot M^{\leq 3})$$

► **Definition 11.** A CSLCG is defined as an n -player arena together with such properties, called objectives, for every players: $\mathcal{G} = (\mathcal{A}, \Phi_1 \dots \Phi_n)$. In particular, when $n = 2$, and $\Phi_1 = \neg \Phi_2$, we say that \mathcal{G} is a 2.5 player zero-sum game. A strategy $\sigma_i \in \mathfrak{S}_i$ such that $\forall \vec{\sigma}', (\vec{\sigma}'_{-i}, \sigma_i), s \models \Phi$ is called a winning strategy of Φ for i .

► **Definition 12.** For a game \mathcal{G} , strategy profile $\vec{\sigma}$, and state s , we define the set of winners and losers by $W_{\mathcal{G}}(\vec{\sigma}, s) = \{i \in \text{Agt} : (\vec{\sigma}, s) \models \Phi_i\}$ and $L_{\mathcal{G}}(\vec{\sigma}, s) = \text{Agt} \setminus W_{\mathcal{G}}(\vec{\sigma}, s)$. When s is the initial state we simply write $W_{\mathcal{G}}(\vec{\sigma})$ and $L_{\mathcal{G}}(\vec{\sigma})$.

4 Verifying the Core

We consider a cooperative equilibrium concept called the core [14, 41, 30]. Analogous to a Nash equilibrium (NE) [38], a member of the core can be characterized by (the absence of) *beneficial deviations*. However, unlike a Nash equilibrium where only one player can deviate, with the core, a *group* or *coalition* of players can deviate together. The notion of beneficial coalitional deviation is formally defined as follows.

► **Definition 13.** For a strategy profile $\vec{\sigma}$, we say that a joint strategy $\vec{\sigma}'_C, C \subseteq \text{Agt}, C \neq \emptyset$, is a *beneficial coalitional deviation* from $\vec{\sigma}$ if $C \subseteq L_{\mathcal{G}}(\vec{\sigma})$ and for all $\vec{\sigma}'_{-C}$, we have $C \subseteq W_{\mathcal{G}}((\vec{\sigma}'_C, \vec{\sigma}'_{-C}))$.

The core of a game \mathcal{G} is defined to be *the set of strategy profiles that admit no beneficial coalitional deviation*³. We write $\text{Core}(\mathcal{G})$ to denote the set of strategy profiles in the core of \mathcal{G} . We focus on two decision problems related to the core: E-CORE and A-CORE. These problems are formally defined below.

Given: (\mathcal{G}, Γ) with game \mathcal{G} and property Γ .


E-CORE: Does there exists some $\vec{\sigma} \in \text{Core}(\mathcal{G})$ such that $\mathcal{G}, \vec{\sigma} \models \Gamma$?

A-CORE: Is it the case that for all $\vec{\sigma} \in \text{Core}(\mathcal{G})$ we have $\mathcal{G}, \vec{\sigma} \models \Gamma$?

Note that, due to the duality of these problems, it is enough to provide a procedure for E-CORE.

► **Example 14.** Consider a transmission system with one channel similar to the one discussed in Example 1. Now, suppose there are three players, S (sender), R (receiver), and A (attacker). S/R decides to send/request a message, a or b . A message is delivered successfully if R reads

³ An alert reader may notice the similarities between the core and *strong NE* [13] and *coalition-proof NE* [18]. The main difference is that these non-cooperative equilibrium concepts do not assume the existence of binding agreements. We refer to [29] for a more detailed discussion on this matter.



■ **Figure 3** A graphical representation of a 3-player arena. Joint actions are ordered by *SRA*, for example, xyz indicates that players S, R, A choose actions x, y, z respectively.

(dequeues) the same type of message as she requested. An attack is successful if A chooses the same type of message being sent, and this turns the message into a c . We model this game as a CSLCG where $\text{Act}_S = \text{Act}_A = \{a, b, -\}$, $\text{Act}_R = \{a, b, \text{dequeue}, \text{reset}\}$. The arena is depicted in Figure 3. The goals of S and R are given as $\Phi_S = \text{AS}(\Box(L \cdot M^{\leq k}))$ for some $k \in \mathbb{N}$ (almost surely the channel never exceeds size k), and $\Phi_R = \text{AS}(\Diamond(l_F \cdot M^*))$ (almost surely a correct message is delivered). The goal of A is $\Phi_A = \neg(\Phi_S \wedge \Phi_R)$.

Consider the following strategy profile: if the channel contains fewer than k messages, S and R play aa and bb uniformly at random. Otherwise, S plays action $-$. This profile satisfies both Φ_S and Φ_R , and is therefore in the core. In fact, *all* strategy profiles in the core satisfy both Φ_S and Φ_R . As a result, if we query E-CORE with $\Gamma = \Phi_A = \neg(\Phi_S \wedge \Phi_R)$, it returns a negative answer. This is in contrast to NE: since with NE S and R do not act as a coalition, there exists a (undesirable) NE in which there is a non-zero probability that the channel will exceed size k or a correct message will never be delivered (i.e., Γ is satisfied). \square

To address E-CORE, we first introduce the notion of *concurrent two-player coalition game* (TPCG) as follows.

► **Definition 15.** Let $\mathcal{G} = (\mathcal{A}, (\Phi_i)_{i \in \text{Agt}})$ be a CSLCG and $C \subseteq \text{Agt}$, with the underlying arena $\mathcal{A} = (\text{Agt}, L, M, \{\text{Act}_i\}_{i \in \text{Agt}}, \text{Tab}, l_0)$. The concurrent two-player coalition game arena is defined as $\mathcal{A}^C = ((1, 2), L, M, \{\text{Act}_i\}_{i \in \text{Agt}}, \text{Tab}', l_0)$ where for all actions $\vec{\alpha}, \vec{\beta} \in \text{Act}$, the transition is determined by the projections on C and $-C = \text{Agt} \setminus C$, respectively for player 1 and 2: $\text{Tab}'(l, (\vec{\alpha}, \vec{\beta})) = \text{Tab}(l, (\vec{\alpha}_C, \vec{\beta}_{-C}))$. The TPCG with respect to \mathcal{G}, C , and objective Ψ_C is thus defined as $\mathcal{G}^{C, \Psi_C} = (\mathcal{A}^C, (\Psi_C, \neg\Psi_C))$.

Observe that the ability of a coalition C to satisfy an objective Ψ_C depends on whether it has a winning strategy in TPCG \mathcal{G}^{C, Ψ_C} from the initial state of the game (thus, the existence of beneficial deviation). With this observation, we restate the characterization of E-CORE from [26] as follows.

► **Proposition 16** ([26]). A pair (\mathcal{G}, Γ) is a yes-instance of E-CORE if and only if there exists $W \subseteq \text{Agt}$ such that

- (a) there exists some $\vec{\sigma}$ such that $\mathcal{G}, \vec{\sigma} \models \Phi_W$, and
 - (b) for all $C \subseteq \text{Agt} \setminus W$, C has no winning strategy in \mathcal{G}^{C, Ψ_C} ,
- where $\Phi_W = \bigwedge_{i \in W} \Phi_i \wedge \bigwedge_{i \notin W} \neg \Phi_i \wedge \Gamma$ and $\Psi_C = \bigwedge_{i \in C} \Phi_i$.

Using Proposition 16, we provide the following procedure for determining whether some (\mathcal{G}, Γ) is a yes-instance of E-CORE.

1. Guess a set of winning players $W \subseteq \text{Agt}$;
2. Check if there is a winning strategy $\vec{\sigma}$ in TPCG $\mathcal{G}^{\text{Agt}, \Phi_W} = (\mathcal{A}, (\Phi_W, \neg\Phi_W))$;
3. Check if there is a coalition $C \subseteq \text{Agt} \setminus W$ with a winning strategy $\vec{\sigma}'_C$ in TPCG $\mathcal{G}^{C, \Psi_C} = (\mathcal{A}^C, (\Psi_C, \neg\Psi_C))$;
4. If the answers to Step 2 is “Yes” and Step 3 is “No”, then (\mathcal{G}, Γ) is a yes-instance of E-CORE. Otherwise, it is not.

Observe that above procedure corresponds to Proposition 16. In particular, Steps 2 and 3 correspond precisely to (a) and (b) in Proposition 16, respectively. Thanks to this procedure, the problem of checking whether (\mathcal{G}, Γ) is a yes-instance of E-CORE can be reduced to solving a collection of concurrent 2.5-player zero-sum games.

Note that if we consider almost-sure LTL objectives, we can construct a CSLCG \mathcal{G} with one player whose goal is $\Phi = \text{AS}((\Box\Diamond R_1) \wedge (\Box\Diamond R_2))$. Then, we can set $\Gamma = \Phi$ and query E-CORE. This reduces to a 1.5-player game over a lossy channel with objective Φ , which is already undecidable [16, Lemma 5.12]. Therefore, we obtain the following result.

► **Proposition 17.** *For a pair (\mathcal{G}, Γ) where players’ objectives and Γ are given by almost-sure LTL formulae, the problems of E-CORE and A-CORE are undecidable.*

In the following sections, we study the technical machineries required to solve concurrent 2.5-player zero-sum games with almost-sure safety and almost-sure reachability objectives. As motivated by Proposition 16 and its corresponding procedure for E-CORE, we further solve concurrent 2.5-player zero-sum games with a conjunction of objectives. These provide the necessary foundation for the main decidability result presented later in Section 7.

5 The Zero-Sum case is Effective

In this section, we focus on solving concurrent 2.5-player zero-sum CSLCG as a crucial step towards our main decidability result for E-CORE. We assume that the CSLCG \mathcal{G} with arena \mathcal{A} is fixed, and R represents a set of states, which may be infinite. While [23] provides an approach for solving concurrent stochastic games with $\text{AS}(\Diamond R)$, $\text{AS}(\Box R)$, $\text{NZ}(\Diamond R)$ and $\text{NZ}(\Box R)$ objectives in the *finite-state* setting, here we extend their approach to the *infinite-state* setting. We note that the infinite-state setting has been previously examined for *turn-based* games in [19, 5, 15].

Our approach can be summarized as follows: First, we provide algorithms to symbolically compute one-step reachability for regular set of states. Then, we prove that the algorithms from [23] remain valid (in terms of termination and correctness) for this class. More precisely, we fix a regular set $R \subseteq (L \cdot M^*)^*$ and a 2.5-player zero-sum CSLCG \mathcal{G} with $\text{Agt} = \{1, 2\}$. For $i \in \text{Agt}$, the objective Φ_i is of the form $\text{NZ}(\Diamond R)$ or $\text{AS}(\Diamond R)$. Furthermore, by instantiating games for the opponent $-i$, setting $R' = S \setminus R$, and by determinacy of such games [35], we also solve the game for objectives of the form $\text{AS}(\Box R')$ or $\text{NZ}(\Box R')$. To this end, we focus on the existence of winning strategies (Definition 11) and the computation of *winning regions*.

► **Definition 18.** *For an objective Φ of player i , the winning region of Φ for player i in \mathcal{G} is given by: $\llbracket \Phi \rrbracket_i(\mathcal{G}) = \{s \in S \mid \exists \sigma_i \in \mathfrak{S}_i, \forall \vec{\sigma}' \in \vec{\mathfrak{S}}, (\vec{\sigma}'_{-i}, \sigma_i), s \models \Phi\}$.*

As previously mentioned, moving from a finite to an infinite state setting presents a challenge in the representation of winning strategies. To address this, in what follows, we provide a novel encoding of different classes of winning strategies necessary for ensuring a win.

Figure 4 illustrates how a FM strategy can be represented finitely. It determines which action to play by running a finite number of automata on the history, reading every location and channel configuration. In contrast, a strategy that depends only on the last state may not be representable in the infinite state case, as infinitely many different decisions might be taken. Therefore, we avoid using the usually synonymous term *memoryless* and instead refer to these strategies as *positional*. A convenient representation is therefore the positional *and* finite memory (PFM) strategies, meaning that the action depends only on the last state, which must belong to one of finitely many regular sets.

5.2 Positive Reachability

In the positive reachability case, player i tries to enforce that *some* finite prefix reaches the target set R . This can be achieved by a backward-reachability algorithm instantiated on well-quasi-ordered sets [6]:

► **Lemma 22.** *Given a regular set $R \subseteq S^*$, the algorithm that computes $\bigcup_{k \geq 0} \text{Pre}_i^k(R)$ converges in a finite number of steps and returns a set $R \cup V$, where V is upward-closed and $R \cup V = \llbracket \text{NZ}(\diamond R) \rrbracket_i$. Moreover, PFM winning strategies are sufficient for both players.*

► **Example 23.** Consider the example from Figure 1. A winning strategy for player 2/Attacker to achieve the objective $\text{NZ}(\diamond l_2 \cdot c)$ (eventually reaching state $l_2 \cdot c$ with positive probability) consists in playing $\mathcal{U}(\{a, b\})$ from any state in $l_0 \cdot M^*$, then b from any state in $l_1 \cdot M^* \cdot c^2$. With positive probability, the state $l_1 \cdot cc$ is reached after 3 steps, regardless of Sender’s strategy, and then $l_2 \cdot c$ is reached.

This algorithm can also be used to compute almost-sure safety of R by applying the determinacy result: $\llbracket \text{AS}(\square R) \rrbracket_i = \llbracket \text{NZ}(\overline{\diamond R}) \rrbracket_{-i}$. As shown in Lemma 22, the optimal strategy is PFM, and in the safety case, it can be further restricted without compromising the safety property. More precisely, we define the most general action restriction for player i that preserves safety as follows:

► **Definition 24.** *For any set $R \subseteq S$, $\text{Stay}_i^{\mathcal{G}}(R)$ is the mapping that restricts the allowed actions on R to stay in R :*

$$s \in R \mapsto \{\alpha \in \text{Act}_i(s) \mid \forall \vec{\sigma}, \mathbb{P}^{(\vec{\sigma}-i, \alpha)}(\bigcirc R) = 1\}$$

5.3 Almost-Sure Reachability

In contrast to the turn-based case in [8], concurrent actions require a careful analysis of the allowed actions. The intuition is as follows: As the player with reachability objective tries to avoid “being trapped” in bad states, he may choose not to play some actions based on the previously defined Stay operator (Definition 24). This limits the available actions and gives more power to his opponent, which then reduces the winning region. This idea was proposed in the algorithm for almost-sure reachability in finite state games described in [23]. We follow here this approach by implementing the algorithm symbolically as depicted in Algorithm 1 which boils down to providing effective procedures for Pre_i and Stay_i , when the input is a regular set of states. As the algorithm in [23] actually solves almost-sure *repeated* reachability, the authors assumed that the target set R is *absorbing*, meaning that once a state in R is reached, the game cannot exit R , regardless of the players’ actions. This assumption is also made here as adjusting the Pre_i operator is sufficient to guarantee this property in our setting. Termination of the algorithm can be shown using the wqo property, similar to that described in [8], resulting in the following lemma:

■ **Algorithm 1** Almost-Sure Reachability.

Input: An arena \mathcal{A} , $i \in \text{Agt}$ and a regular set R

Output: $D_k = \llbracket \text{AS}(\Box \Diamond R) \rrbracket_i$

$k \leftarrow 0$; $D_0 \leftarrow S$; $\mathcal{A}_0 \leftarrow \mathcal{A}$

repeat

$C_k = \llbracket \text{AS}(\Box(D_k \setminus R)) \rrbracket_{-i}(\mathcal{A}_k)$

$Y_k = \llbracket \text{NZ}(\Diamond R) \rrbracket_i(\mathcal{A}_k)$

$D_{k+1} = \llbracket \text{AS}(\Box(Y_k)) \rrbracket_i(\mathcal{A}_k)$

$\mathcal{A}_{k+1} \leftarrow \mathcal{A}_k$ where $\text{Act}_i^{\mathcal{A}_{k+1}} = \text{Stay}_i^{\mathcal{A}_k}(D_{k+1})$

until $D_k = D_{k+1}$

► **Lemma 25.** *The almost-sure reachability algorithm of [23], instantiated symbolically on 2.5-player CSLCG in Algorithm 1, terminates and returns a downward-closed set.*

To prove the correctness of Algorithm 1, we must provide a winning strategy from every state in the computed set D_k , and a strategy for the opponent from every state in the complement set $\overline{D_k}$. At this point, it is important to note that positional strategies may not be sufficient, especially for the opponent player. As illustrated by the *Hide-or-Run game* presented in [23], non-positional may be needed. More precisely, the authors noticed that Markov strategies –where decisions depend on the current state and on the clock value only– are sufficient for winning with positive safety objectives. They further introduced the sufficient subclass of so-called *counting strategies*, where a sequence of probability values is fixed, so that the strategy can be finitely represented. We generalise this notion of counting strategies to the infinite state case as follows:

► **Definition 26.** *For any k , let $p_k = 2^{-1/(2^k)}$. A strategy $\sigma_i \in \mathfrak{S}_i$ is counting (C) if there exist two PFM strategies $\sigma_i^v, \sigma_i^u \in \mathfrak{S}_i$ such that for every $k \in \mathbb{N}$, $h \cdot s \in S^k$ and any $\alpha \in \text{Act}_i(s)$,*

$$\sigma_i(\alpha \mid h \cdot s) = p_k \sigma_i^u(\alpha \mid h \cdot s) + (1 - p_k) \sigma_i^v(\alpha \mid h \cdot s)$$

Note that p_k is fixed a sequence of reals between 0 and 1, such that the infinite product $\prod_{i=1}^{\infty} p_i = p$ is between 0 and 1. This means that the strategy σ_i^u always has some positive probability of being played, but overall cannot be played forever. Since the sequence $(p_k)_k$ is fixed, a counting strategy requires infinite memory but can be finitely represented by two PFM strategies.

► **Example 27.** Consider again the example shown in Figure 1. We provide two examples of winning strategies in the AS case:

- If $\Phi_1 = \text{AS}(\Diamond(l_1 \cdot \{a\}^3 \cup l_2 \cdot M^*))$ –namely Sender can almost-surely eventually force a valid transmission of 3 consecutive a 's, assuming that the game continues forever– Sender has a winning strategy by playing $\mathcal{U}(\{a, b\})$ from all states. For any strategy σ_2 of Attacker, either the game eventually reaches l_2 , or there is a state $s = l \cdot \mu \in \{l_0, l_1\} \cdot M^*$ visited infinitely often. From this state, there is a fixed probability $p > 0$ to produce three consecutive a 's and that all b 's are dropped, so this event eventually happens almost-surely.
- If $\Phi_1 = \text{AS}(\Diamond L \cdot M^* \{a\} M^*)$ –namely a message a is eventually sent almost-surely– we argue that Sender cannot achieve her objective. To observe this, one can exhibit a winning strategy for Attacker, whose objective is then $\Phi_2 = \text{NZ}(\Box L \cdot \{b, c\}^*)$. Such a strategy consists of the counting strategy playing action w with probability p_k at round k and

$\mathcal{U}(\{a, b\})$ otherwise. At any round still in l_0 , Sender cannot risk playing a since there is a small probability for Attacker to scramble the communication and then reach l_2 . The overall probability to stay in $l_0 \cdot b^*$ is therefore $\prod_k p_k > 0$.

On the other hand, any PFM strategy by Attacker can be defeated, which proves that counting strategies are required. Indeed by playing b from all states in $l_0 \cdot b^*$, Sender ensures that she will never lose since either σ_A eventually plays w with probability 1, and Sender can then play a and win, or there is a fixed positive probability (FM) that the game moves to $l_1 \cdot b^*$, which happens almost-surely.

This allows us to conclude on the almost-sure reachability case by adapting the correctness proof of the almost-sure reachability algorithm in [23] mentioned in Lemma 25. As a matter of fact, the finite attractor property seen in Proposition 20 allows us to refer back to the finite state case and derive sufficient strategies for both players. We conclude this subsection with the following:

► **Lemma 28.** *Using Algorithm 1, one can compute:*

- *The winning set $W = \llbracket \text{AS}(\diamond R) \rrbracket_i$;*
- *A PFM winning strategy for i , $\sigma_i : h \cdot s \in S^* \cdot W \mapsto \mathcal{U}(\text{Stay}_i(W))$;*
- *A counting (C) winning strategy for his opponent $-i$ (objective $\text{NZ}(\square \bar{R})$).*

6 Conjunction of Objectives

In this section, we stay in the 2.5-player zero-sum setting, and address the computation of winning regions for objectives composed as a *conjunction* of qualitative objectives. Note that contrary to the previous section, games are in general not determined for such conjunctive objectives [43], so we focus now on the winning strategies for the player whose objective is a conjunction of objectives.

As discussed in Section 5, positive reachability and almost-sure safety/reachability can be achieved using PFM strategies, while positive safety may require the use of counting strategies. The rest of this section is dedicated to proving the following theorem by combining PFM and counting strategy classes.

► **Theorem 29.** *Let Φ be a conjunction of NZ and AS objectives for safety and reachability path specifications. Then the winning region $\llbracket \Phi \rrbracket_i(\mathcal{G})$ is computable.*

We first notice that conjunction of objectives usually enjoys some convexity properties, namely randomization between individual optimal strategies is sometimes sufficient to achieve the conjunctive objective. This is the case when the conjunction Φ does not contain a positive safety objective. Following this observation, we get:

► **Lemma 30.** *If Φ is a conjunctive objective containing only $\text{NZ}(\diamond \cdot)$, $\text{AS}(\square \cdot)$ and $\text{AS}(\diamond \cdot)$ objectives of regular sets, then $\llbracket \Phi \rrbracket_i(\mathcal{G})$ is computable, and there exists a winning FM strategy for player i .*

Sketch.

- If all reachability objective sets are absorbing –that is to say cannot be left once satisfied– we prove that a PFM winning strategy exists for player i .

For every objective $\Psi = \text{AS}(\square R)$ or $\Psi = \text{AS}(\diamond R)$ appearing in Φ , we restrict the set of available actions to $\text{Stay}_i(\llbracket \Psi \rrbracket_i)$, as a winning strategy necessarily plays actions from this set only. States where no actions are available are removed. Finally, we solve the game on the restricted arena for each $\text{NZ}(\diamond \cdot)$ objective, and take the intersection of winning regions. A sufficient winning strategy consists in playing all actions uniformly at random.

- In the general case, we build a new game \mathcal{G}' keeping track of the already satisfied reachability objectives. Each reachability objective can therefore be rewritten to be absorbing by referring to the extra information bit stored in every state. A PFM winning strategy in \mathcal{G}' is then translated back to a FM strategy in \mathcal{G} , at the expense of a 1-bit memory per objective. ◀

The treatment of positive safety objectives requires more attention, as winning strategies may require infinite memory, as seen in Lemma 28. We notice however that any positive safety objective can be replaced by a simple positive reachability property as it is sufficient for a winning strategy to wait for all almost-sure reachability objectives to be met, then play the counting strategy for $\text{NZ}(\square\cdot)$. It is therefore sufficient and necessary, for at least one state in $\llbracket \text{NZ}(\square\cdot) \rrbracket_i$ to be reachable with positive probability. This is summarized by the final lemma below:

► **Lemma 31** (From $\text{NZ}(\square)$ to $\text{NZ}(\diamond)$). *If $\Phi \equiv \Theta \wedge \text{NZ}(\square\bar{R})$, then $\llbracket \Phi \rrbracket_i = \llbracket \Theta \wedge \text{NZ}(\diamond R') \rrbracket_i$ where Θ is some conjunctive condition (as defined in Definition 10) and*

$$R' = \llbracket \text{NZ}(\square\bar{R}) \rrbracket_i \cap \bigcap_{\text{AS}(\diamond R'') \in \Phi} R''.$$

Theorem 29 is then obtained by applying Lemma 31 for every $\text{NZ}(\square\cdot)$ objective, then applying Lemma 30.

7 Decidability of E-Core and A-Core

In this section, we present our main decidability result. We begin by recalling that E-CORE and A-CORE are undecidable when players' goals are given by almost-sure LTL formulae (Proposition 17). To obtain decidability, we focus on two types of objectives: almost-sure reachability and almost-sure safety objectives. Consider a CSLCG $\mathcal{G} = (\mathcal{A}, (\Phi_i)_{i \in \text{Agt}})$ in which for each $\Phi_i = \text{AS}(\varphi_i)$, either $\varphi_i \equiv \diamond R_i$ or $\varphi_i \equiv \square R_i$. Then the formula used in Step 2 of the procedure for solving E-CORE in Section 4 is given as:

$$\Phi_W = \bigwedge_{i \in W} \text{AS}(\varphi_i) \wedge \bigwedge_{i \notin W} \text{NZ}(\neg\varphi_i) \wedge \Gamma.$$

The formula used in Step 3 is given as $\Psi_C = \bigwedge_{i \in C} \text{AS}(\varphi_i)$.

Observe that the TPCG $\mathcal{G}^{\text{Agt}, \Phi_W}$ in Step 2 is, in fact, a 1.5-player game, since player 1 consists of all players in the original game. Thus, solving Step 2 amounts to finding a scheduler that satisfies Φ_W . This problem is decidable if $\Gamma = \text{AS}(\varphi)$ and φ is of the form $\bigwedge_i \diamond R_i$, $\bigwedge_i \square R_i$, or $\bigwedge_i \square \diamond R_i$ [16]. Furthermore, if such a scheduler exists, then there is one that is deterministic and has finite memory. To solve Step 3, we reduce it to solving, for each $C \subseteq \text{Agt}$, a 2.5-player zero-sum game in which the goal of player 1 is Ψ_C . As shown in Theorem 29, this problem is decidable. Therefore, we obtain the following theorem.

► **Theorem 32.** *For a pair (\mathcal{G}, Γ) where players' objectives are almost-sure reachability or almost-sure safety objectives, and property $\Gamma = \text{AS}(\varphi)$ with φ of the form $\bigwedge_i \diamond R_i$, $\bigwedge_i \square R_i$, or $\bigwedge_i \square \diamond R_i$, the problems of E-CORE and A-CORE are decidable.*

8 Concluding Remarks

This paper presents the first result on rational verification of concurrent stochastic games on infinite arenas, specifically for lossy channel games. Our focus is on the decidability of verifying the core of multi-player concurrent stochastic lossy channel games. To this end, we

provide an approach for solving concurrent 2.5-player zero-sum games with a conjunction of almost-sure safety and almost-sure reachability objectives, and infinite state arenas. Our approach extends previous methods, which were limited to either finite state arenas or turn-based games, to work on concurrent settings and infinite state arenas.

It is worth mentioning that most of the algorithms described in this work have non-elementary complexity in the worst case, as they address problems that can be reduced to the reachability problem in lossy channel systems [42]. One might wonder if it is possible to trade stochasticity for non-determinism in the model to obtain an easier problem. Surprisingly, although the winning region for non-deterministic reachability (i.e., equivalent to $NZ(\diamond)$) is computable, the winning region for non-deterministic safety (“does there exist an infinite path . . .”) is non-computable, due to undecidability results by [36].

Finally, an obvious direction for future work is to consider the NE concept and its related variants, such as strong NE and coalition-proof NE. However, addressing questions related to NE in concurrent games with probabilistic qualitative objectives is challenging, even in the finite state case. There are two main challenges: Firstly, because NE are strategy profiles where players do not behave as a coalition, limiting their ability to prevent deviations (see Remark 8), finding the appropriate joint strategies is not straightforward. Secondly, encoding stability against any deviation in the NE setting produces more complex conjunctions of objectives, requiring significant extensions of the techniques presented in Section 6. Reductions from concurrent to turn-based two-players games, have been used in previous work. For instance, in *suspect games* [20] one player proposes a NE while the second attempts to disprove it. Another approach is by sequentializing games in order to compute *punishing strategies/regions* to characterize NE [28, 31]. However, these approaches encounter similar challenges above when stochastic behaviours are introduced. We conjecture that in order to capture the NE concept in its generality, the reduction to 2.5-player games should feature concurrent actions to account for randomized actions and even counting strategies in the NE.

References

- 1 Alessandro Abate, Julian Gutierrez, Lewis Hammond, Paul Harrenstein, Marta Kwiatkowska, Muhammad Najib, Giuseppe Perelli, Thomas Steeples, and Michael Wooldridge. Rational verification: game-theoretic verification of multi-agent systems. *Applied Intelligence*, 51(9):6569–6584, 2021.
- 2 Parosh Aziz Abdulla, C. Aiswarya, and Mohamed Faouzi Atig. Data communicating processes with unreliable channels. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 166–175. ACM, 2016. doi:10.1145/2933575.2934535.
- 3 Parosh Aziz Abdulla, Mohamed Faouzi Atig, Raj Aryan Agarwal, Adwait Godbole, and Shankara Narayanan Krishna. Probabilistic total store ordering. In Ilya Sergey, editor, *Programming Languages and Systems – 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings*, volume 13240 of *Lecture Notes in Computer Science*, pages 317–345. Springer, 2022. doi:10.1007/978-3-030-99336-8_12.
- 4 Parosh Aziz Abdulla, Mohamed Faouzi Atig, Ahmed Bouajjani, and Tuan Phong Ngo. A load-buffer semantics for total store ordering. *Logical Methods in Computer Science*, 14(1), 2018. doi:10.23638/LMCS-14(1:9)2018.

- 5 Parosh Aziz Abdulla, Nathalie Bertrand, Alexander Moshe Rabinovich, and Philippe Schnoebelen. Verification of probabilistic systems with faulty communication. *Inf. Comput.*, 202(2):141–165, 2005. doi:10.1016/j.ic.2005.05.008.
- 6 Parosh Aziz Abdulla, Karlis Cerans, Bengt Jonsson, and Yih-Kuen Tsay. General decidability theorems for infinite-state systems. In *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27-30, 1996*, pages 313–321. IEEE Computer Society, 1996. doi:10.1109/LICS.1996.561359.
- 7 Parosh Aziz Abdulla, Lorenzo Clemente, Richard Mayr, and Sven Sandberg. Stochastic parity games on lossy channel systems. *Log. Methods Comput. Sci.*, 10(4), 2014. doi:10.2168/LMCS-10(4:21)2014.
- 8 Parosh Aziz Abdulla, Noomene Ben Henda, Luca de Alfaro, Richard Mayr, and Sven Sandberg. Stochastic games with lossy channels. In Roberto M. Amadio, editor, *Foundations of Software Science and Computational Structures, 11th International Conference, FOSSACS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29 – April 6, 2008. Proceedings*, volume 4962 of *Lecture Notes in Computer Science*, pages 35–49. Springer, 2008. doi:10.1007/978-3-540-78499-9_4.
- 9 Parosh Aziz Abdulla and Bengt Jonsson. Verifying programs with unreliable channels. In *Proceedings of the Eighth Annual Symposium on Logic in Computer Science (LICS '93), Montreal, Canada, June 19-23, 1993*, pages 160–170. IEEE Computer Society, 1993. doi:10.1109/LICS.1993.287591.
- 10 Parosh Aziz Abdulla and Bengt Jonsson. Model checking of systems with many identical timed processes. *Theor. Comput. Sci.*, 290(1):241–264, 2003. doi:10.1016/S0304-3975(01)00330-9.
- 11 Benjamin Aminof, Marta Kwiatkowska, Bastien Maubert, Aniello Murano, and Sasha Rubin. Probabilistic strategy logic. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 32–38, 2019.
- 12 Benjamin Aminof, Axel Legay, Aniello Murano, Olivier Serre, and Moshe Y Vardi. Pushdown module checking with imperfect information. *Information and Computation*, 223:1–17, 2013.
- 13 Robert J Aumann. Acceptable points in general cooperative n-person games. *Contributions to the Theory of Games*, 4(40):287–324, 1959.
- 14 Robert J. Aumann. The core of a cooperative game without side payments. *Transactions of the American Mathematical Society*, 98(3):539–552, 1961. URL: <http://www.jstor.org/stable/1993348>.
- 15 Christel Baier, Nathalie Bertrand, and Philippe Schnoebelen. A note on the attractor-property of infinite-state Markov chains. *Inf. Process. Lett.*, 97(2):58–63, 2006. doi:10.1016/j.ipl.2005.09.011.
- 16 Christel Baier, Nathalie Bertrand, and Philippe Schnoebelen. Verifying nondeterministic probabilistic channel systems against ω -regular linear-time properties. *ACM Trans. Comput. Log.*, 9(1):5, 2007. doi:10.1145/1297658.1297663.
- 17 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 18 B Douglas Bernheim, Bezalel Peleg, and Michael D Whinston. Coalition-proof Nash equilibria i. concepts. *Journal of economic theory*, 42(1):1–12, 1987.
- 19 Nathalie Bertrand and Philippe Schnoebelen. Model checking lossy channels systems is probably decidable. In Andrew D. Gordon, editor, *Foundations of Software Science and Computational Structures, 6th International Conference, FOSSACS 2003 Held as Part of the Joint European Conference on Theory and Practice of Software, ETAPS 2003, Warsaw, Poland, April 7-11, 2003, Proceedings*, volume 2620 of *Lecture Notes in Computer Science*, pages 120–135. Springer, 2003. doi:10.1007/3-540-36576-1_8.
- 20 Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Pure Nash equilibria in concurrent deterministic games. *Log. Methods Comput. Sci.*, 11(2), 2015. doi:10.2168/LMCS-11(2:9)2015.

- 21 Laura Bozzelli, Aniello Murano, and Adriano Peron. Module checking of pushdown multi-agent systems. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 17, pages 162–171, 2020.
- 22 Taolue Chen, Fu Song, and Zhilin Wu. Global model checking on pushdown multi-agent systems. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 2459–2465. AAAI Press, 2016. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11984>, doi:10.1609/AAAI.V30I1.10124.
- 23 Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. Concurrent reachability games. *Theor. Comput. Sci.*, 386(3):188–217, 2007. doi:10.1016/j.tcs.2007.07.008.
- 24 Javier Esparza, Antonín Kucera, and Stefan Schwoon. Model checking LTL with regular valuations for pushdown systems. *Inf. Comput.*, 186(2):355–376, 2003. doi:10.1016/S0890-5401(03)00139-1.
- 25 Alain Finkel and Philippe Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1-2):63–92, 2001. doi:10.1016/S0304-3975(00)00102-X.
- 26 Julian Gutierrez, Lewis Hammond, Anthony W. Lin, Muhammad Najib, and Michael J. Wooldridge. Rational verification for probabilistic systems. In Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, editors, *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, pages 312–322, 2021. doi:10.24963/kr.2021/30.
- 27 Julian Gutierrez, Paul Harrenstein, and Michael Wooldridge. From model checking to equilibrium checking: Reactive modules for rational verification. *Artificial Intelligence*, 248:123–157, 2017.
- 28 Julian Gutierrez, Paul Harrenstein, and Michael J. Wooldridge. Expressiveness and complexity results for strategic reasoning. In Luca Aceto and David de Frutos-Escrig, editors, *26th International Conference on Concurrency Theory, CONCUR 2015, Madrid, Spain, September 1-4, 2015*, volume 42 of *LIPICs*, pages 268–282. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.CONCUR.2015.268.
- 29 Julian Gutierrez, Szymon Kowara, Sarit Kraus, Thomas Steeples, and Michael Wooldridge. Cooperative concurrent games. *Artificial Intelligence*, 314:103806, 2023.
- 30 Julian Gutierrez, Sarit Kraus, and Michael J. Wooldridge. Cooperative concurrent games. In Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor, editors, *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, pages 1198–1206. International Foundation for Autonomous Agents and Multiagent Systems, 2019. URL: <http://dl.acm.org/citation.cfm?id=3331822>.
- 31 Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael J. Wooldridge. Automated temporal equilibrium analysis: Verification and synthesis of multi-player games. *Artif. Intell.*, 287:103353, 2020. doi:10.1016/j.artint.2020.103353.
- 32 Graham Higman. Ordering by divisibility in abstract algebras. In *Proc. London Math. Soc.*, volume 2, pages 326–336, 1952.
- 33 John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation, 2nd Edition*. Addison-Wesley series in computer science. Addison-Wesley-Longman, 2001.
- 34 Ori Lahav, Egor Namakonov, Jonas Oberhauser, Anton Podkopaev, and Viktor Vafeiadis. Making weak memory models fair. *Proc. ACM Program. Lang.*, 5(OOPSLA):1–27, 2021. doi:10.1145/3485475.
- 35 Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975. URL: <http://www.jstor.org/stable/1971035>.
- 36 Richard Mayr. Undecidable problems in unreliable computations. *Theor. Comput. Sci.*, 297(1-3):337–354, 2003. doi:10.1016/S0304-3975(02)00646-1.

- 37 Aniello Murano and Giuseppe Perelli. Pushdown multi-agent system verification. In Qiang Yang and Michael J. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 1090–1097. AAAI Press, 2015. URL: <http://ijcai.org/Abstract/15/158>.
- 38 John F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(1):48–49, 1950.
- 39 Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October – 1 November 1977*, pages 46–57. IEEE Computer Society, 1977. doi:10.1109/SFCS.1977.32.
- 40 Halsey L. Royden and Patrick M. Fitzpatrick. *Real Analysis*. Prentice Hall, 2010. URL: <https://books.google.fr/books?id=H65bQgAACAAJ>.
- 41 Herbert E Scarf. The core of an n person game. *Econometrica: Journal of the Econometric Society*, pages 50–69, 1967.
- 42 Philippe Schnoebelen. Verifying lossy channel systems has nonprimitive recursive complexity. *Inf. Process. Lett.*, 83(5):251–261, 2002. doi:10.1016/S0020-0190(01)00337-4.
- 43 Tobias Winkler and Maximilian Weininger. Stochastic games with disjunctions of multiple objectives. In Pierre Ganty and Davide Bresolin, editors, *Proceedings 12th International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2021, Padua, Italy, 20-22 September 2021*, volume 346 of *EPTCS*, pages 83–100, 2021. doi:10.4204/EPTCS.346.6.