Verification of Networks of Smart Energy Systems over the Cloud

Alessandro Abate(✉)

Department of Computer Science, University of Oxford, Oxford, UK
aabate@cs.ox.ac.uk

Abstract. This contribution advocates the use of formal methods to verify and certifiably control the behaviour of computational devices interacting over a shared infrastructure. Formal techniques can provide compelling solutions not only when safety-critical goals are the target, but also to tackle verification and synthesis problems on populations of such devices: we argue that alternative solutions based on classical analytical techniques or on approximate computations are prone to errors with global repercussions, and instead propose an approach based on formal abstractions, error-based refinements, and the use of interface functions for the synthesis of abstract controllers and their concrete implementation. Two applicative areas are elucidated, dealing respectively with thermal loads and electricity-generating devices interacting over a smart energy network or over a local power grid. We discuss the aggregation of large populations of thermostatically-controlled loads and of photovoltaic panels, and the corresponding problems of energy management in smart buildings, of demand-response on smart grids, and respectively of frequency stabilisation and grid robustness.

Keywords: Cyber-physical systems · Systems of systems · Internet of things · Hybrid models · Stochastic processes · Nondeterminism · Partial observations · Real-time systems · Security · Model learning · (quantitative) probabilistic verification · Formal abstractions · Bisimulations · Statistical verification · Feedback controllers · Policy and strategy synthesis · Distributed control · Safety and performance · Games · Correct-by design synthesis · Autonomy · Energy and power networks · Electricity demand-response · Thermostatically controlled loads · Smart buildings and smart grids · Photovoltaic panels · Blackouts · Aggregations of large populations

1 Technological Context: Networks of Complex Systems

There is an ever increasing trend to place and integrate computational devices over the cloud. By "cloud" we denote an infrastructure (predominantly with digital features over physical qualities) allowing for seamless (mostly wireless) communication between devices, which thus give form to a network of distinct

S. Bogomolov et al. (Eds.): NSV 2016, LNCS 10152, pp. 1–14, 2017.
DOI: 10.1007/978-3-319-54292-8_1

components[1]. Such devices are nowadays identified as "smart" – this attribute denotes a capability to engage with the environment (over the cloud) that is not purely static: indeed modern devices are clearly changing from having reactive features to active ones; not only do they interact passively with the environment (neighbouring devices, or a population-level feedback from the whole network), but also internally learn from it and indeed actively engage with it by modifying it locally. This interaction, which is either digital or physical depending on the type of interconnection or embedding within the medium, leads to repercussions over adjacent devices. As such, unlike static or purely reactive elements, these devices comprise internal dynamics that are locally actively coupled with neighbouring components.

Such devices are furthermore often "complex", in that they encompass both digital components and (possibly) analogue ones, and are likely to evolve (internally, or spatially within their environment) according to non-trivial dynamics. Digital components may comprise the computational platforms they run on, or the logic-based control architectures that affect their dynamical behaviour, whereas analogue parts may encompass the physical medium they are embedded into, or the continuous components making up the devices themselves.

These devices is often referred to as "cyber-physical systems" (CPS), and the network they are part of is thought of as a "system of systems" (SOS), or as "internet of things" (IoT). The first acronyms seem to be more relevant within engineering contexts, whereas the second depicts a more abstract concept (it does not necessarily distinguish between digital and analogue components, nor automatically emphasises the networking or dynamical aspects) and appears in use within the applied mathematics literature. The third and last seems to be widespread within the computer sciences and in general involves fewer elements of coordination, actuation, and dynamics than CPS.

We are thus facing an engineering platform of multiple, interleaving and interacting complex systems, a true "system of systems" with issues of synchronisation and coordination, feedback from couplings and interactions, and with a global behaviour that is emergent from local dynamics. Such complex systems are thus not monolithic, and entail issues of operational independence, geographical distribution and heterogeneity, and local adaptability. We argue throughout this work that this network ought to be quantitatively analysed, by means of formal methods that are based on mathematical models of the single system components and of their networked interactions. Beyond analysis, autonomy can be established by means of modern feedback control architectures, which ought to be certified and indeed be "correct-by-design".

It is often the case that such complex systems are only partially known, that is they are not exactly nor fully observed, and possibly subject to uncertainty and/or randomness. We intend the latter aspect to be due to the presence of heterogeneity, noise, random or chaotic (and as such not precisely predictable) behaviours (as with, say, weather forecast), or to the presence of human users in

[1] In this context we do not distinguish between computations performed over the cloud (fog computing) or the edge of this platform.

the loop (interfering with - or perhaps supporting - the autonomy aspects of the system). All such elements are in general hard to model deterministically, but can be encompassed by probabilistic terms (whether objectively or subjectively construed). Notice that we distinguish this presence (probability) from that of environmental non-determinism (to be separately discussed in each of the next parts). These assumptions lead to the consideration of stochastic behaviours in the systems under study, either at the level of their dynamics (this is known as process noise), or upon their observations (e.g., sensors uncertainty or noise).

We take this opportunity to clarify and distinguish among two different aspects, both relevant but distinct to the systems under consideration. On the one hand, this setup deals with issues of learning within the single devices comprising the overall CPS. Learning can be directed towards models of the devices, as well as towards specifications (harvesting requirements). On the other hand, the platform comprises issues of computation, communication, and control. Whilst learning issues are data-based, and require bottom-up studies, the latter set of problems naturally require top-down analysis and are classically model-based. The first issue (learning) is internal to the device interacting (via sensing and actuation) with and adapting to the local environment. The second issue (communication) deals with the specifics of the processing and exchange of information within and among devices.

Next, we move from the perspective of the practical engineering setups, to that of their mathematical models. This new perspective allows for the development of quantitative analysis tools, and to obtain results leveraging the area of formal methods.

2 Formal Verification of Complex Models

Quantitative issues such as reliability, safety, dependability, are key in practical applications of the complex engineering setups that we have previously introduced. These issues play a central role either because the components belong to safety-critical applications, or because it is often the case that the complex network of systems is comprised within a safety-critical setup: much like existing complex engineering infrastructures, in an era of increasing inter-connectivity systems are seldom isolated, thus safety criticality depends on coupled behaviours and thus represents a global network issue.

The need for formal methods is therefore key. The use of these techniques is an alternative to approaches based on qualitative analysis, relying on more classical mathematics and often focusing on global, network-level properties, which are seldom useful in a complex context preventing analytical or explicit mathematical results. The use of formal methods is further in juxtaposition to fine-grained agent-based modelling and related simulation-based techniques, or to statistical approaches, which are known to be stymied with a number of limitations: they establish presence of potential faults or errors, but cannot assert their absence in general; they hardly scale when non-determinism and stochasticity play a relevant role in the dynamics, and when the system under study presents

continuous (uncountable) variables interleaved with discrete components (discontinuous). Sample-based techniques further lack formal guarantees, which is in particular a fundamental limitation when single devices or whole systems are to be certified (as opposed to be validated towards quality assurance). On the other hand these approaches should not be completely dismissed, but rather integrated within the use of formal approaches which, as is known, present computational limitations when applied to large-scale and complex models, as is the case of the engineering setups under study.

Formal verification techniques hinge on quantitative models. The verification of quantitative models of complex systems requires handling a mathematical formalism encompassing dynamical variables evolving over hybrid (continuous/discrete) state spaces, comprising probabilistic behaviours, possibly with continuous-time semantics, and under partial observations of the model's variables. Emerging from a broad research initiative on alternative models of computation, a natural modelling framework encompassing all these aspects is that of stochastic hybrid models [5]. Such models are dynamically rich and require modern, tailored techniques for analysis, verification and synthesis.

At the outset, it is easy to realise that the verification of such models is bound to undecidability results, unless the problem at hand admits an analytical (that is, explicit) solutions, which is quite unlikely in view of models complexity and of the possibly rich objectives under consideration [18]. On the other hand, state-of-the-art software tools for automated quantitative model checking of complex models (e.g., the PRISM model checker) are not applicable to models comprising all the aspects discussed above. As a partial attainment of the grand goal of verification of truly complex models, literature has seen probabilistic model checking of finite-state models, model checking of concurrent models with continuous (dense) time semantics, reachability-based verification of (non-probabilistic) hybrid dynamical models, and some timid early attempts to provide partial evaluation of more complex models – all daunted by the sheer complexity of the goal. Notice that the mentioned successful verification instances apply to strict subsets of the target models of interest that have been discussed previously.

This body of work has been extended to deal with the verification of non-deterministic models – such extensions are often bound to conservative outcomes (whenever non-determinism in quantified universally), or to rather different synthesis frameworks (whenever non-determinism in quantified universally). The application of SAT or SMT techniques can be of particular interest for this goal, as well as results on robustness analysis, as classically investigated in control theory or, more recently, in formal methods.

If formal verification tools have not been fully extended and thoroughly applied to complex models, might complex models be compelled to fit existing formal methods tools? The next section elaborates on this goal.

The wider accessibility properties and interconnectivity features that modern smart devices allow carry as a drawback increased pressure towards issues of security. Security deals with coping with interferences of various sorts and nature, often thrusted in surreptitious and hidden manner, which can affect the correct

functioning of the single devices and can potentially lead to global repercussions. Amongst major categories of cyber security threats we (non comprehensively) list those affecting resilience and privacy, as well as malicious intrusions and attacks. Notice that the presence of security attacks requires models semantics that are fundamentally different than those (non-deterministic or probabilistic) used to describe the presence of the environment. In this security context, various frameworks have been put forward: from worst-case non-deterministic approaches, to average-case probabilistic models, to (possibly stochastic) game-theoretical frameworks. Of them, the latter approach appears to encompass the nature of potential attacks and of possible replies, mitigations, or preventions actions against them.

Key aspects at the interface of communication and computation are those dealing with real-time engineering. These, within a networking perspective, relate to important computational issues of inter-operability, synchronisation, and concurrency; and, from the communications perspective, to problems of network theory and issues of data integration within models.

Whilst understood as key in the context of CPS applications, we do not delve into details of such problems any further in this essay.

3 Approximate Model Checking of Stochastic and Hybrid Models

This contribution is underpinned by recent research on stochastic hybrid models [5]. Properties of interest are usually encoded within known and exploited modal logics, such as PCTL or CSL (whether in continuous or discrete time), or just by looking at the likelihood attached to trajectories verifying linear time specifications expressed in LTL or as (e.g., Büchi) automata over infinite strings. Extensions to conditional probabilities have been pursued in recent literature.

Over related frameworks of stochastic and hybrid models, a number of authors have recently investigated the characterisation of basic probabilistic reachability and invariance specifications [5], as well as the extension to reach-avoid (constrained reachability), and to richer properties such as linear-time properties expressed as a DFA or as Büchi automata [19]. With regards to the latter, infinite-horizon properties have been also studied, and involve advanced analytical tools [18].

Beyond characterisation and towards numerical assessment, the properties above have been computed by means of finite abstractions [3,4]. The derivation of formal errors due to such finite abstractions [9] has effectively led to the development of approximate model checking of stochastic hybrid models. Errors have been further extended and refined [12,19], and embedded in the development of a software tool, which feeds complex models to probabilistic model checking software packages.

FAUST^2 [13] is a Matlab-based software tool, which accepts as an input a stochastic process and a formal specification, and generates a finite abstraction that can be fed into a probabilistic model checker such as PRISM or MRMC.

The abstraction is guaranteed to abide by a user-defined error that is required on the satisfaction of the given property of interest. The error is computed based on the underlying dynamics of the SHS and on the given property [9], and encompasses the difference between the probability distributions (in time) of concrete and abstract models. It has been shown that such error induces an approximate probabilistic bisimulation relation between the concrete and the (finite) abstract models [1,15], which can be also of use to study transient dynamical properties [2,11,17]. The overall procedure leads to an anytime algorithm, which sequentially refines coarse model abstractions, based on an update of the computed bound on the current error [9].

Related results have been developed by approximating the concrete stochastic model with a noiseless abstraction [20,22], which is then deterministically verified by means of software tools for dynamical models, such as PESSOA. The outcome allows for the refinement of assertions or of synthesised controllers, in view of a quantified error of the abstraction procedure. Unlike the approach described earlier, the new error encompasses (a moment of) the absolute value of the difference between the solutions of the two models, and is shown to exist under certain contractivity (which is a form of stability) assumptions on the concrete model dynamics.

Cognate research on verification of stochastic (and hybrid) dynamical models has been looking at the use of stochastic SAT modulo theory, the recasting of a verification objective as the solution of a PDE (with associated numerics), or the approximation of the above quantitative verification problems as convex optimisation ones.

The lack of full access to the state variables leads to the setup of partially observed modes, of which a known instance is that of hidden Markov models. In view of undecidability issues, substantial work on heuristics for the analysis of these models within the field of artificial intelligence. Formally, this setup requires the introduction of sufficient statistics and work over a belief space which, with the exception of linear models with additive Gaussian noise and associated corresponding Kalman estimators, is in general prone to lack analytical and computational tractability. Further work, both theoretical and algorithmic, is most definitely needed on this class of models.

The verification of parametric models encompassing non-determinism has not been focus of thorough and practically scalable investigation, regardless of whether seen as internal (to be universally quantified against) or external (to be synthesised over). Perhaps further steps can be attained by principled use of SMT approaches, or of results in robustness analysis.

A known and evident concern on the applicability of model-based quantitative verification techniques is the issue of scalability: these approaches are known to be stymied by state-space explosion, which is particularly relevant for complex CPS models. It is necessary to mitigate this issue by means of a multi-pronged approach: exploiting model modularity and topological distributivity, use of assume-guarantee reasoning, employment of deep compositional results, interfacing with legacy systems (and corresponding models), use of paradigms

of object-oriented programming, and development of state-less abstractions [21]. On the other hand, it is well understood that sample-based techniques (such as in paradigms of runtime verification and testing) have the potential to scale to complex models, whilst on the down side they suffer from lack of tight performance guarantees. Within the context of formal verification, it is then of utmost interest to provide a novel formal integration of sample-based techniques within model-based deductive approaches [14]. As a side comment, learning algorithms can be directed towards models of the devices, as well as towards specifications (harvesting requirements). Towards this direction, the employment of coverage metrics has recently grown much interest.

As much as sample-based techniques or agent-based simulations seldom provide formal performance guarantees, approaches based on (non-formal) approximate computations are stymied by lack of certified behaviour. For instance, employing continuous mean-field limits (which are shown correct exclusively at asymptotic limits) renders an intrinsically probabilistic population a deterministic problem, which prevents the generation of certain allowable fringe behaviours. Further, the use of grid-based techniques with no control of the precision is bound to lead to suboptimal solutions, or errors that accumulate fast with time and which certainly do not meet any certification requirement or formal guarantee.

In conclusion, we argue that model verification based on classical analytical techniques has shown its limits, whereas sample-based results or outcomes based on non-formal approximate computations are prone to generate uncontrollable errors with overreaching global repercussions. We argue that a principled application of formal methods techniques, properly enhanced via computationally-prone approaches, is the way forward particularly for CPS applications.

4 From Verification to Synthesis: Correct-by-Design Control of Complex Models

Beyond quantitative verification, control synthesis also requires proper formalisation over complex models. As discussed, control is a form of external non-determinism, and as such it has to be contrasted with forms of non-determinism that are resolved by the environment, that are due to coarse-grained abstractions (internal non-determinism), or with (static) parametric uncertainty. The semantics of external non-determinism involve a volitive agent, which selects functions of time and of the state, possibly in a randomised manner and accounting for past history. Such selection leads to control laws known as policies or strategies. It is often of interest to focus on memoryless laws, possibly deterministic ones, which limit the computational overhead and do not infringe the Markov property of the closed-loop model. Classical control synthesis deals with performance criteria (introduced either as costs or rewards that are function of the state and/or the action space), over which optimality is sought via (respectively) minimisation or maximisation on a finite- or infinite-time horizon, within a predefined class of allowable policies. Of course such an approach can be applied to

the goal of maximising the likelihood of verifying a quantitative property, such as those in PCTL logics discussed above for stochastic processes. The study of finite-horizon quantitative properties expressed in PCTL boils down to a characterisation via multiplicative cost functions [5], whereas that of infinite-horizon properties can be reduced to reachability computation over a product automaton [19]. The study of infinite horizon reachability is tricky since it involves extended notions of absorbing sets that are related to classical ones of bottom strongly connected components, or of max-end components [18]. Even more so in the case of controller synthesis over infinite-horizon properties [19].

Of course it is of interest to expand the issue of synthesis over both quantitative specifications and over performance: this can be done by resorting to techniques for multi-criteria or lexicographic optimisation. This goal has been recently pursued both within the computer science area, the control theory, and the optimisation literature.

Beyond process noise affecting the state dynamics, lack of exact observations (as partial access to the hidden variables or presence of sensor noise) lead to partially observed models, such as POMDP. As in the previous section, control synthesis for partially observable models brings along a number of technical and computational hurdles which, whilst thoroughly investigated within artificial intelligence and control literature, have only in part led to results that can be deemed satisfactory within the stricter context of formal methods.

Control synthesis problems are further prone to be extended to stochastic (two-and-a-half player) games [8]. Games can be played against the environment (e.g., towards compositional reasoning) or against an adversary (e.g., for applications in security). The author recognises interesting connections between game-theoretical setups in applied mathematics and control theory (for instance, dealing with existential results over uncountable models) and problems that are algorithmically solved for discrete configurations in theoretical computer science. The concept of formal abstractions (discussed in the previous section) could provide a link between results in the two areas.

The digital platform that characterises networks of complex systems, which we called "cloud" earlier, encompasses pervasive elements of wireless communications, and moves away from older, tethered communications, thus allowing for more agile reconfigurations as well as practical mobility of the agents in the network. As much as general communication aspects (real-time issues, protocol design aspects, and the like) need to be dealt with at the level of modelling, the specific deployment of wireless channels require proper handling of packets corruption or losses, and delays. Accepting that controllers are not necessarily embedded in a monolithic plant, but rather separated from it by a communication network, a recent and lively literature [23] has started to investigate issues of communication within control theoretical architectures.

As much as verification algorithms applied to complex CPS models suffer from state-space explosion, control synthesis ones are stymied by Bellman's curse of dimensionality. Techniques aimed at speeding up such formal algorithms (either via abstractions, or via approximations, or through compositionality, or

also by mans of data-driven approaches, such as reinforcement learning) are naturally seen of much priority in this area.

5 Verification of Networks of Smart Energy Systems over the Cloud

In this section we provide a compelling application of the concepts, models and problems elaborated above. We discuss how present-day energy networks and electricity grids are transitioning to become interconnected networks of complex and smart systems, dynamically coupled both physically and over the cloud. We argue that their formal verification and correct-by-design control is relevant in engineering and industrial contexts, as well as for the market opportunities that they have the potential to catalyse. We provide two case studies zooming in on, respectively, the smart grid (investigated from the perspective of smart buildings in [6]) and electricity networks (elaborated in the project [7]).

Smart Buildings over the Smart Grid. Buildings consume more than 40 % of the energy in Europe. In order to sustainably reduce energy consumption by improving their usage and management, an optimal operation and an improved commissioning and maintenance of building management systems (BMS) are seen as key factors by the sector's industry. Efficient automation systems embedded in so-called "smart buildings" can indeed reduce the energy consumption up to an estimated thirty percent in many relevant instances. The objective of a smart building is to deliver useful building services that make occupants comfortable and productive (for instance, providing regulation for thermal comfort and air quality), at the lowest energy costs, over the entire building life cycle. This objective requires adding intelligence to the infrastructure of buildings and utilising information technology during their operation. This enables the connectivity of devices and components in a building (think of home automation devices, smart appliances, an application related to the broad area known as the "Internet of Things"), the interaction of buildings with their occupants and building operators or with their building management systems, as well as (at a higher level) their connection to other buildings or infrastructure components within a smart grid platform.

Such modern features and capabilities have opened new challenges related to the optimised performance of smart buildings, as (components of) networks of complex dynamical systems. Indeed, both the interconnection of smart BMS devices (such as sensorised HVAC modules) within a smart building, and (at a higher level) the local interaction of various buildings within a smart grid, clearly lead to the CPS configurations discussed above.

In the context of a single building, the construction of models that accurately capture the time evolution of its physical variables can be based on data gathered from the buildings. The continuous nature of physical variables, the discrete feature of digital controllers, and the presence of uncertainty originating from the environment and from the users behaviour, render the general framework of

Stochastic Hybrid Systems well-suited for modelling purposes [5]. Formal models enable the solution of engineering problems, such as optimal temperature regulation in a building, which in view of the slow dynamics we argue can be promising to tackle in the loop via formal methods. The stochastic and time-varying nature of the system under study suggests that a data-based update of the model, and more generally an integration of on-line data within the models under study [14], are of interest.

At a higher level, each building can be thought as a node in a network, such as a smart grid, partaking alongside other energy-consuming buildings in its dynamics, as well as interacting with devices generating energy. The connection with the CPS framework is again evident. Due to their flexibility in providing services to occupants, smart buildings can then be engaged in services by energy companies, such as load shifting, peak shaving, and more general in demand-response programs. As a result, whilst we can naturally think of engineering problems such as robustness and resiliency of the energy dynamics within a local smart grid, there is another layer of problems dealing with market design for demand response, with the engagement of consumers over grid markets, such as the electricity one (related to load shifting and consumers' demand response) and that dealing with economy of energy production (as per the concept of prosumers) and storage (zero net-energy buildings).

Within these CPS configurations, it is important to understand how global dynamics emerge from single dynamics and local interactions. One way to look at global dynamics arising from local contributions is to develop aggregation techniques: model aggregations lump together the dynamics of single buildings, providing a global description that is computable and scales well with the size of the problem.

Recent research has developed a procedure based on formal abstractions [10], which generates a finite stochastic dynamical model as an aggregation of the continuous temperature dynamics of a (possibly heterogeneous) population of Thermostatically Controlled Loads (TCL), which are basic models for the dynamics in smart buildings. The temperature of each single TCL is described by a stochastic difference equation and the TCL status by a deterministic switching mechanism – in all, a hybrid model. The procedure is formal as it allows the exact quantification of the error introduced by the abstraction. Research has discussed extensions to the case of controlled TCL, with dynamics affected by an aggregator (which could be a utility company engaged with consumers on demand response schemes). The structure of a control scheme (centralised, decentralised or distributed) hinges on many assumptions on the placement of sensors and on the capability of the actuators (HVAC modules), and in general on the information flow between the central aggregator and the single components of the population. The employment of distributed architectures appears to be relevant to optimise both engineering and economic goals.

We argue that approaches based on formal aggregation can be relevant both at an engineering level (working out precise approaches for safe and optimised energy consumption in a building, or for reliable and robust operation of a smart

grid) as well as at a financial level (understanding fairness in market design between energy providers and customers over demand response schemes [16]). In conjunction with the development of formal models and quantitative verification approaches, abstractions can offer a principled approach to the understanding of the complex dynamics of smart buildings, and to their optimised and certifiable operation within the context of smart grids.

Renewables Generation over the Power Grid. The previous section has suggested how the engagement of consumers over a smart grid had the potential to lead to some paradigm shifts: from a centralised and fuel-based grid operation to a decentralised and renewable-based economy; from passive electricity and gas consumers supplied by energy utility companies, to active electricity prosumers and as such to utility companies partners (if not competitors). The grid becomes an accessible infrastructure to be locally leveraged in both directions by numerous players on the energy market.

Smart buildings have the capability to generate energy by means of renewables, as is the case of photovoltaic panels or wind turbines. In many regions worldwide, renewables are increasingly relied upon for electricity generation. As much as beneficial green energy can be regarded towards a sustainable decrease of greenhouse gases, the de-carbonisation of energy usage, and the long-term goal of zero net-energy buildings, renewables pose challenging engineering problems in view of their distributed engagement within a power grid that was conceived and built for centralised production and distribution, and of the intrinsically volatile quality of the electricity generated, which hinges on meteorological and local grid conditions. The decentralised nature of energy generation and its close connection with its distribution clearly suggest the presence of features of a network of complex systems, with evident CPS modelling opportunities.

Transmission System Operators (TSO) have to ensure the physical balancing of the power grid (the total electricity generation must match the total electricity consumption). In AC electrical grids, the frequency (50 Hz in nominal conditions) indicates whether or not the system is balanced. More and more Photo-Voltaic (PV) Panels are installed in distribution grids and they are not directly controllable by the TSO (unless organised in structured, large PV farms). In order to assess the safe operation and robustness of the whole system, TSO ought to take PV panels dynamical behaviours into account.

Recent research has attempted to formalise and implement a formal study of large populations of PV Panels. We have focused on the modelling of the dynamics of PV panels and their interaction with the power grid as Markov models. Within the broad goal of formal aggregation of large-scale populations of Markov models, and the objective to provide new computational algorithms for the optimal policy synthesis over such a population, we have looked at aggregating and controlling large populations of photovoltaic panels over the power grid.

We plan to again employ techniques from formal methods to generate quantitative abstractions of models of interest (large populations of PV panels), with specific predictive capabilities and a guaranteed error on the quality of

the abstractions. Among other goals, we are investigating issues of decentralised control of such PV, of emergence of global behaviours from local conditions, and in general issues of robustness, dependability, and reliability over the grid.

6 Conclusions

The increase in complexity, adaptability, and inter-connectivity of modern technological devices raises new challenges towards their understanding and that of their complex interaction network. Likewise, their local actuation or the control of the global network they are part of, pose new challenges at engineering and technical levels. We have argued for the necessary development of formal verification approaches to study networks of complex dynamical systems, underpinned by quantitative models that are at the core stochastic and hybrid, built from data and formally reasoned upon.

We envision the development of semi-automatic approaches, based on formal abstractions, for the synthesis of policies around quantitative specifications encompassing formal requirements (e.g., safety, reliability), and trading off against performance, as well as the development of formal verification tools accounting for robustness. Such approaches are deemed "formal" in that they are enhanced by quantitative guarantees on their outcomes, being it an assertion over a formal property or the implementation of a control policy towards a given objective.

Furthermore, in view of the unavoidable connection of the systems of interest with data, we have further advocated more research on a tighter and formal integration of data-driven, sample-based approaches, with model-based deductive techniques: we view this integration as necessary not only to encompass adaptability features for the underlying models towards learning, but also to increase the scalability of formal verification techniques towards reasoning.

In conclusion, the technical challenges related to the existing technological trend towards more complex, adaptable and more integrated devices, can be offset by engineering and economic benefits, provided principled approaches for integrated data-based modelling, quantitative formal verification, and correct-by-design synthesis are embraced.

Acknowledgments. This is a position article accompanying a keynote talk at NSV 2016. It exclusively portrays the perspective of the author, without claims of generality or comprehensiveness. As such, references are limited to own work.

The author would like to thank former and current students involved in this specific area of work, particularly Sofie Haesaert, Sadegh E.S. Soudjani, and Majid Zamani; and collaborators on these topics, particularly Martie Fränzle, Joost-Pieter Katoen, Daniel Kroening, Marta Kwiatkovska, John Lygeros, Rupak Majumdar, Maria Prandini, and Claire Tomlin.

References

1. Abate, A.: Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: a survey. Electron. Notes Theor. Comput. Sci. **297**, 3–25 (2012)
2. Abate, A., D'Innocenzo, A., Di Benedetto, M.D.: Approximate abstractions of stochastic hybrid systems. IEEE Trans. Autom. Control **56**(11), 2688–2694 (2011)
3. Abate, A., Katoen, J.-P., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. Eur. J. Control **6**, 624–641 (2010)
4. Abate, A., Kwiatkowska, M., Norman, G., Parker, D.: Probabilistic model checking of labelled markov processes via finite approximate bisimulations. In: Breugel, F., Kashefi, E., Palamidessi, C., Rutten, J. (eds.) Horizons of the Mind. A Tribute to Prakash Panangaden. LNCS, vol. 8464, pp. 40–58. Springer, Heidelberg (2014). doi:10.1007/978-3-319-06880-0_2
5. Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. Automatica **44**(11), 2724–2734 (2008)
6. Various authors. FP7 AMBI project. http://www.ambi-project.eu
7. Various authors. FP7 MoVeS project. http://www.movesproject.eu
8. Ding, J., Kamgarpour, M., Summers, S., Abate, A., Lygeros, J., Tomlin, C.J.: A stochastic games framework for verification and control of discrete-time stochastic hybrid systems. Automatica **49**(9), 2665–2674 (2013)
9. Esmaeil Zadeh Soudjani, S., Abate, A.: Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. SIAM J. Appl. Dyn. Syst. **12**(2), 921–956 (2013)
10. Esmaeil Zadeh Soudjani, S., Abate, A.: Aggregation of thermostatically controlled loads by formal abstractions. In: European Control Conference, Zurich, Switzerland, pp. 4232–4237, July 2013
11. Esmaeil Zadeh Soudjani, S., Abate, A.: Precise approximations of the probability distribution of a Markov process in time: an application to probabilistic invariance. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. LNCS, vol. 8413, pp. 547–561. Springer, Heidelberg (2014). doi:10.1007/978-3-642-54862-8_45
12. Esmaeil Zadeh Soudjani, S., Abate, A.: Probabilistic reach-avoid computation for partially-degenerate stochastic processes. IEEE Trans. Autom. Control **59**(2), 528–534 (2014)
13. Soudjani, S.E.Z., Gevaerts, C., Abate, A.: FAUST^2: formal abstractions of uncountable-state stochastic processes. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 272–286. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46681-0_23
14. Haesaert, S., Abate, A., Van Den Hof, P.M.J.: Data-driven and model-based verification: a Bayesian identification approach. In: Proceedings of the IEEE Conference on Decision and Control, pp. 6830–6835 (2016)
15. Haesaert, S., Abate, A., Hof, P.M.J.: Verification of general Markov decision processes by approximate similarity relations and policy refinement. In: Agha, G., Houdt, B. (eds.) QEST 2016. LNCS, vol. 9826, pp. 227–243. Springer, Heidelberg (2016). doi:10.1007/978-3-319-43425-4_16
16. Kamgarpour, M., Ellen, C., Esmaeil Zadeh Soudjani, S., Gerwinn, S., Mathieu, J.L., Mullner, N., Abate, A., Callaway, D.S., Fränzle, M., Lygeros, J.: Modeling options for demand side participation of thermostatically controlled loads. In: International Conference on Bulk Power System Dynamics and Control (IREP), pp. 1–15, August 2013

17. Esmaeil Zadeh Soudjani, S., Abate, A.: Quantitative approximation of the probability distribution of a Markov process by formal abstractions. Logical Methods Comput. Sci. **11**(3), 1–29 (2015)
18. Tkachev, I., Abate, A.: Characterization and computation of infinite-horizon specifications over Markov processes. Theoret. Comput. Sci. **515**, 1–18 (2014)
19. Tkachev, I., Mereacre, A., Katoen, J.-P., Abate, A.: Quantitative model checking of controlled discrete-time Markov processes. Inform. Comput. (2016). doi:10.1016/j.ic.2016.11.006
20. Zamani, M., Abate, A.: Symbolic models for randomly switched stochastic systems. Syst. Control Lett. **69**, 38–46 (2014)
21. Zamani, M., Abate, A., Girard, A.: Symbolic models for stochastic switched systems: a discretization and a discretization-free approach. Automatica **55**(5), 183–196 (2015)
22. Zamani, M., Mohajerin Esfahani, P., Majumdar, R., Abate, A., Lygeros, J.: Symbolic control of stochastic systems via approximately bisimilar finite abstractions. IEEE Trans. Autom. Control **59**(12), 2825–2830 (2014)
23. Zamani, M., Mazo, M., Abate, A.: Finite abstractions of networked control systems. In: Proceedings of the 53rd IEEE Conference on Decision and Control, pp. 95–100 (2014)