

# Towards Probabilistic Strategic Timed CTL

## Extended Abstract

Wojciech Jamroga

Institute of Computer Science, PAS  
Warsaw, Poland  
Nicolaus Copernicus University  
Toruń, Poland  
jamroga@ipipan.waw.pl

Marta Kwiatkowska

Institute of Computer Science, PAS  
Warsaw, Poland  
University of Oxford  
Oxford, United Kingdom  
marta.kwiatkowska@cs.ox.ac.uk

Wojciech Penczek

Institute of Computer Science,  
Polish Academy of Sciences  
Warsaw, Poland  
penczek@ipipan.waw.pl

Laure Petrucci

LIPN, CNRS UMR 7030,  
Université Sorbonne Paris Nord  
Villetaneuse, France  
petrucci@lipn.univ-paris13.fr

Teofil Sidoruk

Institute of Computer Science,  
Polish Academy of Sciences  
Warsaw, Poland  
t.sidoruk@ipipan.waw.pl

### ABSTRACT

We define **PSTCTL**, a probabilistic variant of Strategic Timed CTL (**STCTL**), interpreted over stochastic multi-agent systems with continuous time and asynchronous execution semantics. **STCTL** extends **TCTL** with strategic operators in the style of **ATL**. Moreover, we demonstrate the feasibility of verification with irP-strategies.

### KEYWORDS

model checking; probabilistic systems; strategic ability; real time

#### ACM Reference Format:

Wojciech Jamroga, Marta Kwiatkowska, Wojciech Penczek, Laure Petrucci, and Teofil Sidoruk. 2026. Towards Probabilistic Strategic Timed CTL: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/OEXM4130>

## 1 INTRODUCTION

Strategic Timed Computation Tree Logic (**STCTL**) [5] extends the classical branching-time logic **CTL** [11, 12] in two important directions, firstly by adding discrete- or continuous-time representation, and secondly through enabling strategic reasoning in multi-agent systems. As such, it addresses the ongoing needs and challenges in formal verification, offering a specification language richer than comparable formalisms (i.e., Alternating-time Temporal Logic **ATL** [2, 3, 17, 24] and its timed variants [5, 30]) without creating an overhead in computational complexity. However, to adequately capture the diversity of real-world behaviours, which are subject to uncertainty, we need to consider probabilistic models and specifications. Hence, we investigate **PSTCTL**, a probabilistic variant of **STCTL**. On one hand, it constitutes a strategic extension of the well-known formalisms of **PCTL** and **PTCTL** [6, 9, 15]. In contrast to these two logics, typically considered in synchronous or turn-based settings, **PSTCTL** is interpreted over probabilistic continuous-time

asynchronous multi-agent systems (**PCAMAS**), allowing for reasoning about imperfect information strategies. On the other hand, **PSTCTL** is also a natural branching-time counterpart to the linear-time **PTATL** [22], and a missing component in the parallel line of research on branching-time frameworks. We demonstrate the feasibility of model checking of **PSTCTL** in practice. As opposed to [22], we focus on probabilistic (irP), rather than deterministic (irp), strategies of memoryless agents with imperfect information.

**Related Work.** **SCTL** and **STCTL** [5] augment, respectively, **CTL** [11, 12] and its (continuous or discrete) time extension **TCTL** [1] with the strategic modality. **ATL** [2, 3] and **TATL** [26, 30] are the corresponding linear-time formalisms based on **LTL** [33]. **ATL** was further developed in multiple directions, e.g. by the addition of epistemic reasoning [14, 18], strategy contexts [10, 29], or quantification of agents' uncertainty [35]. The closest linear-time counterpart to **PSTCTL** is the recently proposed **PTATL** [22], notably also in the imperfect information setting (which, due to significantly higher complexity of verification, has seldom been considered in conjunction with probabilistic logics and models). **PTATL** is in **2EXPTIME** for irP and in **EXPTIME** for irp-strategies, while untimed **PATL** (studied in [7, 8] also with imperfect information, but interpreted over synchronous models) is in **EXPTIME** and  $\Delta_2^P$ -complete for these respective semantics.

## 2 PRELIMINARIES

In this section, we recall the relevant theoretical background from [5, 22, 23]. Then, we define the syntax and semantics of **PSTCTL**.

**Agents.** Probabilistic Continuous-time Asynchronous Multi-Agent Systems (**PCAMAS**) [22] are networks of Probabilistic Timed Automata (**PTA**), whose components represent individual agents. Following the modeling tradition from the theory of concurrent systems, local (private) transitions are asynchronously interleaved, while synchronisations occur on actions shared by two or more agents [13, 23, 31]. However, the formalism also features an aspect typically present in synchronous frameworks, namely *protocols* that list actions available to agents. The product of components (or **PCAMAS model**) captures the system's global behaviour, and the *concrete model* adds continuous-time representation to the states.



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/OEXM4130>

**Strategies.** Conditional plans of an agent (in a coalition), dictating choices in each possible situation, are called (joint) *strategies* and usually classified [22, 34] based on the agents’ *state information*: perfect (I) vs. imperfect (i), *recall of state history*: perfect (R) vs. no recall (r), and *action selection*: probabilistic (P) vs. deterministic (p).

We focus on memoryless strategies without recall, i.e., of type irP. Formally, an irP-strategy for agent  $i$  is a function that maps from  $i$ ’s local states to a probability distribution over  $i$ ’s actions available in these states. Note that deterministic (irp) strategies are a special case of irP where only point distributions are considered.

**Logic.** The syntax of Probabilistic Strategic Timed CTL (PSTCTL):

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\gamma, \\ \gamma &::= \varphi \mid \neg\gamma \mid \gamma \wedge \gamma \mid \mathbb{P}^{\geq z}\gamma \cup_I \gamma \mid \mathbb{P}^{\geq z}\gamma R_I \gamma, \end{aligned}$$

includes the *strategic modality*  $\langle\langle A \rangle\rangle$  (expressing that agent coalition  $A$  has a strategy to enforce the property that follows), continuous-time intervals  $I \subseteq \mathbb{R}_{0+}$  (restricting the evaluation of temporal operators they are subscribed to), and the PCTL-style *probabilistic path operator*  $\mathbb{P}^{\geq z}$  (stating that the probability of taking a path that satisfies the property that follows is in relation  $\bowtie$  with constant  $z$ ).

Atomic propositions  $p$ , temporal operators  $\cup_I$ ,  $R_I$  (and derived  $F_I$ ,  $G_I$ ), and their Boolean combinations are defined as usual [5, 22]. Below, we give the semantic clause for the case unique to PSTCTL:

- $M, (s, v) \models \langle\langle A \rangle\rangle\mathbb{P}^{\geq z}\psi$  iff there exists a joint strategy  $\sigma_A$  such that for all  $\mu_M((s, v), \sigma_A) \in \text{out}_M((s, v), \sigma_A)$  we have  $\mu_M((s, v), \sigma_A)(\{\rho \mid M, \rho \models \psi\}) \bowtie z$ , where:
  - $M, \rho \models \gamma_1 \cup_I \gamma_2$  iff there is  $r \in I$  such that:  $M, \pi_\rho(r) \models \gamma_2$  and for all  $0 \leq r' < r$ :  $M, \pi_\rho(r') \models \gamma_1$ .
  - $M, \rho \models \gamma_1 R_I \gamma_2$  iff for all  $r \in I$ :  $M, \pi_\rho(r) \models \gamma_2$  or there is  $0 \leq r' < r$ :  $M, \pi_\rho(r') \models \gamma_1$ .

In the above,  $\text{out}_M((s, v), \sigma_A)$  denotes the *outcome* of strategy  $\sigma_A$  in concrete state  $(s, v)$ , i.e., the set of all probability distributions induced by executions from  $(s, v)$  consistent with  $\sigma_A$ . We refer the reader to [22] for formal definitions. Moreover,  $\pi_\rho$  denotes the dense path corresponding to an execution  $\rho$ , see [32].

### 3 EXPERIMENTS

In this section, we discuss practical model checking of PSTCTL with probabilistic strategies (PSTCTL<sub>irP</sub>). The approach is quite different from the algorithm for PTATL<sub>irP</sub>, which combined two verifiers: IMITATOR (handling irp-strategies and timing constraints) and PRISM (handling probabilistic constraints), see [22]. Since irP-strategies can choose different actions in each visit of a local state, the strategy encoding introduced in [22] is no longer suitable, and both time and probabilities can be handled by PRISM (where remembering the actions chosen for irp-strategies is not as easy).

To address this issue, for each local state of the coalition agent(s) in the input PCAMAS, we define one parameter per available action, representing the probability of selecting this particular action in the agent’s irP-strategy, see Fig. 1.

```
const double p, q, r; // parameters
module Controller
...
[] locC=0 -> p : (locC'=1) + q : (locC'=2) + r : (locC'=3) + 1-p-q-r : (locC'=4);
```

**Figure 1: Probabilistic choice in the PRISM model for  $n = 4$ .**

| n        | T=5                              | T=30  | T=100  |
|----------|----------------------------------|-------|--------|
| 1 train  | n/a (requires $n \geq 2$ trains) |       |        |
| 2 trains | 0.1 s                            | 0.1 s | 0.1 s  |
| 3 trains | 0.1 s                            | 0.7 s | 28.6 s |
| 4 trains | 0.4 s                            | 4.1 s | 493 s  |
| 5 trains | 4.1 s                            | 849 s | memout |

**Table 1: Results for the Train-Gate-Controller benchmark with  $n$  trains and time constraint  $T$  in the verified property  $\varphi = \langle\langle C \rangle\rangle\mathbb{P}^{\geq 0.8}F_{[0,T]}(\text{passed}_1 \wedge \text{passed}_2)$ .**

**Benchmark.** The classical Train-Gate-Controller (TGC) scenario [4, 16, 23] was modified to add a timing constraint  $T$  and scaled with the number of trains  $n$ . The  $i$ -th train takes  $n - i + 1$  time units to go through the tunnel, e.g.  $n$  for the first one and 1 for the last. The formula  $\varphi = \langle\langle C \rangle\rangle\mathbb{P}^{\geq 0.8}F_{[0,T]}(\text{passed}_1 \wedge \text{passed}_2)$  states that the controller  $C$  has a strategy to let the first two trains through the tunnel by time  $T$ , with a probability of at least 0.8. Clearly, this property is not satisfied by any irp-strategy, as opposed to irP-strategies that assign sufficient probabilities to actions selecting trains 1 and 2 in  $C$ ’s initial state.

Note that, although PRISM does not support parametric verification for probabilistic timed automata models, this issue can be circumvented by extracting the intermediate model generated by the digital clocks engine using the `-exportdigital` argument. Following [28], under suitable restrictions this model is a discrete-time Markov Decision Process (MDP) that preserves probabilistic reachability values and retains parameter declarations, so can be fed back to PRISM.

**Test Platform.** PRISM was executed in the Windows Subsystem for Linux (WSL) on a 4.0 GHz CPU (8 cores, 16 threads) with 64 GB of RAM. Results in Table 1 combine verification and model generation time (the latter consistently negligible); *memout* indicates termination of the process upon memory usage exceeding 32 GB (i.e., half of the total available, following the default WSL setting).

**Results.** While we were able to successfully verify a strategic property involving the probabilistic choices of an agent for the first time, it is also clear that the scalability of our initial approach is limited. For more than  $n = 5$  trains, PRISM no longer returns any synthesised probability values (or their ranges) for the defined parameters. As such, the cost of encoding probabilistic choices in this manner can be prohibitive for larger models. However, this is expected, given the theoretical complexity of PSTCTL verification, emphasising the need not only for dedicated tools implementing more efficient algorithms, but also complementary techniques such as approximations [21, 27], model abstractions [19, 20, 25], and partial order reduction [23].

### 4 CONCLUSIONS

We proposed PSTCTL, a probabilistic variant of strategic CTL with continuous time, and tackled for the first time the case of irP-strategies in practical model checking. There are several directions for further research, including formally establishing the expressive power and model checking complexity of PSTCTL<sub>irP</sub> and PSTCTL<sub>irP</sub>, as well as investigating the new angle of timed vs. untimed strategies.

## ACKNOWLEDGMENTS

This work was supported by CNRS IRP “Le Trójkąt”, by NCBR Poland & FNR Luxembourg under the PoLux/FNR-CORE project SpaceVote (POLLUX-XI/14/SpaceVote/2023), by the PHC Polonium project MoCca (BPN/BFR/2023/1/00045), and by the ANR-22-CE48-0012 project BISOUS. M. Kwiatkowska contributed while on sabbatical and acknowledges funding from the ERC under the European Union’s Horizon 2020 research and innovation programme (FUN2MODEL, grant agreement No. 834115).

## REFERENCES

- [1] Rajeev Alur, Costas Courcoubetis, and David L. Dill. 1993. Model-Checking in Dense Real-Time. *Inf. Comput.* 104, 1 (1993), 2–34.
- [2] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 1997. Alternating-Time Temporal Logic. In *Proceedings of FOCS’97*. IEEE Computer Society, 100–109.
- [3] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-Time Temporal Logic. *J. ACM* 49 (2002), 672–713.
- [4] Rajeev Alur, Thomas A. Henzinger, and Moshe Y. Vardi. 1993. Parametric Real-time Reasoning. In *Proceedings of STOC ’93* (San Diego, California, USA). ACM, 592–601.
- [5] Jaime Arias, Wojciech Jamroga, Wojciech Penczek, Laure Petrucci, and Teofil Sidoruk. 2023. Strategic (Timed) Computation Tree Logic. In *Proceedings of AAMAS’23*. ACM, 382–390.
- [6] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of Model Checking*. MIT Press.
- [7] Francesco Belardinelli, Wojciech Jamroga, Munyque Mittelmann, and Aniello Murano. 2023. Strategic Abilities of Forgetful Agents in Stochastic Environments. In *Proceedings of KR’23*. 726–731.
- [8] Francesco Belardinelli, Wojciech Jamroga, Munyque Mittelmann, and Aniello Murano. 2024. Verification of Stochastic Multi-Agent Systems with Forgetful Strategies. In *Proceedings of AAMAS’24*. ACM, 160–169.
- [9] Andrea Bianco and Luca de Alfaro. 1995. Model Checking of Probabilistic and Nondeterministic Systems. In *Proceedings of FSTTCS 1995 (Lecture Notes in Computer Science, Vol. 1026)*. Springer, 499–513.
- [10] Thomas Brihaye, Arnaud Da Costa Lopes, François Laroussinie, and Nicolas Markey. 2009. ATL with Strategy Contexts and Bounded Memory. In *Proceedings of LICS’09 (Lecture Notes in Computer Science, Vol. 5407)*. Springer, 92–106.
- [11] E.M. Clarke and E.A. Emerson. 1981. Design and Synthesis of Synchronization Skeletons Using Branching Time Temporal Logic. In *Proceedings of Logics of Programs Workshop (Lecture Notes in Computer Science, Vol. 131)*. 52–71.
- [12] Edmund M. Clarke, Orna Grumberg, Daniel Kroening, Doron A. Peled, and Helmut Veith. 2018. *Model Checking, 2nd Edition*. MIT Press.
- [13] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. 1995. *Reasoning about Knowledge*. MIT Press.
- [14] Dimitar P. Guelev and Catalin Dima. 2012. Epistemic ATL with Perfect Recall, Past and Strategy Contexts. In *Proceedings of CLIMA-XIII (Lecture Notes in Computer Science, Vol. 7486)*. Springer, 77–93.
- [15] Hans Hansson and Bengt Jonsson. 1994. A Logic for Reasoning about Time and Reliability. *Form. Asp. Comp.* 6, 5 (1994), 512–535.
- [16] W. van der Hoek and M. Wooldridge. 2002. Tractable Multiagent Planning for Epistemic Goals. In *Proceedings of AAMAS’02*. ACM Press, 1167–1174.
- [17] Wojciech Jamroga. 2015. *Logical Methods for Specification and Verification of Multi-Agent Systems*. ICS PAS Publishing House.
- [18] Wojciech Jamroga and Nils Bulling. 2011. Comparing Variants of Strategic Ability. In *Proceedings of IJCAI’11*. 252–257.
- [19] Wojciech Jamroga and Yan Kim. 2023. Practical Abstraction for Model Checking of Multi-Agent Systems. In *Proceedings of KR’23*. 384–394.
- [20] Wojciech Jamroga and Yan Kim. 2023. Practical Model Reductions for Verification of Multi-Agent Systems. In *Proceedings of IJCAI’23*. ijcai.org, 7135–7139.
- [21] Wojciech Jamroga, Michał Knapik, Damian Kurpiewski, and Łukasz Mikulski. 2019. Approximate Verification of Strategic Abilities under Imperfect Information. *Artificial Intelligence* 277 (2019).
- [22] Wojciech Jamroga, Marta Kwiatkowska, Wojciech Penczek, Laure Petrucci, and Teofil Sidoruk. 2025. Probabilistic Timed ATL. In *Proceedings of AAMAS’25*. ACM, 1051–1059.
- [23] Wojciech Jamroga, Wojciech Penczek, Teofil Sidoruk, Piotr Dembinski, and Antoni W. Mazurkiewicz. 2020. Towards Partial Order Reductions for Strategic Ability. *J. Artif. Intell. Res.* 68 (2020), 817–850.
- [24] Magdalena Kacprzak and Wojciech Penczek. 2005. Fully Symbolic Unbounded Model Checking for Alternating-time Temporal Logic. *Autonomous Agents and Multi-Agent Systems* 11, 1 (2005), 69–89.
- [25] Yan Kim, Wojciech Jamroga, Wojciech Penczek, and Laure Petrucci. 2025. Practical Abstractions for Model Checking Continuous-Time Multi-Agent Systems. In *Proceedings of AAMAS 2025*. ACM, 1117–1126.
- [26] Michał Knapik, Étienne André, Laure Petrucci, Wojciech Jamroga, and Wojciech Penczek. 2019. Timed ATL: Forget Memory, Just Count. *J. Artif. Intell. Res.* 66 (2019), 197–223.
- [27] Damian Kurpiewski, Wojciech Jamroga, and Yan Kim. 2025. Approximate Verification of Strategic Abilities under Imperfect Information Using Local Models. In *Proceedings IJCAI 2025*. ijcai.org, 143–151.
- [28] Marta Z. Kwiatkowska, Gethin Norman, David Parker, and Jeremy Sproston. 2006. Performance Analysis of Probabilistic Timed Automata Using Digital Clocks. *Formal Methods Syst. Des.* 29, 1 (2006), 33–78.
- [29] François Laroussinie and Nicolas Markey. 2015. Augmenting ATL with Strategy Contexts. *Inf. Comput.* 245 (2015), 98–123.
- [30] François Laroussinie, Nicolas Markey, and Ghassan Oreiby. 2006. Model-Checking Timed ATL for Durational Concurrent Game Structures. In *Proceedings of FOR-MATS’06 (Lecture Notes in Computer Science, Vol. 4202)*. Springer, 245–259.
- [31] Alessio Lomuscio, Wojciech Penczek, and Hongyang Qu. 2010. Partial Order Reductions for Model Checking Temporal-Epistemic Logics over Interleaved Multi-Agent Systems. *Fundam. Informaticae* 101, 1-2 (2010), 71–90.
- [32] Wojciech Penczek and Agata Pólrola. 2006. *Advances in Verification of Time Petri Nets and Timed Automata: A Temporal Logic Approach*. Studies in Computational Intelligence, Vol. 20. Springer.
- [33] Amir Pnueli. 1977. The Temporal Logic of Programs. In *Proceedings of FOCS’77*. 46–57.
- [34] Pierre-Yves Schobbens. 2004. Alternating-Time Logic with Imperfect Recall. In *Proceedings of LCMAS’03*, Vol. 85(2). 1–12.
- [35] Masoud Tabatabaei and Wojciech Jamroga. 2023. Playing to Learn, or to Keep Secret: Alternating-Time Logic Meets Information Theory. In *Proceedings of AAMAS’23*. ACM, 766–774.