Stochastic Timed Games Revisited

S. Akshay[*][1]**, Patricia Bouyer**[†][2]**, Shankara Narayanan Krishna** [‡][3]**,
Lakshmi Manasa**[4]**, and Ashutosh Trivedi**[5]

1   **Department of Computer Science & Engineering, IIT Bombay, Bombay, India**
      `akshayss@cse.iitb.ac.in`
2   **LSV, CNRS & ENS Cachan, Université Paris-Saclay, Paris, France**
      `bouyer@lsv.fr`
3   **Department of Computer Science & Engineering, IIT Bombay, Bombay, India**
      `krishnas@cse.iitb.ac.in`
4   **Department of Computer Science & Engineering, IIT Bombay, Bombay, India**
      `manasa@cse.iitb.ac.in`
5   **University of Colorado, Boulder, USA**
      `ashutosh.trivedi@colorado.edu`

Abstract

Stochastic timed games (STGs), introduced by Bouyer and Forejt, naturally generalize both continuous-time Markov chains and timed automata by providing a partition of the locations between those controlled by two players (Player Box and Player Diamond) with competing objectives and those governed by stochastic laws. Depending on the number of players – 2, 1, or 0 – subclasses of stochastic timed games are often classified as $2\frac{1}{2}$-player, $1\frac{1}{2}$-player, and $\frac{1}{2}$-player games where the $\frac{1}{2}$ symbolizes the presence of the stochastic "nature" player. For STGs with reachability objectives it is known that $1\frac{1}{2}$-player one-clock STGs are decidable for qualitative objectives, and that $2\frac{1}{2}$-player three-clock STGs are undecidable for quantitative reachability objectives. This paper further refines the gap in this decidability spectrum. We show that quantitative reachability objectives are already undecidable for $1\frac{1}{2}$ player four-clock STGs, and even under the time-bounded restriction for $2\frac{1}{2}$-player five-clock STGs. We also obtain a class of $1\frac{1}{2}$, $2\frac{1}{2}$ player STGs for which the quantitative reachability problem is decidable.

1   Introduction

Two-player zero-sum games over finite state-transition graphs are a natural framework for controller synthesis for discrete event systems. In this setting two players – say Player Box and Player Diamond (after necessity and possibility operators) – represent the controller and the environment, and control-program synthesis corresponds to finding a winning (or optimal) strategy of the controller for some given performance objective. Finite graphs, however, often do not satisfactorily model real-time safety-critical systems as they disregard not only the continuous dynamics of the physical environment but also the presence of stochastic behavior. Stochastic behavior in such systems stems from many different sources, e.g., faulty

or unreliable sensors or actuators, uncertainty in timing delays, the random coin flips of distributed communication and security protocols.

Timed automata [2] were introduced as a formalism to model asynchronous real-time systems interacting with a continuous physical environment. Timed automata and their two-player counterparts [3] provide an intuitive and semantically unambiguous way to model non-stochastic real-time systems, and a number of case-studies [23] demonstrate their application in the design and analysis of real-time systems. On the other hand, classical formalisms (discrete-time and continuous-time) Markov decision processes (MDPs) and stochastic games [22, 15] naturally model analysis and synthesis problems for stochastic systems, and have been applied in control theory, operations research, and economics.

For the formal analysis of stochastic real-time systems, a number of recent works considered a combination of stochastic features with timed automata, e.g. probabilistic timed automata [18], continuous probabilistic timed automata [17] and stochastic timed automata [9]. Probabilistic timed automata, respectively continuous probabilistic and stochastic timed automata can be considered as generalizations of timed automata with the features of discrete-time Markov decision processes, respectively continuous-time Markov chains [5] (or even generalized semi-Markov processes [13]). Stochastic timed games [12] form the most general formalism for studying controller-synthesis for stochastic real-time systems. These games can be considered as interactions between three players – Player Box, Player Diamond and the stochastic player (Nature) – such that Player Box and Player Diamond are adversarial and choose their delay and action so as to maximize and minimize probability to reach a given set of target states, while the stochastic player plays according to a given probability distribution. A key verification problem in this setting is that of games with reachability objectives, where the goal of Player Diamond is to reach a set of target states, while the goal of the Player Box is to avoid it.

Related Work. Probabilistic timed automata [18] and games [16] can be considered as subclasses of stochastic timed games where all of the locations controlled by stochastic players are *urgent* (no time delay allowed), while the decision-stochastic timed automata of [10] can be seen as a subclass of $1\frac{1}{2}$-player STGs where the locations of the rational players are urgent. The quantitative reachability problem for probabilistic timed automata is known to be decidable [18] with any number of clocks, while the best known decidability result for the quantitative reachability problem for $1\frac{1}{2}$-player STGs is using a single clock. $\frac{1}{2}$-player STGs, also called stochastic timed automata (STA) [9], have also received considerable attention: an abstraction based on the region abstraction has been proposed, which allows to solve the qualitative reachability problem under a *fairness* assumption on the STA (several subclasses of STAs have been proven to be fair). For quantitative reachability, the only decidability result is for a subclass of single-clock STA [8], but a recent approximability result has been shown in [7] for the class of *fair STA*.

Other variants of stochastic timed automata have been studied in the past. The model in [17] uses "countdown clocks" (which decrease from a set value) unlike the more timed-automata style of clock variables used in our model. The model in [11] (which is also called stochastic timed automata; we shall refer to them here as Modest-STA) is very general and encompasses most models with time and probabilities (and in particular the STA of [9]). However, Modest-STA is more aimed at capturing general languages (and providing a tool-set to simulate their runs) and less with decidability issues, and hence is orthogonal to our approach.

Table 1 Results in bold are contributions from this paper. "Conj" are conjectures.

Model		Qualitative Results	Quantitative Results
$\frac{1}{2}$ player	1 clock	Dec. [4]	Dec. (some restrictions) [8]
	n clocks	Open in general Dec. (fair) [9]	Open in general Approx. (fair) [7]
$1\frac{1}{2}$ player	1 clock	Dec. [12]	**Dec. (Initialized, Theorem 8)**
	n clocks	Open	**Undec. (Theorem 3)** Conj: **Undec.** (Time bounded)
$2\frac{1}{2}$ player	1 clock	Conj: **Dec.**	**Dec. (Initialized, Corollary 9)**
	n clocks	Open	Undec [12] **Undec. (Time bounded, Theorem 6)**

Contributions.   The scope of this paper is to investigate decidability of the reachability problem in STGs as defined in [12], for which the decidability picture is far from complete. In [12], the authors showed the decidability of qualitative reachability problem on 1-clock $1\frac{1}{2}$-player STGs, and the undecidability of quantitative reachability problem on STGs (with $2\frac{1}{2}$-players). This leaves a wide gap in the decidability horizon of STGs. In this paper, we study $1\frac{1}{2}$, $2\frac{1}{2}$-player games and contribute to a better understanding of the decidability status of STGs with quantitative reachability objectives.

Table 1 summarizes the results presented in this paper. We show that the quantitative reachability problem is already undecidable for $1\frac{1}{2}$-player games for systems with 4 or more clocks and for $2\frac{1}{2}$-player games the quantitative reachability problem remains undecidable even under the time-bounded restriction with 5 or more clocks. Another key contribution of this paper is the characterization of a previously unexplored subclass of stochastic timed games for which we recover decidability of quantitative reachability game for $1\frac{1}{2}$ (and even $2\frac{1}{2}$)-player stochastic timed games. We call a 1-clock stochastic timed game *initialized* if (i) all the transitions from non-stochastic states to stochastic states reset the clock, and (ii) in every bounded cycle, the clock is reset. The definition can be generalized to multiple clocks using the notion of strong reset where one resets all the clocks together. For some of the gaps in this spectrum, we provide our best conjectures as justified in the Discussion section:–the undecidability of time-bounded quantitative reachability for $1\frac{1}{2}$-player STG, and the decidability of qualitative reachability of 1-clock $2\frac{1}{2}$-player STG. Due to lack of space, details of some proofs can be found in [1].

2   Stochastic Timed Games

We use standard notations for the set of reals ($\mathbb{R}$), rationals ($\mathbb{Q}$), and integers ($\mathbb{Z}$), and add subscripts to indicate additional constraints (for instance $\mathbb{R}_{\geq 0}$ is for the set of non-negative reals). Let $\mathcal{C}$ be a finite set of real-valued variables called *clocks*. A *valuation* on $\mathcal{C}$ is a function $\nu : \mathcal{C} \to \mathbb{R}_{\geq 0}$. We assume an arbitrary but fixed ordering on the clocks and write x_i for the clock with order i. This allows us to treat a valuation ν as a point $(\nu(x_1), \nu(x_2), \ldots, \nu(x_n)) \in \mathbb{R}_{\geq 0}^{|\mathcal{C}|}$. Abusing notations slightly, we use a valuation on $\mathcal{C}$ and a point in $\mathbb{R}_{\geq 0}^{|\mathcal{C}|}$ interchangeably. For a subset of clocks $X \subseteq \mathcal{C}$ and valuation $\nu \in \mathbb{R}_{\geq 0}^{|\mathcal{C}|}$, we write $\nu[X:=0]$ for the valuation where $\nu[X:=0](x) = 0$ if $x \in X$, and $\nu[X:=0](x) = \nu(x)$ otherwise. For $t \in \mathbb{R}_{\geq 0}$, write $\nu + t$ for the valuation defined by $\nu(x) + t$ for all $x \in X$. The valuation $\mathbf{0} \in \mathbb{R}_{\geq 0}^{|\mathcal{C}|}$ is a special valuation such that $\mathbf{0}(x) = 0$ for all $x \in \mathcal{C}$. A clock constraint over $\mathcal{C}$ is

a subset of $\mathbb{R}_{\geq 0}^{|\mathcal{C}|}$ defined by a (finite) conjunction of constraints of the form $x \bowtie k$, where $k \in \mathbb{Z}_{\geq 0}$, $x \in \mathcal{C}$, and $\bowtie \in \{<, \leq, =, >, \geq\}$. We write $\varphi(\mathcal{C})$ for the set of clock constraints. For a constraint $g \in \varphi(\mathcal{C})$, and a valuation ν, we write $\nu \models g$ to represent the fact that valuation ν satisfies constraint g (defined in a natural way). A timed automaton (TA) [2] is a tuple $\mathcal{A} = (L, \mathcal{C}, E, \mathcal{I})$ such that (i) L is a finite set of locations, (ii) $\mathcal{C}$ is a finite set of clocks, (iii) $E \subseteq L \times \varphi(\mathcal{C}) \times 2^{\mathcal{C}} \times L$ is a finite set of edges, (iv) $\mathcal{I} : L \to \varphi(\mathcal{C})$ assigns an invariant to each location. A state s of a timed automaton is a pair $s = (\ell, \nu) \in L \times \mathbb{R}_{\geq 0}^{|\mathcal{C}|}$ such that $\nu \models \mathcal{I}(\ell)$ (the clock valuation should satisfy the invariant of the location). If $s = (\ell, \nu)$, and $t \in \mathbb{R}_{\geq 0}$, we write $s + t$ for the state $(\ell, \nu + t)$. A transition (t, e) from a state $s = (\ell, \nu)$ to a state $s' = (\ell', \nu')$ is written as $s \xrightarrow{t,e} s'$ if $e = (\ell, g, C, \ell') \in E$, such that $\nu + t \models g$, and for every $0 \leq t' \leq t$ we have $\nu + t' \models \mathcal{I}(\ell)$ and $\nu' = \nu + t[C{:=}0](x)$. A run is a finite or infinite sequence of transitions $\rho = s_0 \xrightarrow{t_1,e_1} s_1 \xrightarrow{t_2,e_2} s_2 \ldots$ of states and transitions. An edge e is enabled from s whenever there is a state s' such that $s \xrightarrow{0,e} s'$. Given a state s of $\mathcal{A}$ and an edge e, we define $I(s, e) = \{t \in \mathbb{R}_{\geq 0} \mid s \xrightarrow{t,e} s'\}$ for some s' and $I(s) = \bigcup_{e \in E} I(s, e)$. We say that $\mathcal{A}$ is non-blocking iff for all states s, $I(s) \neq \emptyset$. Now we are ready to introduce stochastic timed games.

▶ **Definition 1** (Stochastic Timed Games [12])**.** A *stochastic timed game (STG)* is a tuple $\mathcal{G} = (\mathcal{A}, (L_\square, L_\diamond, L_\bigcirc), \omega, \mu)$ where

- $\mathcal{A} = (L, \mathcal{C}, E, \mathcal{I})$ is a timed automaton;
- $L_\square, L_\diamond$, and $L_\bigcirc$ form a partition of L characterizing the set of locations controlled by players $\square$ and $\diamond$ and the stochastic player, respectively;
- $\omega : E(L_\bigcirc) \to \mathbb{Z}_{>0}$ assigns some positive weight to each edge originating from $L_\bigcirc$ (notation $E(L_\bigcirc)$);
- μ is a function assigning a measure over $I(s)$ to all states $s \in L_\bigcirc \times \mathbb{R}_{\geq 0}^{|\mathcal{C}|}$ satisfying the properties that $\mu(s)(I(s)) = 1$ and for Lebesgue measure λ, if $\lambda(I(s)) > 0$ then for each measurable set $B \subseteq I(s)$ we have $\lambda(B) = 0$ if and only if $\mu(s)(B) = 0$.

The timed automaton $\mathcal{A}$ is said equipped with uniform distributions over delays if for every state s, $I(s)$ is bounded, and $\mu(s)$ is the uniform distribution over $I(s)$. The timed automaton $\mathcal{A}$ is said equipped with exponential distributions over delays whenever, for every state s, either $I(s)$ has Lebesgue measure zero, or $I(s) = \mathbb{R}_{\geq 0}$ and for every location l, there is a positive rational α_l such that $\mu(s)(I(s)) = \int_{t \in I} \alpha_l e^{-\alpha_l t} dt$. For $s \in L_\bigcirc \times \mathbb{R}_{\geq 0}^{|\mathcal{C}|}$, both delays and discrete moves will be chosen probabilistically: from s, a delay t is chosen following the probability distribution over delays $\mu(s)$. Then, from state $s + t$, an enabled edge is selected following a discrete probability distribution that is given in a usual way with the weight function w: in state $s + t$, the probability of edge e (if enabled), denoted $p(s + t)(e)$ is $w(e)/\sum_{e'} \{w(e') \mid e' \text{ is enabled in } s + t\}$. This way of probabilizing behaviours in timed automata has been presented in [9].

If $L_\square = \emptyset$ then the STGs are called $1\frac{1}{2}$ STGs or $1\frac{1}{2}$-player STGs while STGs with $L_\square = L_\diamond = \emptyset$ are called $\frac{1}{2}$ STGs or $\frac{1}{2}$-player STGs or STAs. We often refer to $l \in L_\bigcirc$ as stochastic nodes, $l \in L_\square$ as box (or $\square$) nodes and $l \in L_\diamond$ as diamond (or $\diamond$) nodes.

Fix a STG $\mathcal{G} = (\mathcal{A}, (L_\square, L_\diamond, L_\bigcirc), \omega, \mu)$ with $\mathcal{A} = (L, \mathcal{C}, E, \mathcal{I})$ for the rest of this section.

Strategies, Profiles, and Runs. A strategy for Player $\square$ (resp. $\diamond$) is a function that maps a finite run $\rho = s_0 \xrightarrow{t_0,e_0} s_1 \xrightarrow{t_1,e_1} \ldots s_n$ to a pair (t, e) such that $s_n \xrightarrow{t,e} s'$ for some state s', whenever $s_n = (\ell_n, \nu_n)$ and $\ell_n \in L_\square$ (resp. $\ell_n \in L_\diamond$). In this work we focus on deterministic strategies, though randomized strategies could also make sense; nevertheless understanding the case of deterministic strategies is already challenging. A strategy profile

is a pair $\Lambda = (\lambda_\diamond, \lambda_\square)$ where $\lambda_\diamond, \lambda_\square$ respectively are strategies of players $\diamond$ and $\square$. In order to measure probabilities of certain sets of runs, the following measurability condition is imposed on strategy profiles $\Lambda = (\lambda_\diamond, \lambda_\square)$: for every finite sequence of edges $e_1, \ldots, e_n$ and every state s, the function $\kappa_s : (t_1, \ldots, t_n) \to (t, e)$ defined by $\kappa_s(t_1, \ldots, t_n) = (t, e)$ iff $\Lambda(s \xrightarrow{t_1, e_1} s_1 \xrightarrow{t_2, e_2} s_2 \ldots \xrightarrow{t_n, e_n} s_n) = (t, e)$, should be measurable.

Given a finite run ρ ending in state s_0, and a strategy profile Λ, define $Runs(\mathcal{G}, \rho, \Lambda)$ (resp. $Runs^\omega(\mathcal{G}, \rho, \Lambda)$) to be the set of all finite (resp. infinite) runs generated by Λ after prefix ρ; that is, the set of all runs of the automaton satisfying the following condition: If $s_i = (\ell_i, \nu_i)$ and $\ell_i \in L_\diamond$ (resp. $\ell_i \in L_\square$), then $\lambda_\diamond$ (resp. $\lambda_\square$) returns (t_{i+1}, e_{i+1}) when applied to $\rho \xrightarrow{t_1, e_1} s_1 \xrightarrow{t_2, e_2} \ldots \xrightarrow{t_i, e_i} s_i$. Given a finite sequence $e_1, \ldots, e_n$ of edges, a symbolic path $\pi_\Lambda(\rho, e_1 \ldots e_n)$ is defined as

$$\pi_\Lambda(\rho, e_1 \ldots e_n) = \{\rho' \in Runs(\mathcal{G}, \rho, \Lambda) \mid \rho' = \rho \xrightarrow{t_1, e_1} s_1 \xrightarrow{t_2, e_2} s_2 \ldots \xrightarrow{t_n, e_n} s_n, \text{ with } t_i \in \mathbb{R}_{\geq 0}\}.$$

When Λ is clear, we simply write $\pi(\rho, e_1 \ldots e_n)$.

Probability Measure of a Strategy Profile. Given a strategy profile $\Lambda = (\lambda_\diamond, \lambda_\square)$, and a finite run ρ ending in $s = (\ell, \nu)$, a measure $\mathcal{P}_\Lambda$ can be defined on the set $Run(\mathcal{G}, \rho, \Lambda)$, following [12]: First, for the empty sequence ϵ, $\mathcal{P}_\Lambda(\pi(\rho, \epsilon)) = 1$, and

- If $\ell \in L_\diamond$ (resp. $\ell \in L_\square$), and $\lambda_\diamond(\rho) = (t, e)$ (resp. $\lambda_\square(\rho) = (t, e)$), then
  $\mathcal{P}_\Lambda(\pi(\rho, e_1 \ldots e_n))$ equals 0 if $e_1 \neq e$ and equals $\mathcal{P}_\Lambda(\pi(\rho \xrightarrow{t, e} s', e_2 \ldots e_n))$, otherwise.
- If $\ell \in L_\bigcirc$, $\mathcal{P}_\Lambda(\pi(\rho, e_1 \ldots e_n)) = \int_{t \in I(s, e_1)} p(s + t)(e_1) \cdot \mathcal{P}_\Lambda(\pi(\rho \xrightarrow{t, e_1} s', e_2 \ldots e_n)) \, d\mu(s)(t)$
  where $s \xrightarrow{t, e_1} s'$ for every $t \in I(s, e_1)$.

The cylinder generated by a symbolic path is defined as follows: an infinite run ρ'' is in the cylinder generated by $\pi_\Lambda(\rho, e_1, \ldots, e_n)$ denoted $\mathsf{Cyl}(\pi_\Lambda(\rho, e_1, \ldots, e_n))$ if ρ'' is in $Runs^\omega(\mathcal{G}, \rho, \Lambda)$ and there is a finite prefix ρ' of ρ'' such that $\rho' \in \pi_\Lambda(\rho, e_1, \ldots, e_n)$. It is routine to extend the above measure $\mathcal{P}_\Lambda$ to cylinders, and thereafter to the generated σ-algebra; extending [9], one can show this is a probability measure over $Runs^\omega(\mathcal{G}, \rho, \Lambda)$.

Example. An example of a STG is shown in the adjoining figure. In this example all the locations belong to stochastic player (this is an $\frac{1}{2}$ STG) and there is only one clock named x.

We explain here the method for computing probabilities. We assume uniform distribution over delays at all states, and initial state $s_0 = (A, 0)$. Let $d\mu_{(A,0)}$ be the uniform distribution over $[0, 1]$ and $d\mu_{(B,0)}$ uniform distribution over $[0, 2]$. Then $\mathcal{P}(\pi((A, 0), e_1 e_2))$ equals $\frac{1}{8}$:

$$\int_0^1 \frac{\mathcal{P}(\pi((B, 0), e_2))}{2} d\mu_{(A,0)}(t) = \int_0^1 \frac{1}{2} \left(\int_1^2 \frac{1}{2} d\mu_{(B,0)}(u) \right) d\mu_{(A,0)}(t) = \frac{1}{2} \int_0^1 \left(\int_1^2 \frac{1}{2} \frac{1}{2} du \right) dt)$$

Reachability Problem. We study the reachability problem for STGs, stated as follows. Given a STG $\mathcal{G}$ with a set T of target locations, an initial state s_0 and a threshold $\bowtie p$ with $p \in [0, 1] \cap \mathbb{Q}$, decide whether there is a strategy $\lambda_\diamond$ for Player $\diamond$ such that for every strategy $\lambda_\square$ for Player $\square$, $\mathcal{P}_\Lambda(\{\rho \in Run(\mathcal{G}, s_0, \Lambda) \mid \rho \text{ visits } T\}) \bowtie p$, with $\Lambda = (\lambda_\diamond, \lambda_\square)$. There are two categories of reachability questions:

1. **Quantitative reachability**: The constraint on probability involves $0 < p < 1$.
2. **Qualitative reachability**: The constraint on probability involves $p \in \{0, 1\}$.

The key results of the paper are the following:

▶ **Theorem 2.** *The quantitative reachability problem is*
1. *Undecidable for $1\frac{1}{2}$ STGs with 4 or more clocks;*
2. *Undecidable for $2\frac{1}{2}$ STGs with 5 or more clocks even under the time-bounded semantics;*
3. *Decidable for $1\frac{1}{2}$ and $2\frac{1}{2}$ initialized STGs with one clock.*

Mentioned restrictions (time-bounded semantics and initialized) will be introduced when needed. In Section 3, we deal with the quantitative reachability problem, where we show strengthened undecidability results. In Section 4, we explore a new model of STGs with a single clock and an initialized restriction to recover decidability for the quantitative reachability problem. In Section 5, we discuss the intrinsic difficulties and challenges ahead, summarize our key contributions and conjectures.

3 Undecidability Results for Quantitative Reachability

In this section, we focus on the quantitative reachability problem for STGs. We strengthen the existing undecidability result, which holds for $2\frac{1}{2}$ STGs [12], in two distinct directions. First, we show the undecidability of the quantitative reachability problem in $1\frac{1}{2}$ STGs, improving from $2\frac{1}{2}$. Second, we show the undecidability of the quantitative reachability problem for $2\frac{1}{2}$ STGs even in the time-bounded setting.

For both results, given a two-counter machine, we construct respectively, $1\frac{1}{2}$ and $2\frac{1}{2}$ STGs whose building blocks are the modules for the instructions in the two-counter machine. The objective of player $\diamond$ is linked to a faithful simulation of various increment, decrement and zero-test instructions of the two-counter machine by choosing appropriate delays to adjust the clocks to reflect changes in counter values. However, the two proofs differ in how this verification is done and even in the problem from which the reduction is done, i.e., halting/non-halting for two-counter machines. This results in two quite different and non-trivial reductions as described in Subsection 3.1 and Subsection 3.2 respectively.

3.1 Quantitative reachability for $1\frac{1}{2}$ STGs

As mentioned above, in the case of $1\frac{1}{2}$ STGs we improve the corresponding result of [12] for $2\frac{1}{2}$ STGs. But unlike in [12], we reduce from the *non-halting problem* for two-counter machines to the existence of a winning strategy for Player $\diamond$ with the desired objective. This crucial difference makes it possible for the probabilistic player to verify the simulation performed by player $\diamond$.

▶ **Theorem 3.** *The quantitative reachability problem is undecidable for $1\frac{1}{2}$ STGs with ≥ 4 clocks.*

Let $\mathcal{M}$ be a two-counter machine. Our reduction uses a $1\frac{1}{2}$ player STG $\mathcal{G}$ with four clocks and uniform distributions over delays, and a set of target locations T such that player $\diamond$ has a strategy to reach T with probability $\frac{1}{2}$ iff $\mathcal{M}$ does not halt. Each instruction (increment, decrement and test for zero value) is specified using a module. The main invariant in our reduction is that upon entry into a module, we have that $x_1 = \frac{1}{2^{c_1}}, x_2 = \frac{1}{2^{c_2}}, x_3 = x_4 = 0$, where c_1 (resp. c_2) is the value of counter C_1 (resp. C_2) in $\mathcal{M}$.

We outline the simulation of an increment instruction « ℓ_i : increment counter C_1, goto ℓ_j » in Figure 1 (top). The module is entered with values $x_1 = \frac{1}{2^{c_1}}, x_2 = \frac{1}{2^{c_2}}, x_3 = x_4 = 0$. A time $1 - \frac{1}{2^{c_1}}$ is spent at location ℓ_i, so that at location B we have $x_1 = 0$, $x_2 = \frac{1}{2^{c_2}} + 1 - \frac{1}{2^{c_1}}$ (or $\frac{1}{2^{c_2}} - \frac{1}{2^{c_1}}$, if $c_2 > c_1$ – we write in all cases $\frac{1}{2^{c_2}} + 1 - \frac{1}{2^{c_1}}$ mod 1), $x_3 = 1 - \frac{1}{2^{c_1}}$, $x_4 = 0$.
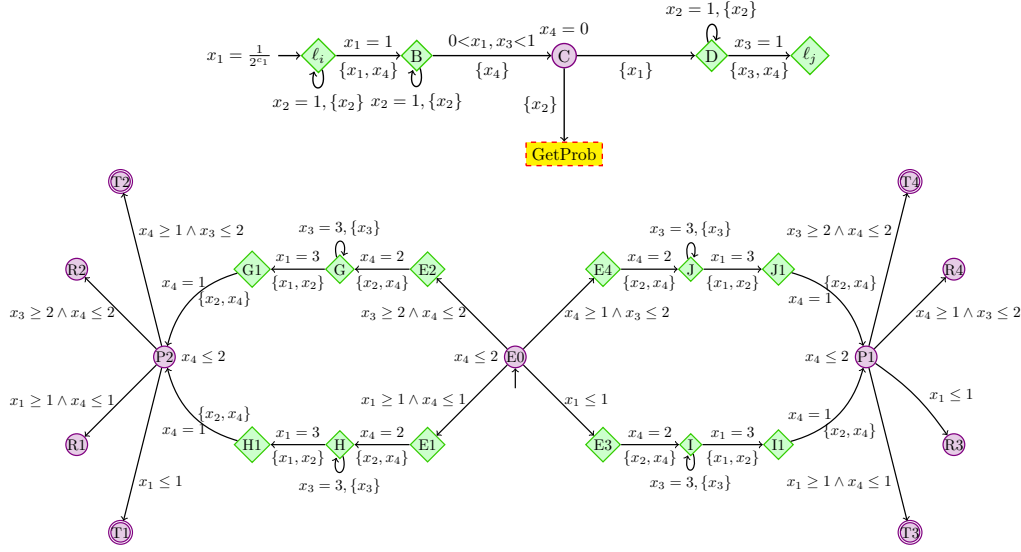
Figure 1 The Increment c_1 module on the top and the GetProb gadget below

An amount of time $t \in (0, \frac{1}{2^{c_1}})$ is spent at B, which is decided by Player $\diamond$. We rewrite this as $t = \frac{1}{2^{c_1+1}} \pm \epsilon$ for $-\frac{1}{2^{c_1+1}} < \epsilon < \frac{1}{2^{c_1+1}}$. This is because, ideally we want t to be $\frac{1}{2^{c_1+1}}$ and want to consider any deviation as an error.

Now at C, we have $x_1 = t$, $x_2 = \frac{1}{2^{c_2}} + 1 - \frac{1}{2^{c_1}} + t \bmod 1$, $x_3 = 1 - \frac{1}{2^{c_1}} + t$, $x_4 = 0$. The computation proceeds to D with probability $\frac{1}{2}$, and the location ℓ_j corresponding to the next instruction ℓ_j is reached with $x_1 = \frac{1}{2^{c_1}} - t$, $x_2 = \frac{1}{2^{c_2}}$, $x_3 = x_4 = 0$. On the other hand, with probability $\frac{1}{2}$, the gadget $GetProb$ is reached. The gadget $GetProb$ has 4 target locations $T1, T2, T3, T4$, which we will show are reached with probability $\frac{1}{2}$ from the start location $E0$ of $GetProb$ iff $t = \frac{1}{2^{c_1+1}}$. Thus, in this case when $t = \frac{1}{2^{c_1+1}}$, we reach ℓ_j with the values $x_1 = \frac{1}{2^{c_1+1}}$, $x_2 = \frac{1}{2^{c_2}}$, $x_3 = x_4 = 0$ which implies that c_1 has been incremented correctly according to our encoding. We now look at the gadget $GetProb$.

▶ **Lemma 4.** *For any value $\epsilon \in (-\frac{1}{2^{c_1+1}}, \frac{1}{2^{c_1+1}})$, the probability to reach a target location in $GetProb$ from $E0$ is $\frac{1}{2}(1 - 4\epsilon^2)$ $(\leq \frac{1}{2})$. Further this probability is equal to $\frac{1}{2}$ iff $\epsilon = 0$.*

Proof. Note that when the start location $E0$ of $GetProb$ is reached, we have $x_1 = \frac{1}{2^{c_1+1}} + \epsilon$, $x_2 = 0$, $x_3 = 1 - \frac{1}{2^{c_1+1}} + \epsilon$, $x_4 = 0$. A total of 2 time units can be spent at $E0$. It can be seen that transitions to $E3$ and $E4$ are respectively enabled with the time intervals $[0, 1 - \frac{1}{2^{c_1+1}} - \epsilon]$ and $[1, 1 + \frac{1}{2^{c_1+1}} - \epsilon]$. Similarly, reaching $E1$ and $E2$ are enabled by the time intervals $[1 - \frac{1}{2^{c_1+1}} - \epsilon, 1]$ and $[1 + \frac{1}{2^{c_1+1}} - \epsilon, 2]$. The sum of probabilities of reaching either $E3$ or $E4$ is thus $\frac{1}{2}(1 - 2\epsilon)$. Similarly, the sum of probabilities for reaching $E1$ or $E2$ is $\frac{1}{2}(1 + 2\epsilon)$. The locations $P1, P2$ are then reached with the values $x_1 = \frac{1}{2^{c_1+1}} + \epsilon$, $x_2 = 0$, $x_3 = 1 - \frac{1}{2^{c_1+1}} + \epsilon$, $x_4 = 0$. The probability of reaching the target locations $T3$ or $T4$ (i.e., through $P1$) from $E0$ is hence $\frac{1}{2}(1+2\epsilon)\frac{1}{2}(1-2\epsilon) = \frac{1}{4}(1-4\epsilon^2)$, while the probability of reaching a target location $T1$ or $T2$ (i.e., through $P2$) from $E0$ is $\frac{1}{2}(1 + 2\epsilon)\frac{1}{2}(1 - 2\epsilon) = \frac{1}{4}(1 - 4\epsilon^2)$. Thus, the probability of reaching a target location (one of $T1, T2, T3, T4$) in $GetProb$ is, $\frac{1}{2}(1 - 4\epsilon^2)$, which is always $\leq \frac{1}{2}$. This completes the first statement of the lemma. Further, from the expression, we immediately have that the probability to reach a target location in $GetProb$ from $E0$ is $\frac{1}{2}$ iff $\epsilon = 0$. ◀

The decrement c_1, increment c_2 as well as decrement c_2 modules are similar and these as well as the zero test modules can be found in [1].

▶ **Lemma 5.** *Player $\diamond$ has a strategy to reach the (set of) target locations in $\mathcal{G}$ with probability $\frac{1}{2}$ iff the two-counter machine does not halt.*

Proof. Suppose the two-counter machine halts (say in k steps). Then there are two cases: (a) the simulations of all instructions are correct in $\mathcal{G}$. In this case, the target location can be reached in either of the first k steps. By Lemma 4, the probability of reaching a target location in the first k steps is the summation $\frac{1}{2} \cdot \frac{1}{2} + (\frac{1}{2})^2 \cdot \frac{1}{2} + (\frac{1}{2})^3 \cdot \frac{1}{2} + \cdots + (\frac{1}{2})^k \cdot \frac{1}{2} < \frac{1}{2}$. (b) Player $\diamond$ made an error in the computation in the first k steps. But then again by Lemma 4, the finite sum obtained is $< \frac{1}{2}$ (since in the error step(s), the probability to reach target locations is $\frac{1}{2} - 4\epsilon^2 < \frac{1}{2}$). Thus, if the two-counter machine halts, under any strategy of $\diamond$ player, the probability to reach the target locations is $< \frac{1}{2}$.

On the other hand, suppose the two-counter machine does not halt. Then, if Player $\diamond$ chooses the strategy which faithfully simulates all instructions of the two-counter machine, the probability to reach the (set of) target locations is given by the infinite sum $\sum_{i=0}^{\infty} (\frac{1}{2})^i \frac{1}{2} = \frac{1}{2}$. Any other strategy of Player $\diamond$ corresponds to performing at least one error in the simulation. In this case, the infinite sum obtained has at least one term of the form $(\frac{1}{2})^k (\frac{1}{2} - 4\epsilon^2)$, for $\epsilon^2 > 0$. Clearly, such an infinite sum does not sum to $\frac{1}{2}$. This concludes the proof. ◀

The previous proof can be changed for other thresholds and to use unbounded intervals and exponential distributions.

3.2 Time-bounded quantitative reachability for $2\frac{1}{2}$ STGs

In this section, we tackle the *time-bounded* version of the quantitative reachability problem. This strengthens the definition of reachability by considering a given time bound Δ, and requiring that $\mathcal{P}_\sigma(\{\rho \in Run(\mathcal{G}, s_0, \sigma) \mid \rho$ visits T within Δ time units$) \bowtie p$.

In this new framework, we show the undecidability of the quantitative reachability problem for $2\frac{1}{2}$ STGs. We reduce from the *halting* problem for two-counter machines (unlike in the previous section, where our reduction was from the *non-halting* problem), using Player $\square$ to verify the correctness of the simulation. The complication here is that the total time spent should be bounded and hence we cannot allow arbitrary time elapses. We will in fact show a global time bound of $\Delta = 5$ for this reduction.

▶ **Theorem 6.** *The time-bounded quantitative reachability problem is undecidable for $2\frac{1}{2}$ STGs with ≥ 5 clocks.*

Proof. Let $\mathcal{M}$ be a two-counter machine. We construct an STG with 5 clocks such that the two-counter machine $\mathcal{M}$ halts iff Player $\diamond$ has a strategy to reach some desired locations with probability $\frac{1}{2}$, whatever Player $\square$ does, and such that the total time spent is bounded by $\Delta = 5$ units.

The main idea behind the proof is that the total time spent in the simulation of the k^{th} instruction will be $\frac{1}{2^k}$. We thus get a decreasing sequence of times $\frac{1}{2}, \frac{1}{4}, \frac{1}{8} \ldots$ for simulating the instructions $1, 2 \ldots$ and so on. In total, we will use five clocks x_1, x_2, z, a and b. The clocks x_1 and x_2 are used encode the counter values (along with the current instruction number) such that at the end of the k^{th} instruction, if k is even the values are encoded in x_1 and if k is odd they are encoded in x_2 as follows:

(enc_{x_1}) k is even and $x_1 = \frac{1}{2^{k+c_1}3^{k+c_2}}$, $x_2 = 0$, $z = 1 - \frac{1}{2^k}$, $a = b = 0$;
(enc_{x_2}) k is odd and $x_2 = \frac{1}{2^{k+c_1}3^{k+c_2}}$, $x_1 = 0$, $z = 1 - \frac{1}{2^k}$, $a = b = 0$;

Figure 2 Module for incrementing C_1 (after an even number of steps)

Figure 3 Widgets 'Check z' (left) and 'Check x_2' (right).

We start the simulation with $x_1 = 1, x_2 = z = 0 = a = b$ corresponding to the initial instruction ($k = 0$) and the fact that the values of C_1, C_2 are 0. Moreover, if $x_1 = \frac{1}{2^{k+c_1}3^{k+c_2}}$ at the end of the kth instruction, and if the $(k+1)$th instruction is an increment C_1 instruction, then at the end of the $(k+1)$th instruction, $x_2 = \frac{1}{2^{k+c_1+2}3^{k+c_2+1}}$. Clock z keeps a separate track of the number of instructions simulated so far, by having a value $1 - \frac{1}{2^k}$ after completing the simulation of k instructions. Clocks a and b are auxiliary clocks that we need for the simulation. We assume uniform distribution over delays in probabilistic locations. If no weight is written on an edge, it is assumed to be 1. We outline the simulation of a increment instruction « ℓ_i : increment counter C_1, goto ℓ_j » in Figure 2, assuming this is the $(k+1)$th instruction, where k is even. Thus, at the end of the k first instructions, we have $x_1 = \frac{1}{2^{k+c_1}3^{k+c_2}}$, $z = 1 - \frac{1}{2^k}$ and $a = b = x_2 = 0$ (the other case of odd k, i.e., (enc_{x_2}) encoding is symmetric). At the end of this $(k+1)$th instruction's simulation, the value of clock z should be $z = 1 - \frac{1}{2^{k+1}}$ to mark the end of the $(k+1)^{th}$ instruction. Also, we must obtain $x_2 = \frac{x_1}{2^2 \cdot 3} = \frac{x_1}{12}$, marking the successful increment of C_1.

Player $\diamond$ elapses times t_1, t_2 in locations ℓ_i, B. When the player $\square$ location *Check* is reached, we have $a = t_1 + t_2 = t$ and $x_2 = t_2$, $z = 1 - \frac{1}{2^k} + t_1 + t_2$. Player $\square$ has three possibilities : (1) to continue the simulation going to ℓ_{k+2}, (2) verify that $t_2 = \frac{1}{2^{k+c_1+2}3^{k+c_2+1}}$ by going to the widget 'Check x_2' or (3) verify that $t_1 + t_2 = \frac{1}{2^{k+1}}$ by going to the widget 'Check z'. These widgets are given in Figure 3. The probability of reaching a target location in widget 'Check z' is $\frac{1}{2}(1-t) + \frac{1}{4}\frac{1}{2^k} = \frac{1}{2}$ iff $t = \frac{1}{2^{k+1}}$. In widget 'Check x_2', the transitions from $F1$ to $C1$ and $F1$ to $C2$ are taken with probability $\frac{1}{12}$ and $\frac{11}{12}$, respectively since the weights of edges connecting F1,C1 and F1,C2 are respectively 1 and 11. With this, for $n = \frac{1}{2^{k+c_1}3^{k+c_2}}$, the probability of reaching a target location in 'Check x_2' is $\frac{1}{2}(1-t_2) + \frac{n}{24} = \frac{1}{2}$ iff $t_2 = \frac{n}{12}$.

Time elapse for Increment. If player $\square$ goes ahead with the simulation, the time elapse for the $(k+1)$th instruction is $t_1 + t_2 = \frac{1}{2^{k+1}}$. Consider the case when player $\square$ goes in to 'Check z'. The time elapse till now is $\frac{1}{2} + \cdots + \frac{1}{2^{k+1}}$. The time spent in the 'Check z' widget is as follows: one unit is spent at location $B0$, one unit at location $F0$, and $1 - t$ units at location

$E0$. Thus, ≤ 3 units are spent at the 'Check z' widget. Similarly, the time spent in the 'Check x_2' widget is one unit at $B1$, $1 - t$ units at $C1$, $1 - n$ units at $D1$ and one unit at $E1$. Thus a time ≤ 4 is spent in 'Check x_2'. Thus, the time spent till the $(k + 1)$th instruction is $\leq \frac{1}{2} + \cdots \frac{1}{2^{k+1}} + 4$ if player $\square$ goes in for a check, and otherwise it is $\frac{1}{2} + \cdots + \frac{1}{2^{k+1}}$.

Other increment, decrement, zero-check Instructions. The main module corresponding to *increment* C_2 and decrement C_1, C_2 is the same as in Figure 2. The only change needed is in the 'Check x_2' widget. While incrementing c_2, we need $x_2 = \frac{x_1}{2 \cdot 3^2} = \frac{x_1}{18}$. This is done by changing the weights on the outgoing edges from $F1$ to $C1$ and $C2$ to 1 and 17 respectively. Similarly, while *decrementing* C_1, we need $x_2 = \frac{x_1}{3}$. This is done by changing the weights on the outgoing edges of $F1$ to $1, 2$ respectively. Lastly, to *decrement* C_2, we need $x_2 = \frac{x_1}{2}$, and in this case the weights are 1 each.

The zero check module is a bit more complicated. The broad idea is that we use a diamond node to guess whether the current clock (say C_1) value is zero and branch into two sides (zero and non-zero). Then we use a box node on each branch to verify that the guess was correct. If correct, we proceed with the next instruction, if not, we check this by going to a special widget. In this widget, we can reach a target node with probability $\frac{1}{2}$ iff the guess is correct. The details of this widget and the proof that all these simulations can be done in time bounded by $\Delta \leq 5$ units is given in [1]. ◀

4   Decidability results for quantitative reachability

We have seen in the previous section that the quantitative reachability problem is undecidable in $1\frac{1}{2}$ STGs with ≥ 4 clocks. In this section we study the *quantitative reachability* problem in the setting of $1\frac{1}{2}$ STGs *with a single clock*. In [8], the quantitative reachability problem in $\frac{1}{2}$ STGs with a single clock, under certain restrictions, was shown to be decidable by reducing it to the quantitative reachability problem for finite Markov chains. In our case, we lift this to $1\frac{1}{2}$ STGs with a single clock, under similar restrictions, by reducing to the quantitative reachability problem in finite Markov decision processes (MDPs in short).

For the rest of this section, we consider a $1\frac{1}{2}$ STG $\mathcal{G} = (\mathcal{A}, (L_\diamond, L_\bigcirc), \omega, \mu)$ with a single clock denoted x. We write $c_{\max}$ for the maximal constant appearing in a guard of $\mathcal{G}$. We assume w.l.o.g. that target locations belong to player $\diamond$ (a slight modification of the construction can be done if this is not the case). In the following, when we talk about regions, we mean the clock regions from the classical region construction for timed automata [2, 19]: since $\mathcal{G}$ has a single clock, regions in this case are simply either singletons $\{c\}$ with $c \in \mathbb{Z}_{\geq 0} \cap [0; c_{\max}]$, or open intervals $(c, c + 1)$ with $c \in \mathbb{Z}_{\geq 0} \cap [0; c_{\max} - 1]$, or the unbounded interval $(c_{\max}; +\infty)$. While region automata are standardly finite automata, we build here from $\mathcal{G}$ a *region STG* $\mathcal{G}_\mathcal{R}$, which has only clock constraints defined by regions (that is, either $x = c$ or $c < x < c + 1$ or $x > c_{\max}$), and such that each location of $\mathcal{G}_\mathcal{R}$ is indeed a pair (ℓ, R) where ℓ is a location of $\mathcal{G}$ and R a region (region R is for the region which is hit when entering the location). While it is not completely standard, this kind of construction has been already used in [9, 8, 12], and questions asked on $\mathcal{G}$ can be equivalently asked (and answered) on $\mathcal{G}_\mathcal{R}$. Now, we make the following restrictions on $\mathcal{G}_\mathcal{R}$ (which yields restrictions to $\mathcal{G}$), which we denote $(\star)$:

1. The TA $\mathcal{A}_\mathcal{R}$ is assumed to be structurally non-Zeno: any bounded cycle of $\mathcal{A}_\mathcal{R}$ (a cycle in which all edges have a non-trivial upper-bound) contains at least one location whose associated region is the zero region (i.e., edge leading to it, resets the clock).

Figure 4 An initialized $1\frac{1}{2}$ player STG $\mathcal{G}$, its region game graph $\mathcal{G}_\mathcal{R}$ and the MDP abstraction $M_\mathcal{G}$.

2. For every state $s = ((\ell, r), \nu)$ of $\mathcal{G}_\mathcal{R}$ such that $\ell \in L_\bigcirc$, $I(s) = \mathbb{R}_{\geq 0}$, and μ_s is an exponential distribution; Furthermore the rate of μ_s only depends on location ℓ.

3. $\mathcal{G}_\mathcal{R}$ is *initialized*, that is, any edge from a non-stochastic location to a stochastic location resets the clock x.

While the first two assumptions are already made in [8], even in the $\frac{1}{2}$ player case, the third condition is new. In the following we denote **0** for the region $\{0\}$ and ∞ for the unbounded region $(c_{\max}; +\infty)$. We now show how to obtain an MDP from the STG $\mathcal{G}_\mathcal{R}$. The construction is illustrated on Figure 4. A node (ℓ, R) of $\mathcal{G}_\mathcal{R}$ with $\ell \in L_\bigcirc$ is *deletable* if R is neither the region 0 nor the region ∞. In Figure 4, $(B, (0, 1))$ and $(A, (0, 1))$ in $\mathcal{G}_\mathcal{R}$ are what we call deletable nodes. Then, we recursively remove all deletable nodes $\mathcal{G}_\mathcal{R}$ while labelling remaining paths with (finite) sequences of edges; each surviving edge is labelled by the probability of the (provably) finitely many sequences of edges appearing in the label. One can prove that this object is actually an MDP, which we denote $M_\mathcal{G}$. Target states in $M_\mathcal{G}$ are defined as the pairs (ℓ, R) where ℓ is a target location in $\mathcal{G}$. We can prove (see [1]) that:

▶ **Lemma 7.** *If $\mathcal{G}$ is an $1\frac{1}{2}$ player STG with one clock satisfying the hypotheses $(\star)$, then $M_\mathcal{G}$ is an MDP such that: (a) for every strategy $\lambda_\diamond$ of player $\diamond$ in $\mathcal{G}$, we can construct a strategy $\sigma_\diamond$ of player $\diamond$ in $M_\mathcal{G}$ such that the probability of reaching a target location in $\mathcal{G}$ is the same as the probability of reaching a target state in $M_\mathcal{G}$; and (b) for every strategy $\sigma_\diamond$ of player $\diamond$ in $M_\mathcal{G}$, we can construct a strategy $\lambda_\diamond$ of player $\diamond$ in $\mathcal{G}$ such that the probability of reaching a target location in $M_\mathcal{G}$ is the same as the probability of reaching a target state in $\mathcal{G}$.*

This lemma allows to reduce the quantitative reachability problem from the $1\frac{1}{2}$ STG $\mathcal{G}$ to the MDP $M_\mathcal{G}$.

As an example, in Figure 4, we show a $1\frac{1}{2}$ player STG $\mathcal{G}$, its region game graph $\mathcal{G}_\mathcal{R}$ (guards omitted for readability) and the MDP abstraction $M_\mathcal{G}$. Note that all $\diamond$ nodes remain, while only those stochastic nodes with regions **0** and ∞ are retained in $M_\mathcal{G}$. The stochastic nodes $(B, (0, 1))$ as well as $(C, (0, 1))$ are deleted in $M_\mathcal{G}$. On deleting nodes from the region graph, the probability on the edges of $M_\mathcal{G}$ is the probability of the respective paths from the region graph. For example, the edge from $(A, 0)$ to $(D, (0, 1))$ is labelled with $e_4 e_7$ by deleting $(B, (0, 1))$.

Thus, the remaining thing that has to be addressed now is how to compute the probabilties and compare them with a rational threshold. The first thing to note is that the edges of the MDP are all labelled with polynomials over exponentials obtained using the delays from the underlying game with rational coefficients. For example, in Figure 4, in the MDP in the rightmost picture, we obtain: $\mathcal{P}(e_1) = \mathcal{P}(e_2) = \mathcal{P}(e_5) = e^{-1}$, $\mathcal{P}(e_6, e_7, e_8) = 1 - e^{-1}$,

$\mathcal{P}(e_4e_5)=e^{-1}-e^{-2}$, $\mathcal{P}(e_4e_7)=1-2e^{-1}$, $\mathcal{P}(e_3e_4e_7)=2-5e^{-1}+e^{-2}$, $\mathcal{P}(e_3e_4e_5)=1-e^{-1}+e^{-2}$, $\mathcal{P}(e_3e_1)=\frac{1}{2}(1-e^{-2})$. It can be seen that we can write each of these probabilities as a polynomial in e^{-1}. More generally, for any MDP with differing rates (of the exponential distribution) in each state, we get a set of rational functions in $e^{-\frac{1}{q}}$ for some $q \in \mathbb{Z}_{>0}$, where q is obtained as a function of the rates in each state. Thus, using standard algorithms for MDPs [6], and as done for Markov chains in [8], we get that we can compute expressions for the probability of reaching the targets, and decide the threshold problem.

▶ **Theorem 8.** *Quantitative reachability for 1-clock $1\frac{1}{2}$-player STGs satisfying $(\star)$ is decidable.*

We can lift this construction to include $\square$ player nodes, keeping the same initialized restriction with $\square$ nodes as well. Then the region game graph $\mathcal{G}_{\mathcal{R}}$ includes $\square$ nodes in the obvious way, and we consider strategy profiles of $\square$ and $\diamond$. The question then is to check if $\diamond$ has a strategy to reach a target with probability $\sim c$ against all possible strategies of $\square$ in $M_{\mathcal{G}}$. Hence we have that

▶ **Corollary 9.** *Quantitative reachability for 1-clock $2\frac{1}{2}$ player STGs satisfying $(\star)$ is decidable.*

5 Discussion

In this paper, we have refined the decidability boundaries for STGs as summarized in the table in Introduction. The significance of our undecidability results for quantitative reachability (via different two-counter machine reductions) lies in the fact that they introduce ideas which could potentially help in settling other open problems. We highlight these below:

- for $1\frac{1}{2}$ player games, the crux is to cleverly encode the error ϵ made by player $\diamond$ in such a way that it reflects as $\frac{1}{2} - \epsilon^2$ in the resulting probability. This ensures that the $\diamond$ player can never cheat and the probability will be $< \frac{1}{2}$ as soon as there is an error (even when simulating a non-halting run of the two-counter machine). Indeed, this is why the reduction is from the non-recursively enumerable non-halting problem.

- for $2\frac{1}{2}$ player games in the *time-bounded setting*, we obtain undecidability by showing a reduction from halting problem for two-counter machines. This is surprising, as time-boundedness restores decidability in several classical undecidable problems like the inclusion problem in timed automata [20, 21]. In the case of priced timed games [14], time-boundedness gives undecidability; however, this can be attributed to the fact that price variables are not clocks, and can grow at different rates in different locations. Somehow, the combination of simple clocks and probabilities achieves the same.

Combining these ideas might allow us, for eg., to improve Theorem 6 by showing undecidability of time bounded, quantitative reachability in $1\frac{1}{2}$ player STGs with a larger number of clocks. The main challenge is to replace $\square$ player nodes by stochastic nodes, and adapt the gadgets in such a way that, within a global time bound, the probability of reaching a target is $\frac{1}{2}$ iff all simulations are correct and the two-counter machine does not halt. As another example, if in the first item above, we obtain a probability of $1 - \epsilon^2$ (rather than $\frac{1}{2} - \epsilon^2$), this would settle the (currently open) *qualitative* reachability problem for $2\frac{1}{2}$ games [12].

Coming to decidability results, we have for the first time characterized a family of $1\frac{1}{2}, 2\frac{1}{2}$ player STGs for whom the quantitative reachability is decidable. The use of exponential distributions is mandatory to get a closed form expression for the probability. It is unclear if this construction can be extended to some larger classes of STGs. Figure 9 in [9] shows an example of a two-clock $\frac{1}{2}$ player game for which the region abstraction fails to give any relevant information on the real "probabilistic" behaviour of the system (lack of so-called fairness); in particular it cannot be used for qualitative, and therefore quantitative, analysis

of reachability properties. The decidability of qualitative reachability in $1\frac{1}{2}, 2\frac{1}{2}$, multi-clock STG seems then hard due to the same problem of unfair runs.

References

1   S. Akshay, P. Bouyer, S. N. Krishna, L. Manasa, and A. Trivedi. Stochastic timed games revisited. `http://www.cse.iitb.ac.in/~krishnas/TR16.pdf`.

2   R. Alur and D.L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

3   E. Asarin, O. Maler, A. Pnueli, and J. Sifakis. Controller synthesis for timed automata. In *Proc. of IFAC Symposium on System Structure and Control*, pages 469–474. Elsevier, 1998.

4   C. Baier, P. Bouyer, T. Brihaye, and M. Größer. Almost-sure model checking of infinite paths in one-clock timed automata. In *Proc. 23rd Annual Symposium on Logic in Computer Science (LICS'08)*, pages 217–226. IEEE Computer Society Press, 2008.

5   C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(7):524–541, 2003.

6   C Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.

7   N. Bertrand, P. Bouyer, T. Brihaye, and P. Carlier. Analysing decisive stochastic processes. In *Proc. 43rd International Colloquium on Automata, Languages and Programming (IC-ALP'16) – Part II*, Leibniz International Proceedings in Informatics. Leibniz-Zentrum für Informatik, July 2016. To appear.

8   N. Bertrand, P. Bouyer, T. Brihaye, and N. Markey. Quantitative model-checking of one-clock timed automata under probabilistic semantics. In *Proc. 5th International Conference on Quantitative Evaluation of Systems (QEST'08)*. IEEE Computer Society Press, 2008.

9   N. Bertrand, P. Bouyer, T. Brihaye, Q. Menet, M. Größer, and M. Jurdziński. Stochastic timed automata. *Logical Methods in Computer Science*, 10(4):1–73, 2014.

10 N. Bertrand, T. Brihaye, and B. Genest. Deciding the value 1 problem for reachability in 1-clock decision stochastic timed automata. In *Proc. 11th International Conference on Quantitative Evaluation of Systems (QEST'14)*, pages 313–328. IEEE Computer Society Press, 2014.

11 H.C. Bohnenkamp, P.R. D'Argenio, H. Hermanns, and J.-P. Katoen. MODEST: A compositional modeling formalism for hard and softly timed systems. *IEEE Transactions on Software Engineering*, 32(10):812–830, 2006.

12 P. Bouyer and V. Forejt. Reachability in stochastic timed games. In *Proc. 36th International Colloquium on Automata, Languages and Programming (ICALP'09)*, volume 5556 of *LNCS*, pages 103–114. Springer, 2009.

13 Tomáš Brázdil, Jan Krčál, Jan Křetínský, and Vojtěch Řehák. Fixed-delay events in generalized semi-Markov processes revisited. In *Proc. 22nd International Conference on Concurrency Theory (CONCUR'11)*, volume 6901 of *LNCS*, pages 140–155. Springer, 2011.

14 T. Brihaye, G. Geeraerts, S. N. Krishna, L. Manasa, B. Monmege, and A. Trivedi. Adding negative prices to priced timed games. In *Proc. 25th International Conference on Concurrency Theory (CONCUR'14)*, LIPIcs, pages 560–575. Leibniz-Zentrum für Informatik, 2014.

15 J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.

16 V. Forejt, M. Kwiatkowska, G. Norman, and A. Trivedi. Expected reachability-time games. In *Proc. 8th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'10)*, volume 6246 of *LNCS*, pages 122–136. Springer, 2010.

17 M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Verifying quantitative properties of continuous probabilistic timed automata. In *Proc. of 11th International Conference on*

Concurrency Theorey, (CONCUR'00), volume 1877 of *LNCS*, pages 123–137. Springer, 2000.

18 M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282(1):101–150, June 2002.

19 F. Laroussinie, N. Markey, and P. Schnoebelen. Model checking timed automata with one or two clocks. In *Proc. 15th International Conference on Concurrency Theory (CONCUR'04)*, volume 3170 of *LNCS*, pages 387–401. Springer, 2004.

20 J. Ouaknine, A. Rabinovich, and J. Worrell. Time-bounded verification. In *Proc. 20th International Conference on Concurrency Theory (CONCUR'09)*, volume 5710 of *LNCS*, pages 496–510. Springer, 2009.

21 J. Ouaknine and J. Worrell. Towards a theory of time-bounded verification. In *Proc. 37th International Colloquium on Automata, Languages and Programming (ICALP'10)*, volume 6199 of *LNCS*, pages 22–37. Springer, 2010.

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

23 Uppaal case-studies. `http://www.it.uu.se/research/group/darts/uppaal/examples.shtml`.