Advances and Challenges of Quantitative Verification and Synthesis for Cyber-Physical Systems

Marta Kwiatkowska
Department of Computer Science, University of Oxford, UK

***Abstract*—We are witnessing a huge growth of cyber-physical systems, which are autonomous, mobile, endowed with sensing, controlled by software, and often wirelessly connected and Internet-enabled. They include factory automation systems, robotic assistants, self-driving cars, and wearable and implantable devices. Since they are increasingly often used in safety- or business-critical contexts, to mention invasive treatment or biometric authentication, there is an urgent need for modelling and verification technologies to support the design process, and hence improve the reliability and reduce production costs. This paper gives an overview of quantitative verification and synthesis techniques developed for cyber-physical systems, summarising recent achievements and future challenges in this important field.**

I. Introduction

In today's world, a multitude of sensor-enabled, software-controlled computing devices are autonomously making decisions on our behalf. These cyber-physical systems (CPS) are 'smart', wirelessly connected and Internet-enabled, able to monitor and control physical processes, and organised into communities, networks and ecosystems. Examples range from smartphones equipped with miniature sensors, e.g. accelerometers and GPS, self-parking and self-driving cars, environmental and wildlife monitoring, as well as implantable devices such as glucose monitors and cardiac pacemakers. Future potential developments in this area are endless, with nanotechnology and molecular sensing devices already envisaged.

As the number of CPS on the market and in deployment has escalated, unfortunately so has the number of device recalls [1] and remote hacking attacks [2]. These have naturally prompted a surge of interest in methodologies for ensuring their safety, integrity and reliability. Model-based design and automated verification technologies offer a number of advantages, particularly with regard to embedded software controllers: they enable rigorous software engineering methods such as model checking in addition to testing, and have the potential to reduce the development effort through code generation and software reuse via product lines. Models can be extracted from high-level design notations or even source code, represented as finite-state abstractions, and systematically analysed to establish if, e.g., the executions never violate a given property. These technologies provide means to automatically analyse properties such as "the probability that an airbag fails to activate on impact is tolerably low" (safety) and "the smartphone software always correctly authenticates the user" (reliability).

CPS devices possess quantitative characteristics: they not only involve decisions and discrete mode switching, but also features such as real-time delays and constrained resources. They often autonomously monitor and control physical processes, for example chemical concentration, and exhibit stochasticity, which arises from randomisation, for example used in distributed coordination; uncertainty, due to sensor noise, imprecision of localisation and partial observability; and stochastic dynamics, such as user mobility or physical motion. Quantitative, probabilistic modelling is an established methodology to provide safety and software quality assurance, and employs a variety of analysis methods drawn from automated verification, numerical solution techniques and statistics. Quantitative verification [3] techniques, in particular, aim to establish quantitative properties, for example, calculating the real-time deadline, probability of an event of interest, or expected cumulated cost of some process. Tools such as the real-time model checker UPPAAL [4] and the probabilistic model checker PRISM [5] are widely used for this purpose in several application domains.

However, CPS devices present new challenges, requiring extensions of model-based quantitative verification techniques to ensure correctness of their embedded software. Autonomy, in particular, forces us to look beyond verification, and to develop controller (also called strategy) synthesis techniques from quantitative specifications. It also results in the need to consider cooperative, competitive and adversarial behaviour due to conflicting goals, which are well suited to game-theoretic techniques. Sensor-enabledness, in conjunction with mobility, Internet connectivity and aspects such as monitoring and control of physical processes, for example electrical signal in the heart or concentration of glucose in the blood, incur the need for stochastic hybrid models, often with non-linear dynamics. Finally, CPS are increasingly often interacting with humans, as for example in semi-autonomous driving, leading to the need to consider the human-in-the-loop problem. In addition, increasing reliance of CPS on privacy/security assurance and their take up in safety-critical contexts significantly raise the prospect of unacceptable, or even life-endangering, risks.

Below we give a number of concrete example of desirable quantitative properties that one might like to verify for CPS. "the autonomous car has a strategy to safely cross the junction, irrespective of action of other road users" (competitive); "the autonomous agents have a collective strategy to execute a task, even if a single agent is unable to" (cooperative); "the probability that biometric security of the smartphone software is compromised is tolerably low, under any adversarial scenario" (adversarial); "the expected time to select a network service is within a specified interval, assuming wireless communication failure rate is within specified tolerance" (stochastic); "the probability that the wearable device software fails to identify a

Supported in part by the ERC Advanced Grant Grant VERIWARE and EPSRC Programme Grant EP/M019918/01.

dangerous level of glucose in the blood is tolerably low, assuming the dynamics is within specified bounds" (nonlinear hybrid dynamics); and "the probability that the semi-autonomous car causes an accident is tolerably low, assuming the driver is attentive" (human-in-the-loop).

II. Summary of Recent Advances

Conventional verification via model checking inputs a description of a model, representing a state-transition system, and a specification, typically a formula in some temporal logic, and return yes or no, indicating whether or not the model satisfies the specification. In quantitative verification [3], [6], [7], the models can be viewed as collections of states, where the transitions between them can be discrete, probabilistic, or continuous flows. A probability space induced on the system behaviours enables the calculation of likelihood of the occurrence of certain events of interest during the execution of the system, assuming resolution of nondeterministic decisions. This in turn allows one to make quantitative statements about the system, in addition to qualitative statements such as reachability or invariance. Probabilities are captured via probabilistic operators that extend conventional (timed or hybrid) temporal logics, which allow one to express probabilistic specifications such as the probability of a the smartphone software being infiltrated is below a specified bound. Models can be additionally annotated with quantities that represent, e.g., time passage or energy consumption, for which expectations are typically considered.

An important recent trend is a shift towards synthesis, where the ultimate goal is to construct a model from a specification, where techniques based on sketches and search-based method have been introduced. In the quantitative setting, simpler variants of this problem have been considered to date, including correct-by-construction controller synthesis from a given quantitative temporal logic specification, and parameter synthesis, see below, where parameter valuations are selected to optimise a given quantitative objective.

Quantitative verification employs a variety of graph-theoretic and symbolic methods drawn from conventional model checking, together with probabilistic analysis. The latter involves numerical computation, such as solving linear or differential equations or optimisation problems, where answers can be computed with specified precision, or statistical model checking, based on simulating execution runs and applying statistical techniques such as hypothesis testing to estimate the probability or expectation of some event holding up to a given confidence interval. Controller synthesis proceeds by first computing the optimal value and then extracting the strategy. An important consideration is ensuring scalability of the techniques, which is typically achieved through symbolic methods, quantitative abstraction refinement, or some suitable combination with stochastic search.

There are a number of quantitative verification tools that have been developed, to mention UPPAAL [4] and MRMC [8]; we focus here on the probabilistic model checker PRISM [5] and its recent extension PRISM-games [9], [10]. PRISM is based on symbolic BDD-based techniques [11], [12] that provide compact storage for probabilistic models and ensure efficiency of (approximate) computation of the probability. It supports five probabilistic models, discrete- and continuous time Markov chains (DTMCs, CTMCs), Markov decision processes (MDPs), probabilistic timed automata (PTAs) and stochastic games (SMGs), for both verification and strategy synthesis. Applications of probabilistic model checking using PRISM have spanned multiple fields, from wireless protocols and security analysis, through debugging DNA computing designs, to smart energy grids and strategy synthesis for autonomous urban driving.

Below we describe a number of recent advances centred on quantitative verification with PRISM that are relevant for CPS.

Parameter synthesis. A parametric model is one where some value, for example transition probability or timing delay, is given as a parameter, and the probability of an event of interest is then a function of the parameters. The parameter synthesis problem aims to find an optimal value of the parameter that guarantees the satisfaction of a quantitative temporal logic property. Parameter synthesis techniques have been developed for discrete-time Markov chains and probabilistic parameters [13]–[15] and for continuous-time Markov chains, where optimal rates can be synthesised for time-bounded specifications [16], [17]. The techniques are based on region refinement in conjunction with sampling, and have been recently improved through a GPU adaptation. They have been applied to a range of case studies, including synthesising optimal repair rates for the Google file system [18]. Approximate parameter synthesis is also available for MDP models.

Stochastic games. Stochastic game models generalise Markov decision processes in that they allow to model systems composed of a number of players, and are thus particularly well suited to modelling competitive scenarios, such as sharing network bandwith, where users try to selfishly maximise their own utility and pricing mechanisms have to be designed to disincentivise such behaviour. Stochastic games allow one to reason about strategic decisions of coalitions of agents competing or collaborating to achieve some quantitative objective, for example total expected reward or longrun average [19]–[21]. They are supported by the tool PRISM-games, and have been used in a variety of case studies, including energy smartgrid [22], [23] and aircraft power distribution [20], [21].

Multiobjective properties. When verifying correctness of systems, it is often necessary to consider not just a single objective, but instead simultaneous satisfaction of several objectives, for example minimise expected fuel consumption and maximise the probability to reach a goal. In general, Boolean combinations of quantitative objectives can be considered [19], [24], [25], with the respective tradeoffs represented via a Pareto curve [26], where each point on the curve corresponds to the objectives combined with different weights that can be selected by the designer. Multiobjective verification and strategy synthesis are supported in PRISM for MDPs and in PRISM-games for SMGs, and have been used, for example, to synthesise strategies for autonomous urban driving [27].

Compositionality. Complex designs usually comprise multiple components operating in parallel. Verification of such system is aided by compositional assume-guarantee reasoning, where verification of a property for a composed system can be reduced to checking certain properties on the components. Compositional verification is supported in PRISM for a range

of assume-guarantee rules [28] for linear-time temporal logic for a variant of MDPs. The method is based on reduction to multiobjective MDP verification. For SMG, compositional assume-guarantee strategy synthesis from multiobjective specifications [20], [21] is implemented in PRISM-games, which considerably improves scalability at a cost of expressiveness of the controllers.

Controller/strategy synthesis. Controller synthesis from high-level single and multiobjective specifications in quantitative temporal logic [7], [19] is supported for MDPs and SMGs respectively in PRISM and PRISM-games. This is relevant for motion planning in robotics [29], where temporal logic and reactive synthesis based on game-theoretic methods are being taken up for high-level planning. The strategies may need to be randomised and history-dependent. Recently, it was shown that performance of controller synthesis for MDPs and SMGs can be improved by incorporating machine learning [30], with which one can obtain guarantees on accuracy while exploring only a portion of the state space.

Approximate methods. Quantitative verification suffers from state-space explosion and may become infeasible for large systems. There are a number of approaches that reduce the state space, including partial order reduction [31], symmetry reduction [32], [33], bisimulation quotient [34] and state space aggregation [35]. For CTMC models of chemical reaction networks, the linear-noise approximation is particularly attractive in that it is independent of the population size and polynomial in the number of species, thus ameliorating state space explosion. These models are discrete state. Labelled Markov processes, which are continuous-space models, can be approximated as MDPs by employing approximate bisimulations [36].

Self-adaptation and learning. Today's software is increasingly expected to adapt to the changing demands of workload, or even requirements, while ensuring dependability. This is challenging to automate and, in [37], a framework has been proposed that employs learning and quantitative verification at runtime within a monitor, analyse, plan and execute (MAPE) loop. The methods has been applied to service-based systems using PRISM [38] and, more recently, also using single and multiobjective verification for stochastic multi-player games based on PRISM-games [39], [40]. The case studies included defending a system against an external attack and developing adaptive middleware to manage sensor networks.

Real-time and hybrid models. CPS models exhibit a range of quantitative features, including real-time and nonlinear continuous flows, in addition to stochasticity, and can be captured by the general model of stochastic hybrid systems. Though this model is not supported in full generality, its submodels have been analysed using PRISM and related methods [41], [42]. Probabilistic timed automata, which are timed automata with transitions generalised to discrete probability distributions, are directly supported in PRISM, as is parameter synthesis for timing delays to optimise the probability of a given property. Timed I/O automata with data, a subclass of hybrid linear automata, are supported by a suite a techniques that integrate SMT and evolutionary computation techniques. The techniques support optimal timing delay synthesis from quantitative specifications [43], [44], but their probabilistic aspects are limited. Hybrid automata can be analysed using simulation-based techniques [45]. These techniques can provide guarantees on accuracy and have been developed for the analysis of cardiac pacemakers, heart models and personalisation of wearable devices [46]–[48].

Human-in-the-loop problem. CPS increasingly often interact with human operators, to mention robotic assistants and self-parking cars. Therefore, quantitative verification and synthesis needs to consider models which are compositions of CPS and operator models, with the latter modelling cognitive processes and social interactions. [49] studied controller synthesis from multiobjective specifications for UAVs interacting with human operators, using both PRISM and PRISM-games, with the latter allowing to model adversarial behaviour. Early work towards predictive models of semi-autonomous driving that employed PRISM [50] is encouraging, but more effort is necessary.

III. Challenges

Clearly, there has been much progress towards quantitative verification and aspects of quantitative synthesis for CPS, with a wide variety of relevant case studies serving as proof of concept, including energy smartgrid, UAV control, autonomous driving and implantable cardiac pacemakers. However, in order to tackle CPS in full generality, a number of significant challenges have to be overcome. We briefly review a selection of these below.

Model expressiveness. Although certain classes of stochastic hybrid models are currently supported, the techniques often rely on discretisation and their continuous and stochastic dynamics are limited. Approximate techniques that can handle the challenging interaction between nondeterminism, nonlinear continuous and stochastic dynamics are needed.

Partial information. Quantitative verification techniques have so far mostly been concerned with complete information systems. This restriction is not applicable to many CPS scenarios, where agents in the system only have partial information. Partial observability raises a number of algorithmic challenges that need to be tackled.

Model learning and adaptation from data. Quantitative verification has so far mainly focused on modelling system dynamics, but the behaviour of many CPS is data-driven. Techniques that integrate model learning from data in order to inform adaptation in real-time and analysis of the dynamics are needed.

Model synthesis from specifications. Though correct-by-construction synthesis of strategies has been studied, model synthesis from quantitative specifications has received little attention to date. A possible approach is combining template-based and parameter synthesis methods, but more work is required in this direction.

Scalability and efficiency. Compositional approaches aid scalability, but existing methods are limited in that they cannot handle hybrid models and learning assumptions for compositional assume-guarantee reasoning is difficult to automate. Another direction are symbolic techniques, as employed in verification and synthesis algorithms, whose efficiency has recently been improved by employing judicious combination with statistical or stochastic search methods, but more progress

is required. Methodologies combining induction, deduction and machine learning have potential here.

Modelling social interactions. CPS are increasingly often employed to assist and interact with humans, and operator models have to be taken into account. Though some progress has been made, models of cognitive processes and social interactions, such as those based on trust, are needed.

Certification. Ultimately, model-based quantitative verification and model synthesis will be used as part of quality assurance processes to certify products. To this end, techniques and tools to evaluate safety and dependability are needed.

IV. Conclusion

As CPS are becoming an integral part of our society, their failure carries potentially unacceptable and life-endangering risks. Rigorous model-based verification technologies incorporated within the design process can improve CPS safety and reliability and reduce development costs. This paper has briefly summarised quantitative verification and synthesis techniques developed for CPS and future research challenges in this field.

References

[1] "U.S. Food and Drug Admin., List of Device Recalls," http://www.fda.gov/MedicalDevices/Safety/ListofRecalls/default.htm.

[2] "Researchers hack into driverless car system, take control of vehicle," http://www.nationaldefensemagazine.org/archive/2015/May/Pages/ResearchersHackIntoDriverlessCarSystemTakeControlofVehicle.aspx.

[3] M. Kwiatkowska, "Quantitative verification: Models, techniques and tools," in *Proc. ESEC/FSE'07*. ACM Press, September 2007, pp. 449–458.

[4] T. Amnell, G. Behrmann, J. Bengtsson, P. R. D'Argenio, A. David, A. Fehnker, T. Hune, B. Jeannet, K. G. Larsen, M. O. Möller, P. Pettersson, C. Weise, and W. Yi, "Uppaal - Now, Next, and Future," in *Modelling and Verification of Parallel Processes*, ser. Lecture Notes in Computer Science Tutorial, F. Cassez, C. Jard, B. Rozoy, and M. Ryan, Eds., no. 2067. Springer–Verlag, 2001, pp. 100–125.

[5] M. Kwiatkowska, G. Norman, and D. Parker, "PRISM 4.0: Verification of probabilistic real-time systems," in *Proc. CAV'11*, ser. LNCS, G. Gopalakrishnan and S. Qadeer, Eds., vol. 6806. Springer, 2011, pp. 585–591.

[6] ——, "Stochastic model checking," in *Formal Methods for the Design of Computer, Communication and Software Systems: Performance Evaluation (SFM'07)*, ser. LNCS (Tutorial Volume), M. Bernardo and J. Hillston, Eds., vol. 4486. Springer, 2007, pp. 220–270.

[7] V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker, "Automated verification techniques for probabilistic systems," in *Formal Methods for Eternal Networked Software Systems (SFM'11)*, ser. LNCS, M. Bernardo and V. Issarny, Eds., vol. 6659. Springer, 2011, pp. 53–113.

[8] J.-P. Katoen, I. Zapreev, E. M. Hahn, H. Hermanns, and D. Jansen, "The ins and outs of the probabilistic model checker MRMC," in *Proc. 6th Int. Conf. Quantitative Evaluation of Systems (QEST'09)*. IEEE Computer Society Press, 2009, pp. 167–176.

[9] T. Chen, V. Forejt, M. Kwiatkowska, D. Parker, and A. Simaitis, "PRISM-games: A Model Checker for Stochastic Multi-Player Games," in *Proc. of Tools and Algorithms for the Construction and Analysis of Systems TACAS*, ser. LNCS, vol. 7795, 2013, pp. 185–191.

[10] M. Kwiatkowska, D. Parker, and C. Wiltsche, "PRISM-games 2.0: A Tool for Multi-Objective Strategy Synthesis for Stochastic Games," in *Proc. of Tools and Algorithms for the Construction and Analysis of Systems TACAS*, 2016, to appear.

[11] M. Kwiatkowska, G. Norman, and D. Parker, "Probabilistic symbolic model checking with PRISM: A hybrid approach," in *Proc. 8th Int. Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS'02)*, ser. LNCS, J.-P. Katoen and P. Stevens, Eds., vol. 2280. Springer, 2002, pp. 52–66.

[12] ——, "Probabilistic symbolic model checking with PRISM: A hybrid approach," *International Journal on Software Tools for Technology Transfer (STTT)*, vol. 6, no. 2, pp. 128–142, 2004.

[13] E. M. Hahn, H. Hermanns, B. Wachter, and L. Zhang, "PARAM: A model checker for parametric Markov models," in *Proc. 22nd Int. Conf. Computer Aided Verification (CAV'10)*, ser. LNCS, vol. 6174. Springer, 2010, pp. 660–664.

[14] T. Chen, E. M. Hahn, T. Han, M. Kwiatkowska, H. Qu, and L. Zhang, "Model repair for Markov decision processes," in *Proc. 7th Int. Symp. Theoretical Aspects of Software Engineering (TASE'13)*. IEEE Computer Society Press, 2013, pp. 85–92.

[15] C. Dehnert, S. Junges, N. Jansen, F. Corzilius, M. Volk, H. Bruintjes, J.-P. Katoen, and E. Ábrahám, "PROPhESY: A PRObabilistic ParamEter SYnthesis tool," in *Proc. 27th Int. Conf. Computer Aided Verification (CAV'15)*, ser. LNCS, vol. 9206. Springer, 2015, pp. 214–231.

[16] M. Ceska, F. Dannenberg, M. Kwiatkowska, and N. Paoletti, "Precise parameter synthesis for stochastic biochemical systems," in *Proc. 12th Conference on Computational Methods in Systems Biology (CMSB'14)*, ser. LNCS, vol. 8859. Springer, 2014, pp. 86–89.

[17] M. Ceska, F. Dannenberg, N. Paoletti, M. Kwiatkowska, and L. Brim, "Precise parameter synthesis for stochastic biochemical systems," *Acta Informatica, to appear*, 2016.

[18] M. Ceska, P. Pilar, N. Paoletti, L. Brim, and M. Kwiatkowska, "PRISM-PSY: Precise GPU-accelerated parameter synthesis for stochastic systems," in *22nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), to appear*, ser. LNCS. Springer, 2016.

[19] T. Chen, V. Forejt, M. Z. Kwiatkowska, A. Simaitis, and C. Wiltsche, "On Stochastic Games with Multiple Objectives," in *Proc. of Mathematical Foundations of Computer Science MFCS*, 2013, pp. 266–277.

[20] N. Basset, M. Z. Kwiatkowska, and C. Wiltsche, "Compositional Controller Synthesis for Stochastic Games," in *Proc. of Concurrency Theory CONCUR*, 2014, pp. 173–187.

[21] N. Basset, M. Z. Kwiatkowska, U. Topcu, and C. Wiltsche, "Strategy Synthesis for Stochastic Games with Multiple Long-Run Objectives," in *Proc. of Tools and Algorithms for the Construction and Analysis of Systems TACAS*, 2015, pp. 256–271.

[22] T. Chen, V. Forejt, M. Kwiatkowska, D. Parker, and A. Simaitis, "Automatic Verification of Competitive Stochastic Systems," in *Proc. of Tools and Algorithms for the Construction and Analysis of Systems TACAS*, ser. LNCS, vol. 7214, 2012, pp. 315–330.

[23] ——, "Automatic verification of competitive stochastic systems," in *Proc. 18th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'12)*, ser. LNCS, vol. 7214. Springer, 2012, pp. 315–330.

[24] K. Etessami, M. Kwiatkowska, M. Vardi, and M. Yannakakis, "Multi-objective model checking of Markov decision processes," *Logical Methods in Computer Science*, vol. 4, no. 4, pp. 1–21, 2008.

[25] V. Forejt, M. Kwiatkowska, G. Norman, D. Parker, and H. Qu, "Quantitative multi-objective verification for probabilistic systems," in *Proc. TACAS'11*, ser. LNCS, P. Abdulla and K. Leino, Eds., vol. 6605. Springer, 2011, pp. 112–127.

[26] V. Forejt, M. Kwiatkowska, and D. Parker, "Pareto curves for probabilistic model checking," in *Proc. ATVA'12*, ser. LNCS, S. Chakraborty and M. Mukund, Eds., vol. 7561. Springer, 2012, to appear.

[27] T. Chen, M. Z. Kwiatkowska, A. Simaitis, and C. Wiltsche, "Synthesis for Multi-objective Stochastic Games: An Application to Autonomous Urban Driving," in *Proc. of Quantitative Evaluation of Systems QEST*, 2013, pp. 322–337.

[28] M. Kwiatkowska, G. Norman, D. Parker, and H. Qu, "Compositional probabilistic verification through multi-objective model checking," *Information and Computation*, vol. 232, pp. 38 – 65, 2013.

[29] J. Wang, X. C. Ding, M. Lahijanian, I. C. Paschalidis, and C. Belta, "Temporal logic motion control using actor-critic methods," *Int. J. of Rob. Res.*, vol. 34, no. 10, pp. 1329–1344, 2015.

[30] T. Brázdil, K. Chatterjee, M. Chmelík, V. Forejt, J. Křetínský, M. Kwiatkowska, D. Parker, and M. Ujma, "Verification of Markov decision processes using learning algorithms," in *Proc. 12th Int. Symp. Automated Technology for Verification and Analysis (ATVA'14)*, ser. LNCS, vol. 8837. Springer, 2014, pp. 98–114.

[31] C. Baier, M. Groesser, and F. Ciesinski, "Partial order reduction for probabilistic systems," in *Proc. QEST'04*. IEEE, 2004.

[32] A. Donaldson and A. Miller, "Symmetry reduction for probabilistic model checking using generic representatives," in *Proc. 4th Int. Symp. Automated Technology for Verification and Analysis (ATVA'06)*, ser. Lecture Notes in Computer Science, S. Graf and W. Zhang, Eds., vol. 4218. Springer, 2006, pp. 9–23.

[33] M. Kwiatkowska, G. Norman, and D. Parker, "Symmetry reduction for probabilistic model checking," in *Proc. 18th Int. Conf. Computer Aided Verification (CAV'06)*, ser. LNCS, T. Ball and R. Jones, Eds., vol. 4114. Springer, 2006, pp. 234–248.

[34] C. Dehnert, J.-P. Katoen, and D. Parker, "SMT-based bisimulation minimisation of Markov models," in *Proc. 14th Int. Conf. Verification, Model Checking, and Abstract Interpretation (VMCAI'13)*, ser. LNCS, R. Giacobazzi, J. Berdine, and I. Mastroeni, Eds., vol. 7737. Springer, 2013, pp. 28–47.

[35] A. Abate, L. Brim, M. Ceska, and M. Kwiatkowska, "Adaptive aggregation of markov chains: Quantitative analysis of chemical reaction networks," in *27th International Conference on Computer Aided Verification (CAV)*, ser. LNCS. Springer, 2015.

[36] A. Abate, M. Kwiatkowska, G. Norman, and D. Parker, "Probabilistic model checking of labelled markov processes via finite approximate bisimulations," in *Horizons of the Mind – P. Panangaden Festschrift*. Springer Verlag, 2014, pp. 40–58. [Online]. Available: ./publications/bcAKNP14.pdf

[37] R. Calinescu, C. Ghezzi, M. Kwiatkowska, and R. Mirandola, "Self-adaptive software needs quantitative verification at runtime," *Communications of the ACM*, vol. 55, no. 9, pp. 69–77, Sep. 2012.

[38] R. Calinescu, L. Grunske, M. Kwiatkowska, R. Mirandola, and G. Tamburrelli, "Dynamic QoS management and optimisation in service-based systems," *IEEE Transactions on Software Engineering*, vol. 37, no. 3, pp. 387–409, 2011.

[39] T. Glazier, J. Camara, B. Schmerl, and D. Garlan, "Analyzing Resilience Properties of Different Topologies of Collective Adaptive Systems," in *Proc. of Self-Adaptive and Self-Organizing Systems Workshops SASOW*, 2015, pp. 55–60.

[40] J. Cámara, D. Garlan, B. Schmerl, and A. Pandey, "Optimal Planning for Architecture-based Self-adaptation via Model Checking of Stochastic Games," in *Proc. of Symposium on Applied Computing SAC*, 2015, pp. 428–435.

[41] Z. Jiang, M. Pajic, and R. Mangharam, "Cyber–physical modeling of implantable cardiac medical devices," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 122–137, 2012.

[42] M. Pajic, R. Mangharam, O. Sokolsky, D. Arney, J. Goldman, and I. Lee, "Model-driven safety analysis of closed-loop medical systems," *Industrial Informatics, IEEE Transactions on*, pp. 3–16, 2014.

[43] M. Diciolla, C. H. P. Kim, M. Kwiatkowska, and A. Mereacre, "Synthesising optimal timing delays for timed i/o automata," in *14th International Conference on Embedded Software (EMSOFT'14)*, 2014.

[44] M. Kwiatkowska, A. Mereacre, N. Paoletti, and A. Patanè, "Synthesising robust and optimal parameters for cardiac pacemakers using symbolic and evolutionary computation techniques," in *Proceedings of the 4th International Workshop on Hybrid Systems and Biology (HSB 2015)*, ser. LNCS/LNBI, vol. 9271. Springer, 2015, pp. 119–140.

[45] Z. Huang, C. Fan, A. Mereacre, S. Mitra, and M. Kwiatkowska, "Invariant verification of nonlinear hybrid automata networks of cardiac cells," in *Proc. Computer Aided Verification*. Springer, 2014.

[46] T. Chen, M. Diciolla, M. Kwiatkowska, and A. Mereacre, "Quantitative verification of implantable cardiac pacemakers over hybrid heart models," *Information and Computation*, vol. 236, pp. 87–101, 2014.

[47] M. Kwiatkowska, A. Mereacre, and N. Paoletti, "On quantitative software quality assurance methodologies for cardiac pacemakers," in *Proc. 6th International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (ISoLA)*, ser. LNCS, vol. 8803. Springer, 2014, pp. 365–384.

[48] B. Barbot, M. Kwiatkowska, A. Mereacre, and N. Paoletti, "Estimation and verification of hybrid heart models for personalised medical and wearable devices," in *13th International Conference on Computational Methods in Systems Biology (CMSB 2015)*, ser. LNCS, vol. 9308. Springer, 2015, pp. 3–7.

[49] L. Feng, C. Wiltsche, L. Humphrey, and U. Topcu, "Controller Synthesis for Autonomous Systems Interacting with Human Operators," in *Proc. of Int. Conf. on Cyber-Physical Systems ICCPS*, 2015, pp. 70–79.

[50] D. Sadigh, K. Driggs-Campbell, A. Puggelli, W. Li, V. Shia, R. Bajcsy, A. L. Sangiovanni-Vincentelli, S. S. Sastry, and S. A. Seshia, "Data-driven probabilistic modeling and verification of human driver behavior," in *Formal Verification and Modeling in Human-Machine Systems, AAAI Spring Symposium*, March 2014.