




# Probabilistic Model Checking: Applications and Trends

Marta Kwiatkowska<sup>1</sup> , Gethin Norman<sup>1,2</sup> , and David Parker<sup>1</sup> () 

<sup>1</sup> Department of Computer Science, University of Oxford, Oxford, UK  
`{marta.kwiatkowska,david.parker}@cs.ox.ac.uk`

<sup>2</sup> School of Computing Science, University of Glasgow, Glasgow, UK  
`gethin.norman@glasgow.ac.uk`

**Abstract.** Probabilistic model checking is an approach to the formal modelling and analysis of stochastic systems. Over the past twenty five years, the number of different formalisms and techniques developed in this field has grown considerably, as has the range of problems to which it has been applied. In this paper, we identify the main application domains in which probabilistic model checking has proved valuable and discuss how these have evolved over time. We summarise the key strands of the underlying theory and technologies that have contributed to these advances, and highlight examples which illustrate the benefits that probabilistic model checking can bring. The aim is to inform potential users of these techniques and to guide future developments in the field.

## 1 Introduction

Probabilistic model checking [10] is a technique for the formal verification of stochastic systems. Properties to be verified are specified in temporal logic and then algorithmically checked against a model of the system. Some of the earliest work in the field dates from the 1980s [76,31], where algorithms were developed for computing the probability that linear temporal logic specifications are satisfied by sequential or concurrent probabilistic programs. The primary motivation was to establish the correctness of randomised algorithms, which had proved to be difficult, particularly in the context of concurrency.

Further verification techniques for these models, which are now typically referred to as discrete-time Markov chains (DTMCs) and Markov decision processes (MDPs), were soon developed. The now widely used temporal logic PCTL was proposed for DTMCs [42] and MDPs [18], expanding the range of properties that could be specified. Cited motivations included verifying the performance and reliability of, e.g., computer networks.

An extension of PCTL to the continuous-time setting was proposed in the late 1990s, called CSL [6]. This was designed for continuous-time Markov chains (CTMCs), a well established model for assessing performance and dependability

properties of computing and communication systems. Efficient model checking algorithms were also developed [8]. The further integration of costs and rewards, well-known notions from Markov modelling, into the probabilistic model checking framework led to additional temporal logics [44], allowing a modeller to reason about other quantitative characteristics such as power consumption. The flexibility of these formalisms meant that they applied equally well in non-traditional application domains such as biological systems.

For the model of MDPs, interest grew in synthesising “correct-by-construction” controllers or policies, using temporal logic to specify the desired behaviour of the system under control. Existing temporal logics and model checking algorithms were extended to incorporate various cost- and reward-based measures and multi-objective specifications [37]. This opened up new applications, including a variety of scheduling and planning problems. Probabilistic variants of the timed automata formalism, extending MDPs with clocks, allowed modelling of stochastic systems with real-time constraints and delays [64].

Fast forwarding to the present day, we observe a further significant expansion to the range of stochastic models for which temporal logics and model checking algorithms have been developed [57]: partially observable variants of MDPs, widely used in AI and planning, allow modelling of autonomous systems with unreliable sensors or of security protocols deploying secret keys; stochastic games provide reasoning about agents operating either competitively or collaboratively in probabilistic settings; and uncertain (or robust) Markov models formally capture epistemic uncertainty resulting, e.g., from data driven modelling.

To back up these advances in modelling formalisms and property specification languages, tool support has evolved and is now readily available. This comes in the form of both mature, general-purpose probabilistic model checkers, such as PRISM [56], Storm [47] or the Modest toolset [43], and various specialised verification tools [5]. A multitude of larger tool-chains and software frameworks have also been built, as probabilistic model checkers have become more robust and offer a variety of file formats and APIs with which to interface. The availability of this software, combined with the breadth of available formalisms, have seen probabilistic model checking used in a large and diverse range of applications, both within and outside computer science.

In this paper, we provide an overview of the broad range of problems to which probabilistic model checking has been applied, with a view to showcasing the current state of the field and giving some insight into its history. We identify a number of key areas in which the techniques have been particularly successful and consider how these have evolved over time, from the early work in the field to the present day. We also draw out the particular characteristics of the methods that make it well suited to each application area, and mention some of the key developments in theories and technologies that have been leveraged. Throughout, we highlight various case studies, illustrating the different ways in which probabilistic model checking can be deployed, the types of insight it can bring and the diversity of its usage. We hope that this paper therefore also serves as a guide to both current and potential users when selecting a formalism and analysis method for a given target application.

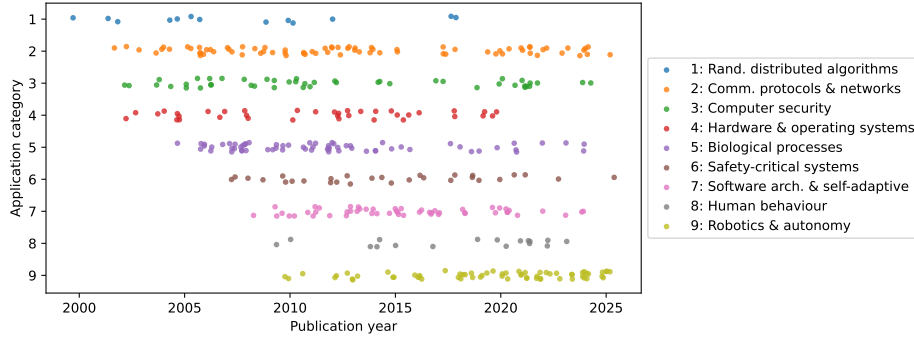


Fig. 1: Categorisation of approximately 400 application-oriented publications from the PRISM bibliography [81], illustrating how applications have evolved.

## 2 Applications of Probabilistic Model Checking

In the following sections we discuss, in loosely chronological order, a selection of key application areas for probabilistic model checking. For context, Figure 1 presents a graphical illustration of how these areas have evolved over time. For this visualisation, we considered application-oriented publications from the on-line bibliography of the PRISM model checker [81], in which many of the examples mentioned in this paper can be found. We categorise them according to the areas discussed below and show the publications that appeared in each year.

### 2.1 Randomised distributed algorithms

Motivated by the success of existing model checking techniques in identifying bugs and unexpected behaviour in concurrent systems, a fruitful initial area of study for probabilistic verification was *randomised distributed algorithms*. In this setting, the use of randomisation is a crucial tool for breaking symmetry between concurrent processes, in order to establish correct termination or fairness between participants. However, the subtle interaction between the parallelisation of processes and their probabilistic behaviour makes it challenging to formally reason about their correctness or evaluate their efficiency.

Probabilistic model checking was used to analyse the correctness and efficiency of various randomised distributed algorithms, including for consensus, byzantine agreement, leader election and self-stabilisation [62]. MDPs (and the closely related model of probabilistic automata) are ideally suited to modelling the mixture of stochasticity and concurrency that arises here. In some cases (e.g., fully synchronous settings), DTMCs are also a useful model.

The scalability of verification techniques to these applications was enhanced significantly by the development of *symbolic* implementations of probabilistic model checking, based on (multi-terminal) binary decision diagrams [7, 66]. While the size of the state spaces of the models grows rapidly, the high degree of regularity and the need for a mixture of both qualitative (“consensus is eventually

established with probability 1”) and quantitative (“what is the worst-case number of algorithm rounds required for termination?”) properties made these applications well placed to benefit from the advantages of symbolic model checking.

In a notable example, in 2005, McIver and Morgan studied a self-stabilization algorithm due to Herman [48], establishing results on its worst-case runtime [60]. They also made a conjecture about the initial configurations of the network that would yield the worst-case expected time to achieve stabilisation, and validated this empirically through probabilistic model checking. However, a proof turned out to be elusive: a series of subsequent papers gradually established tighter bounds until the conjecture was finally proved more than 10 years later [19].

## 2.2 Communications and networks

**Communication protocols.** A natural evolution from early work on distributed algorithms was to verify *communication protocols*. Again, these often deploy randomisation to break symmetry, for example in combination with exponential back-off schemes to prevent collisions between data transmissions; the Ethernet protocol is a classic example. Moreover, particularly in wireless communication settings, it is desirable for a formal analysis of a protocol’s reliability and efficiency to also incorporate a stochastic model of the underlying communication medium’s unreliability, e.g., message delays or losses.

As above, both DTMCs and MDPs are useful for modelling communication protocols, with the latter in particular allowing an analysis of the *worst-case* expected performance or runtime over any possible concurrent scheduling. Verification efforts in this area, for example of the CSMA/CA (carrier-sense multiple access with collision avoidance), FireWire and Zigbee protocols, also benefited from the development of *probabilistic timed automata* (PTAs) [64], providing scalable modelling and verification techniques for analysing probabilistic systems with time-outs and deadlines. Models of randomised protocols with complex interactions but a degree of regularity also proved to be amenable to symbolic model checking. A good example is the verification of the Bluetooth protocol [34], which built and analysed multiple DTMC models with more than  $10^{10}$  states in order to perform a detailed analysis of the impact of differences between variants of the protocol specification, and of the underlying model assumptions used.

Interest in this application area has continued to this day, with a wide variety of more modern wireless protocols being verified, for example, in the context of ad-hoc or vehicular networks. There has been particular interest in wireless sensor networks, see e.g., [24], where the small scale of the hardware deployed makes analysing unreliability particularly important.

**Computer and communication networks.** More broadly, probabilistic model checking has been used to formally evaluate the effectiveness of designs for communication and computer networks. PTA models, in particular, have been shown to be useful for quantifying the reliability or timeliness of transmissions across a network in the context of uncertain delays or message losses caused by individual network components. Examples include exploring the effectiveness of

publish-subscribe messaging system [46], optimising the quality of networked automation systems for control problems [41] and comparing wireless token-passing schemes for networks in safety-critical systems [33].

This application domain also showcases the benefits of another important model class: *continuous-time Markov chains* (CTMCs). These provide an alternative model of stochastic systems in which the delays between each event are represented by negative exponential distributions. These can be very effective in modelling the timing of, for example, the rates with which jobs arrive or are serviced in a network queue, or with which failures occur in hardware components.

Outside formal verification, the well established fields of *performance evaluation* and *performability* use CTMCs to model and evaluate the performance and dependability of computer and communication networks (as well as others, such as manufacturing systems). The coming together of these field with probabilistic model checking [9] showed that temporal logics for CTMCs like CSL and its extensions offer a powerful means of formally specifying a wide range of important properties across both performance and reliability (e.g., “the probability of a server response taking more than 1 second is at most 0.02”, “in the long-run, the availability of the network is at least 98%” or “the mean time to failure is at least 200 hours”). Example applications include comparing the quality-of-service of alternative approaches to traffic shaping in wireless networks [14] and assessing the dependability of novel topologies for optical networks [72].

### 2.3 Computer security

Computer security is another area where formal verification has long been an important tool, due to the need for strong guarantees on the resilience of systems to attack and the existence of unexpected weaknesses caused by subtleties in protocol designs. Probability is again a crucial modelling tool since the use of randomisation is widespread, e.g., for the generation of keys or session identifiers, or to prevent buffer overflows or DNS cache poisoning attacks.

It is natural to use MDPs for the analysis of security systems, with nondeterminism representing the potential actions of an adversary and probability used to model randomisation or other stochastic aspects. The work of [74] represents an early example of using MDP policy synthesis within probabilistic model checking. It generates optimal PIN block attacks, which are unexpected sequences of interactions with an API intended to determine the value of a PIN.

This approach can be generalised by using *stochastic games*, which provide formal modelling of the behaviour of multiple agents with differing objectives, such as an attacker and defender in a secure network. Applications of stochastic game model checking in this area include performing threat analysis for real-world scenarios modelling using information from technical knowledge bases [75] and verifying the resilience of collective adaptive systems to attacks [40].

Other types of models are also widely used. For example, [2] uses CTMC model checking to study the well-known Kaminsky DNS cache-poisoning attack, performing a detailed analysis of the efficacy of randomisation-based fixes that

have been proposed to counter the attack. Furthermore, in many security systems, there is an inherent trade-off between security and performance, in which case the expressivity and flexibility offered by probabilistic temporal logics is valuable. A good example is the work in [15], which performs a detailed analysis of a certified e-mail message delivery protocol, analysing in particular the relation between error rate and transmission cost.

## 2.4 Biological processes

CTMCs have other important modelling applications beyond their usage for performance or reliability evaluation discussed above. A notable example is for modelling the stochastic dynamics of *biological processes*, from the level of cellular reactions, such as protein-protein interactions or gene expression, up to population-level models, for example of disease spreading.

This has proved to be a popular application of probabilistic model checking and has sparked adoption and collaboration with fields such as computational and systems biology. Many biological systems are relatively straightforward to represent in the PRISM modelling language and automated translations have been developed from custom languages [50] or process algebras [29].

As with other uses of CTMCs, the temporal logic CSL and its extensions are expressive enough to specify key properties of interest for biological systems (e.g., “what is the expected number of molecules of protein  $X$  after 5 minutes?” or “what is the probability that more than 10% of the population is infected within 2 weeks?”). However, logics such as LTL have also been used to characterise more complex temporal aspects such as oscillatory behaviour; see, e.g., [13].

From a computational perspective, verification tools implement various efficient CTMC solution methods. However, this area, in particular, has spurred the development of *statistical model checking* [79], which uses discrete-event simulation to provide approximate results to verification queries along with statistical guarantees as to their accuracy. This class of methods works particularly well when, as here, systems are modelled as Markov chains (i.e., without nondeterminism) and properties of interest are often over a finite time horizon.

In contrast to many other applications of probabilistic model checking, which feature human-engineered systems, the mechanisms underlying biological processes are often poorly understood, and these techniques provide a means to validate a hypothesised model. One example is an investigation of the impact of components of the fibroblast growth factor (FGF) signalling pathway [55], with the outputs of verification later validated against laboratory experiments. Other examples include an analysis of the efficacy of hyperthermia treatment [70], which corroborated prior beliefs as to the effectiveness of combining cancer treatment strategies, and a study of concentration-based navigation of sperm cells [53], again validating previous experimental observations. Similarly, probabilistic model checking can support *synthetic biology*, in which biological circuits are engineered; a good example here is DNA computing [58].

## 2.5 Safety-critical systems

Formal verification can be particularly beneficial for *safety-critical systems*, i.e., those where failures could result in death, serious injury or other significant damage. Probabilistic model checking can be used to rigorously quantify the likelihood with which high-risk failures or malfunctions could occur. This can be used for the assessment of *safety integrity levels* (SILs), which specify acceptable probabilities or rates of failure, and which underlie various technical standards that are either recommended or mandated for safety-critical systems.

Typically, systems are modelled as DTMCs or CTMCs, with the latter being more common if the model is being used to quantify the *rate* of failure over a specified time period, rather than the probability of a failure at any point. Safety specifications typically amount to relatively simple PCTL or CSL queries, although the analysis can still be expensive if probabilities are very low or the model comprises a large number of components. It is worth noting that probabilistic verification tools typically feature a variety of different methods for computing probabilities, trading off the accuracy required (which can be high in this setting) against computational cost [21].

An early example of the use of probabilistic model checking in this area was its integration into the *failure mode and effects* analysis of an airbag system [3]. Two contrasting designs were checked against relevant safety standards and the precise causes of safety violations were then identified. Connections have also been established between probabilistic model checking and *fault tree analysis* [77] and applied for example to reliability analysis of railway infrastructure [78] and safety analysis of vehicle guidance systems [39].

Other safety-critical applications of probabilistic model checking include those from *medical* and *aerospace* domains: verifying pacemaker designs by integrating probabilistic models of heart behaviour and quantifying the resulting probability of correct device operation [26]; analysing the reliability, availability and maintainability of satellite systems [49]; and verifying the reliability of spacecraft designs [28]. We also discuss below (in Section 2.7) other applications that may be safety-critical, such as autonomous driving and human-robot interaction.

## 2.6 Hardware and operating systems

Many of the instances of probabilistic verification surveyed above model systems at a relatively high level of abstraction, e.g., the interactions between participants in a protocol or components in a network. However, formal probabilistic modelling also provides value at the *hardware* or *operating system* level.

One application is to verify that hardware circuits function reliably even in the presence of failures of individual system components. This typically requires model checking of a DTMC in a fashion not dissimilar to that described above for quantifying risk in safety-critical systems. An early illustration of this was the verification of NAND multiplexing [63], a fault-tolerant design motivated by manufacturing defects found in the field of nanotechnology. Probabilistic model



checking illustrated the impact of slight variations in component reliability and, in doing so, identified a flaw in an earlier analysis of the design.

A much broader range of system properties than just reliability can also be studied, including those relating to *timing*, *power usage* and *energy consumption*. This makes probabilistic model checking a powerful tool to explore design spaces and investigate trade-offs between competing characteristics. For example, more recently, work in [68] presented a detailed illustration of the integration of probabilistic model checking into the early design process of circuit designs using reconfigurable transistors. They computed guarantees on delay, power dissipation and energy consumption per operation, including various properties that would be difficult to analyse with simulation.

A similar spectrum of quantitative properties can be analysed in operating systems applications. For example, [71] uses MDPs for thermal modelling of multi-core systems and then analyses performance-reliability trade-offs. This is done using probabilistic model checking with *energy-utility quantiles* [12,11], which incorporate conditional probabilities and ratio constraints between cost and reward measures to investigate the interplay between multiple objectives.

## 2.7 Robotics and autonomous systems

In recent years, a clear trend within the use of probabilistic model checking has been a shift towards synthesising “correct-by-construction” controllers or policies for a system, based on a formal specification of its desired behaviour. Although this differs from the more classical approach of verifying a fixed system model to determine whether it satisfies a specification, probabilistic model checking is already well suited to this task, notably through the use of MDPs, with temporal logic used for controller specifications. In parallel, there has been a growing synergy between verification and fields of artificial intelligence such as planning and reinforcement learning, which often also use MDPs.

An application domain where these shifts have been particularly apparent is *robotics and autonomous systems*. The use of stochastic modelling is essential here since robots frequently operate in uncertain and dynamic environments, due to the presence of humans or other unknown obstacles, and imperfect or unreliable sensors and actuators. Probabilistic model checking offers a wealth of ways to formally specify desired behaviour in this setting, and the means to provide a probabilistic guarantee on the behaviour of generated policies. These can be critical for safety purposes, e.g., for robots operating in the presence of humans or high risk environments, or to guarantee performance and reliability sometimes under tight resource (e.g., battery) constraints.

The use of linear temporal logic (LTL) is common here, in order to capture more complex temporal specifications (e.g., “maximise the probability of inspecting all three sensors, in any order, whilst avoiding  $X$ ”). This is often combined with multi-objective model checking [37], which allows the generation of policies that trade-off between several, competing objectives (e.g., battery life vs. mission execution time) or analysis of the corresponding Pareto front.



Robotics provides good examples of probabilistic model checking being embedded in real-world systems. For example, [45] describes a major *long-term autonomy* project deploying mobile robots in everyday environments such as offices and care homes. Multi-objective model checking on MDPs is repeatedly applied within the robots’ control software to perform tasks effectively and reliably as the map of the environment is updated over time. In other recent work, probabilistic model checking is used within a planning system for *autonomous underwater vehicles* (AUVs) to retrieve data from sensor networks [20]. The resulting policies are deployed in a real-world trial and shown to outperform existing hand-designed policies.

Robotics and autonomous systems applications also highlight further challenges arising in verification today. One is the integration of components that deploy *machine learning*. In [22], a probabilistic model checking framework is proposed that incorporates deep-learning perception techniques. The uncertainty arising from the learning-driven components is factored into the verification of the overall system; this is applied to controllers for mobile robots and a driver-attentiveness management system for shared-control autonomous cars.

## 2.8 Software architectures and self-adaptive systems

The performance and reliability of *software architectures* has also been investigated quite extensively with probabilistic model checking. In similar fashion to other applications already discussed, such as computer and communication networks or hardware designs, formal modelling and analysis of Markov chains can be used to rigorously establish how the likelihood of failures or delays in individual components impacts the overall reliability and performance of a complex system of components. In an early illustration of this, researchers from ABB used DTMC models to evaluate the reliability of a large-scale industrial control system with more than 100 components [52].

In recent years, there has been particular interest in the modelling and verification of *cloud computing* services, where there is a need to maintain high Quality-of-Service levels, even under variable and unpredictable workloads. Moreover, it is often essential to meet precisely specified *service level agreements*, relating to response time or availability, which can be conveniently expressed and verified using probabilistic model checking; see e.g., [54]. Perhaps more importantly, these techniques can also support run-time decision making in cloud systems such as for auto-scaling (adjusting resource levels to meet demand) and load balancing. As discussed for autonomous systems above, MDP-based model checking provides an effective way to do this, particular when reasoning about trade-offs between different objectives or metrics; see e.g., [65].

More generally, *self-adaptive* systems, which monitor their environment and adapt their behaviour accordingly at run-time, have also proved to be a popular application domain for these techniques [61]. In fact, a variety of more advanced methods have been used effectively here. This includes model checking of *stochastic games*, which have been used to model worst-case assumptions about the environment of the cloud system [23]. Another example is the use of

*parametric model checking* [51], which can in this context be used to apply a form of sensitivity analysis, quantifying the impact that unknown system parameters have on overall performance metrics [1]. Lastly, methods for *families* of probabilistic models [27] have been used to efficiently verify multiple configurations supported by a self-adaptive system [67].

## 2.9 Human behaviour modelling

The modelling and analysis of *human behaviour* represents another challenge for formal verification, and one where probabilistic techniques play a key role. An example is the evaluation of diabetes patients’ behaviour when using insulin pumps [25]. Machine learning is used to extract representative patient behavioural patterns from a clinical dataset, and probabilistic model checking is deployed to analyse how different behavioral patterns impact an individual’s glucose physiology. The results demonstrate that switching behaviour types can significantly improve a patient’s glycemic control outcomes, boosting the effectiveness of diabetes patient education and peer support. The work of [80] uses DTMCs and probabilistic model checking, together with a cognitive reliability and error analysis method to transform expert estimates of relevant environmental and cognitive factors into human error rates, to assess the reliability of the procedures of a pharmacy and also analyse the effects of potential alternatives.

Another interesting application is the analysis of user interactive systems [4]. The approach is based on first inferring DTMC models of users’ activity patterns by applying machine learning to logged user traces. Probabilistic model checking is then applied to the DTMCs to express hypotheses about user behaviour and relationships within and between the activity patterns. The approach has been applied to a real-world case study of a deployed app with thousands of users and the analysis performed revealed insights into real-life app usage.

Other applications to human modelling include air traffic control [69], controller synthesis of UAVs interacting with human operators [38], human in the loop self-adaptive systems [59], and the design of correct-by-construction Advanced Driver Assistance Systems (ADAS) interacting with a human driver model built using the cognitive architecture ACT-R [35].

## 2.10 Further application areas

In the above, we have identified some of the most well-studied categories of applications for probabilistic model checking. This is of course non-exhaustive, and various other applications have been considered multiple times. These include: *smart grids*, with a focus on performance, resilience or security; *quantum computing*, for example, key distribution algorithms; *blockchain*-based systems, such as bitcoin; and *business process* modelling. See [81] for details.

We conclude this discussion by highlighting a few of the more diverse applications that have appeared in recent years, adding support for our argument that the techniques have very broad applicability.

**Sports tactics.** One example is the work in [30], which develops a framework for reasoning about effectiveness of team strategies in professional football. The first phase is to learn an MDP model, from event stream data, capturing the probability of moving between areas of the pitch and executing actions, such as passes or shots. Then probabilistic model checking is applied, with LTL used to express outcomes of interest. The analysis yields insights, such as when and where passing and shooting is more effective.

**Court interactions.** Another example is an analysis of interactions between participants in cases in the US Supreme Court [32]. Based on a dataset of court transcripts, a DTMC model is constructed that models the dynamics of interactions between people over the course of a trial. Probabilistic model checking is then used to analyse a wide range of properties, from the expected timing of court processes, to the likelihood of various events, such as the decisions taken by judges, and various trends are identified.

Other interesting applications include estimating political affinities through opinion diffusion on Twitter [73], studying the impact of regulations on the performance of public transport [17], planning railway infrastructure through capacity evaluation [36] and assessing graphical user interfaces [16].

### 3 Conclusions and Outlook

We have taken a retrospective look at some of the successful applications of probabilistic model checking over the last 25 years, identifying commonly studied application domains and discussing why and how they are amenable to this approach. Commonalities exist between the different problem domains, but there is considerable breadth and diversity to the set of applications.

Broadly speaking, we observe that this is largely because probabilistic model checking offers ease of rigorous modelling for many different types of stochastic behaviour, combined with extremely flexible logical formalisms to specify their quantitative characteristics. These go well beyond the classical notions of correctness typically used in formal verification. This is backed up by a wide range of flexible analysis techniques and stable tool support. These tools are continually evolving, incorporating new advances in modelling and verification techniques, which we anticipate leading to further application domains in the future.

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

### References

1. Alasmari, N., Calinescu, R., Paterson, C., Mirandola, R.: Quantitative verification with adaptive uncertainty reduction. *Journal of Systems and Software* **188**, 111275 (2022)

2. Alexiou, N., Basagiannis, S., Katsaros, P., Deshpande, T., Smolka, S.A.: Formal analysis of the Kaminsky DNS cache-poisoning attack using probabilistic model checking. In: Proc. 9th IEEE International Symposium on High Assurance Systems Engineering (HASE'10). pp. 94–103. IEEE Computer Society (2010)
3. Aljazzar, H., Fischer, M., Grunske, L., Kuntz, M., Leitner, F., Leue, S.: Safety analysis of an airbag system using probabilistic FMEA and probabilistic counter examples. In: Proc. 6th International Conference on Quantitative Evaluation of Systems (QEST'09). pp. 299–308. IEEE Computer Society (2009)
4. Andrei, O., Calder, M., Higgs, M., Girolami, M.: Probabilistic model checking of DTMC models of user activity patterns. In: Proc. 11th International Conference on Quantitative Evaluation of Systems (QEST'14). LNCS, vol. 8657, pp. 138–153. Springer (2014)
5. Andriushchenko, R., Bork, A., Budde, C.E., Ceska, M., Grover, K., Hahn, E.M., Hartmanns, A., Israelsen, B., Jansen, N., Jeppson, J., Junges, S., Köhl, M.A., Könighofer, B., Kretínský, J., Meggendorfer, T., Parker, D., Pranger, S., Quatmann, T., Ruijters, E., Taylor, L., Volk, M., Weininger, M., Zhang, Z.: Tools at the frontiers of quantitative verification. In: TOOLympics Challenge 2023. pp. 90–146. Springer (2024)
6. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.: Verifying continuous time Markov chains. In: Proc. 8th International Conference on Computer Aided Verification (CAV'96). LNCS, vol. 1102, pp. 269–276. Springer (1996)
7. Baier, C., Clarke, E., Hartonas-Garmhausen, V., Kwiatkowska, M., Ryan, M.: Symbolic model checking for probabilistic processes. In: Proc. 24th International Colloquium on Automata, Languages and Programming (ICALP'97). LNCS, vol. 1256, pp. 430–440. Springer (1997)
8. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering* **29**(6), 524–541 (2003)
9. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Performance evaluation and model checking join forces. *Communications of the ACM* **53**(9), 76–85 (2010)
10. Baier, C., de Alfaro, L., Forejt, V., Kwiatkowska, M.: Model checking probabilistic systems. In: Handbook of Model Checking, pp. 963–999. Springer (2018)
11. Baier, C., Daum, M., Dubslaff, C., Klein, J., Klüppelholz, S.: Energy-utility quantiles. In: NASA Formal Methods. LNCS, vol. 8430, pp. 285–299. Springer (2014)
12. Baier, C., Dubslaff, C., Klüppelholz, S.: Trade-off analysis meets probabilistic model checking. In: Proc. 23rd International Conference on Computer Science Logic and the 29th Symposium on Logic In Computer Science (CSL-LICS'14). pp. 1:1–1:10. ACM (2014)
13. Ballarini, P., Guerriero, M.: Query-based verification of qualitative trends and oscillations in biochemical systems. *Theoretical Computer Science* **411**(20), 2019–2036 (2010)
14. Ballarini, P., Ben-Othman, J., Mokdad, L.: Quantitative verification of WiMAX traffic shaping solutions. In: Proc. 7th International Symposium on Intelligent System Techniques for Ad hoc and Wireless Sensor Networks (IST-AWSN) (2012)
15. Basagiannis, S., Petridou, S.G., Alexiou, N., Papadimitriou, G.I., Katsaros, P.: Quantitative analysis of a certified e-mail protocol in mobile environments: A probabilistic model checking approach. *Computer Security* **30**(4), 257–272 (2011)
16. Bertolini, C., Mota, A.: Using probabilistic model checking to evaluate GUI testing techniques. In: SEFM. pp. 115–124. IEEE Computer Society (2009)

17. Bertrand, N., Bordais, B., Hérouët, L., Mari, T., Parreaux, J., Sankur, O.: Performance evaluation of metro regulations using probabilistic model-checking. In: *RSSRail. LNCS*, vol. 11495, pp. 59–76. Springer (2019)
18. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: *Proc. 15th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'95). LNCS*, vol. 1026, pp. 499–513. Springer (1995)
19. Bruna, M., Grigore, R., Kiefer, S., Ouaknine, J., Worrell, J.: Proving the Herman-protocol conjecture. In: *Proc. 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016). LIPIcs*, vol. 55, pp. 104:1–104:12 (2016)
20. Budd, M., Salavasidis, G., Karnarudzaman, I., Harris, C.A., Phillips, A.B., Duckworth, P., Hawes, N., Lacerda, B.: Probabilistic planning for AUV data harvesting from smart underwater sensor networks. In: *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'22)*. pp. 12051–12057 (2022)
21. Budde, C.E., Hartmanns, A., Klauck, M., Křetínský, J., Parker, D., Quatmann, T., Turrini, A., Zhang, Z.: On correctness, precision, and performance in quantitative verification: QComp 2020 competition report. In: *Proc. 9th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'20). LNCS*, vol. 12479, pp. 216–241. Springer (2020)
22. Calinescu, R., Imrie, C., Mangal, R., Rodrigues, G.N., Pasareanu, C.S., Santana, M.A., Vázquez, G.: Controller synthesis for autonomous systems with deep-learning perception components. *IEEE Transactions on Software Engineering* **50**(6) (2024)
23. Cámara, J., Garlan, D., Schmerl, B.R., Pandey, A.: Optimal planning for architecture-based self-adaptation via model checking of stochastic games. In: *Proc. 30th ACM Symposium on Applied Computing (SAC'15)*. pp. 428–435. ACM (2015)
24. Chen, M., Mokdad, L., Ben-Othman, J., Fourneau, J.M.: Probabilistic performance evaluation of the class-A device in LoRaWAN protocol on the MAC layer. *Performance Evaluation* **166**(C) (2024)
25. Chen, S., Feng, L., Rickels, M.R., Peleckis, A., Sokolsky, O., Lee, I.: A data-driven behavior modeling and analysis framework for diabetic patients on insulin pumps. In: *International Conference on Healthcare Informatics*. pp. 213–222 (2015)
26. Chen, T., Diciolla, M., Kwiatkowska, M.Z., Mereacre, A.: Quantitative verification of implantable cardiac pacemakers. In: *Proc. 33rd IEEE Real-Time Systems Symposium (RTSS'12)*. pp. 263–272 (2012)
27. Chrszon, P., Dubslaff, C., Klüppelholz, S., Baier, C.: Profeat: feature-oriented engineering for family-based probabilistic model checking. *Formal Aspects of Computing* **30**(1), 45–75 (2018)
28. Chrszon, P., Maurer, P., Saleip, G., Müller, S., Fischer, P.M., Gerndt, A., Felderer, M.: Model checking of spacecraft operational designs: a scalability analysis. *Software and Systems Modeling* (2025)
29. Ciocchetta, F., Hillston, J.: Bio-PEPA for epidemiological models. In: *Proc. 4th International Workshop on Practical Applications of Stochastic Modelling (PASM'09). ENTCS*, vol. 261, pp. 43–69. Elsevier (2008)
30. Clijmans, J., Roy, M., Davis, J.: Looking beyond the past: Analyzing the intrinsic playing style of soccer teams. In: *Machine Learning and Knowledge Discovery in Databases. LNCS*, vol. 13718, pp. 370–385 (2023)
31. Courcoubetis, C., Yannakakis, M.: Verifying temporal properties of finite state probabilistic programs. In: *Proc. 29th Symposium on Foundations of Computer Science (FOCS'88)*. pp. 338–345. IEEE CS Press (1988)

32. Das, S., Sharma, A.: Probabilistic model checking of temporal interaction dynamics in the supreme court. In: Proc. 11th International Symposium From Data to Models and Back (DataMod'23) (2023)
33. Dombrowski, C., Junges, S., Katoen, J.P., Gross, J.: Model-checking assisted protocol design for ultra-reliable low-latency wireless networks. In: Proc. IEEE 35th Symposium on Reliable Distributed Systems (SRDS'16). pp. 307–316 (2016)
34. Duflo, M., Kwiatkowska, M., Norman, G., Parker, D.: A formal analysis of Bluetooth device discovery. *International Journal on Software Tools for Technology Transfer* **8**(6), 621–632 (2006)
35. Eiras, F., Lahijanian, M., Kwiatkowska, M.: Correct-by-construction advanced driver assistance systems based on a cognitive architecture. In: Proc. IEEE 2nd Connected and Automated Vehicles Symposium (CAVS'19). pp. 1–7 (2019)
36. Emunds, T., Nießen, N.: Evaluating railway junction infrastructure: A queueing-based, timetable-independent analysis. *Transportation Research Part C: Emerging Technologies* **165** (2024)
37. Etessami, K., Kwiatkowska, M., Vardi, M., Yannakakis, M.: Multi-objective model checking of Markov decision processes. *Logical Methods in Computer Science* **4**(4), 1–21 (2008)
38. Feng, L., Wilsche, C., Humphrey, L., Topcu, U.: Controller synthesis for autonomous systems interacting with human operators. In: Proc. ACM/IEEE 6th International Conference on Cyber-Physical Systems (ICCPS'15). pp. 70–79. ACM (2015)
39. Ghadhab, M., Junges, S., Katoen, J., Kuntz, M., Volk, M.: Model-based safety analysis for vehicle guidance systems. In: Proc. 36th International Conference on Computer Safety, Reliability and Security (SAFECOMP'17). LNCS, vol. 10488, pp. 3–19. Springer (2017)
40. Glazier, T.J., Cámara, J., Schmerl, B.R., Garlan, D.: Analyzing resilience properties of different topologies of collective adaptive systems. In: SASO Workshops. pp. 55–60. IEEE Computer Society (2015)
41. Greifeneder, J., Frey, J.: Optimizing quality of control in networked automation systems using probabilistic models. In: Proc. 11th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'06). pp. 372–379 (2006)
42. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Aspects of Computing* **6**(5), 512–535 (1994)
43. Hartmanns, A., Hermanns, H.: The Modest toolset: An integrated environment for quantitative modelling and verification. In: Proc. 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'14). LNCS, vol. 8413, pp. 593–598. Springer (2014)
44. Haverkort, B., Cloth, L., Hermanns, H., Katoen, J.P., Baier, C.: Model checking performability properties. In: Proc. International Conference on Dependable Systems and Networks (DSN'02). IEEE CS Press (2002)
45. Hawes, N., Burbridge, C., Jovan, F., Kunze, L., Lacerda, B., Mudrová, L., Young, J., Wyatt, J.L., Hebesberger, D., Körtner, T., Ambrus, R., Bore, N., Folkesson, J., Jensfelt, P., Beyer, L., Hermans, A., Leibe, B., Aldoma, A., Faulhammer, T., Zillich, M., Vincze, M., Chinellato, E., Al-Omari, M., Duckworth, P., Gatsoulis, Y., Hogg, D.C., Cohn, A.G., Dondrup, C., Fentanes, J.P., Krajník, T., Santos, J.M., Duckett, T., Hanheide, M.: The STRANDS project: Long-term autonomy in everyday environments. *IEEE Robotics Autom. Mag.* **24**(3), 146–156 (2017)
46. He, F., Baresi, L., Ghezzi, C., Spoletini, P.: Formal analysis of publish-subscribe systems by probabilistic timed automata. In: Proc. Formal Techniques for Net-

- worked and Distributed Systems (FORTE'07). LNCS, vol. 4574, pp. 247–262. Springer (2007)
47. Hensel, C., Junges, S., Katoen, J.P., Quatmann, T., Volk, M.: The probabilistic model checker Storm. *International Journal on Software Tools for Technology Transfer* **24**(4), 589–610 (2022)
48. Herman, T.: Probabilistic self-stabilization. *Information Processing Letters* **35**(2), 63–67 (1990)
49. Hoque, K.A., Mohamed, O.A., Savaria, Y.: Towards an accurate reliability, availability and maintainability analysis approach for satellite systems based on probabilistic model checking. In: *Proc. Design, Automation & Test in Europe Conference & Exhibition (DATE'15)*. pp. 1635–1640. ACM (2015)
50. Hucka, M., Finney, A., Sauro, H., Bolouri, H., Doyle, J., Kitano, H., Arkin, A., Bornstein, B., Bray, D., Cornish-Bowden, A., Cuellar, A., Dronov, S., Gilles, E., Ginkel, M., Gor, V., Goryanin, I., Hedley, W., Hodgman, T., Hofmeyr, J., Hunter, P., Juty, N., Kasberger, J., Kremling, A., Kummer, U., Le Novère, N., Loew, L., Lucio, D., P, M., Minch, E., Mjolsness, E., Nakayama, Y., Nelson, M., Nielsen, P., Sakurada, T., Schaff, J., Shapiro, B., Shimizu, T., Spence, H., Stelling, J., Takahashi, K., Tomita, M., Wagner, J., Wang, J., SBML Forum: The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models. *Bioinformatics* **1**(19), 524–531 (2003)
51. Junges, S., Ábrahám, E., Hensel, C., Jansen, N., Katoen, J., Quatmann, T., Volk, M.: Parameter synthesis for Markov models: covering the parameter space. *Formal Methods System Design* **62**(1), 181–259 (2024)
52. Koziolok, H., Schlich, B., Bilich, C.: A large-scale industrial case study on architecture-based software reliability analysis. In: *Proc. IEEE 21st International Symposium on Software Reliability Engineering (ISSRE'10)*. pp. 279–288 (2010)
53. Kromer, J.A., Märcker, S., Lange, S., Baier, C., Friedrich, B.M.: Decision making improves sperm chemotaxis in the presence of noise. *PLoS Comput. Biol.* **14**(4) (2018)
54. Krotsiani, M., Kloukinas, C., Spanoudakis, G.: Validation of service level agreements using probabilistic model checking. In: *Proc. 14th IEEE International Conference on Services Computing (SCC'17)*. pp. 148–155 (2017)
55. Kwiatkowska, M., Heath, J.: Biological pathways as communicating computer systems. *Journal of Cell Science* **122**(16), 2793–2800 (2009)
56. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*. LNCS, vol. 6806, pp. 585–591. Springer (2011)
57. Kwiatkowska, M., Norman, G., Parker, D.: Probabilistic model checking and autonomy. *Annual Review of Control, Robotics, and Autonomous Systems* **5**, 385–410 (2022)
58. Lakin, M., Parker, D., Cardelli, L., Kwiatkowska, M., Phillips, A.: Design and analysis of DNA strand displacement devices using probabilistic model checking. *Journal of the Royal Society Interface* **9**(72), 1470–1485 (2012)
59. Li, N., Adepu, S., Kang, E., Garlan, D.: Explanations for human-on-the-loop: a probabilistic model checking approach. In: *Proc. IEEE/ACM 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS'20)*. pp. 181–187. ACM (2020)
60. McIver, A., Morgan, C.: An elementary proof that Herman's ring is  $\Theta(N^2)$ . *Information Processing Letters* **94**(2), 79–84 (2005)



61. Moreno, G.A., Cámara, J., Garlan, D., Schmerl, B.R.: Proactive self-adaptation under uncertainty: a probabilistic model checking approach. In: Proc. 10th Joint Meeting on Foundations of Software Engineering (FSE'15). pp. 1–12. ACM (2015)
62. Norman, G.: Analysing randomized distributed algorithms. In: Validation of Stochastic Systems: A Guide to Current Research. LNCS (Tutorial Volume), vol. 2925, pp. 384–418. Springer (2004)
63. Norman, G., Parker, D., Kwiatkowska, M., Shukla, S.: Evaluating the reliability of NAND multiplexing with PRISM. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **24**(10), 1629–1637 (2005)
64. Norman, G., Parker, D., Sproston, J.: Model checking for probabilistic timed automata. *Formal Methods in System Design* **43**(2), 164–190 (2013)
65. Nuraishah, S., Jawaddi, A., Ismail, A., Sulaiman, M.S., Cardellini, V.: Analyzing energy-efficient and Kubernetes-based autoscaling of microservices using probabilistic model checking. *Journal of Grid Computing* **23**(3) (2024)
66. Parker, D.: Implementation of Symbolic Model Checking for Probabilistic Systems. Ph.D. thesis, University of Birmingham (2002)
67. Päßler, J., ter Beek, M.H., Damiani, F., Johnsen, E.B., Tapia Tarifa, S.L.: Analysing self-adaptive systems as software product lines. *Journal of Systems and Software* **222**, 112324 (2025)
68. Raitza, M., Märcker, S., Trommer, J., Heinzig, A., Klüppelholz, S., Baier, C., Kumar, A.: Quantitative characterization of reconfigurable transistor logic gates. *IEEE Access* **8**, 112598–112614 (2020)
69. Rungta, N., Brat, G., Clancey, W.J., Linde, C., Raimondi, F., Seah, C., Shafto, M.: Aviation safety: modeling and analyzing complex interactions between humans and automated systems. In: Proc. 3rd International Conference on Application and Theory of Automation in Command and Control Systems (ATACCS'13). pp. 27–37. ACM (2013)
70. Rybinski, M., Szymanska, Z., Lasota, S., Gambin, A.: Modelling the efficacy of hyperthermia treatment. *Journal of The Royal Society Interface* **10**(88) (2013)
71. Sardar, M.U., Dubslaff, C., Klüppelholz, S., Baier, C., Kumar, A.: Performance evaluation of thermal-constrained scheduling strategies in multi-core systems. In: Proc. 16th European Workshop on Computer Performance Engineering (EPEW'19). LNCS, vol. 12039, pp. 133–147. Springer (2019)
72. Siddique, U., Hoque, K.A., Johnson, T.T.: Formal specification and dependability analysis of optical communication networks. In: Proc. Design, Automation & Test in Europe Conference & Exhibition (DATE'17) (2017)
73. Stamatelatos, G., Gyftopoulos, S., Drosatos, G., Efraimidis, P.S.: Revealing the political affinity of online entities through their twitter followers. *Information Processing & Management* **57**(2), 102172 (2020)
74. Steel, G.: Formal analysis of PIN block attacks. *Theoretical Computer Science* **367**(1-2), 257–270 (2006)
75. Tavalato, P., Luh, R., Eresheim, S.: Formalizing real-world threat scenarios. In: Proc. 8th International Conference on Information Systems Security and Privacy (ICISSP'22). pp. 281–289. SciTePress (2022)
76. Vardi, M.: Automatic verification of probabilistic concurrent finite state programs. In: Proc. 26th Symposium on Foundations of Computer Science (FOCS'85). pp. 327–338. IEEE CS Press (1985)
77. Volk, M., Sher, F., Katoen, J.P., Stoelinga, M.: SAFEST: Fault tree analysis via probabilistic model checking. In: Proc. 70th Reliability and Maintainability Symposium (RAMS'24). pp. 1–7 (2024)

78. Weik, N., Volk, M., Katoen, J., Nießen, N.: DFT modeling approach for operational risk assessment of railway infrastructure. *International Journal on Software Tools for Technology Transfer* **24**(3), 331–350 (2022)
79. Younes, H., Simmons, R.: Probabilistic verification of discrete event systems using acceptance sampling. In: *Proc. 14th International Conference on Computer Aided Verification (CAV'02)*. LNCS, vol. 2404. Springer (2002)
80. Zheng, X., Bolton, M.L., Daly, C., Bilekoff, E.: The development of a next-generation human reliability analysis: Systems analysis for formal pharmaceutical human reliability (SAFPH). *Reliability Engineering & System Safety* **202** (2020)
81. PRISM bibliography, [prismmodelchecker.org/bib.php](https://prismmodelchecker.org/bib.php)