# Model-based Evaluation of Energy Saving Systems

Davide Basile, Felicita Di Giandomenico, and Stefania Gnesi

**Abstract** Nowadays, there is a great attention towards cautious usage of energy sources to be employed in disparate application domains, including critical infrastructures, to save both in financial terms and in environmental impact. This chapter focuses on stochastic model-based as a support to the analysis of energy saving systems, in combination with other non functional properties, such as reliability, safety and availability. We discuss general guidelines to build a model-based framework to analyse critical cyber-physical systems, where effective energy consumption is required, while assuring imposed levels of resilience. Also, an overview of the most commonly employed methodologies and tools for model-based analysis is provided, and extensive literature is indicated as pointers to relevant research activities performed on this attractive topic over the last decades. Finally, in order to corroborate the proposed framework, a case study in the railway domain is proposed. By adopting the Stochastic Activity Networks formalism, the framework is instantiated to analyse effective trade-offs between energy consumption and satisfaction of other dependability related requirements.

**Key words:** energy-saving, reliability, quality models, stochastic analysis

## 1 Introduction

*Energy management* and the related environmental issues are emerging as a relevant technological, political, and societal concern. Reducing energy consumption is nowadays one of the primary goals in the design of IT systems [19]. The benefits of adopting strategies of energy management are many, as for example: *reduction of costs*, by adopting strategies of energy management most organisations can save in economic terms; *mitigating adverse environmental effects*, which helps organi-

Institute of Information Science and Technologies (ISTI), National Research Council (CNR), Pisa, e-mail: {davide.basile, felicita.digiandomenico, stefania.gnesi}@isti.cnr.it

zations in gaining a good reputation and avoid penalties of governments aiming at reducing the environmental pollution; *reduction of risks* due to energy shortages and prices fluctuations.

For being effective, policies of energy saving need to be integrated in the design of an energy-saving system. When critical systems are considered, that is employed in sectors where *dependability* requirements are stringent, the challenge of energy saving needs to be mediated with satisfaction of properties such as *reliability*, *safety*, *availability*. Unfortunately, reducing the energy consumption and increasing reliability are in general contrasting requirements. Hence, it is important to find a good trade-off between the two aspects. To this purpose, resorting to *quantitative assessment* of policies devised for energy management is extremely helpful to support the choice of the one that better suits the overall set of system requirements. Stochastic model based analysis, which is widely recognized as a powerful approach for quantitatively assessing non functional properties since the early phases of a system development, is addressed in this work.

In particular, this chapter offers the following contributions. First, guidelines to the set up of a stochastic model-based framework, suitable to evaluate trade-offs between energy consumption and dependability-related properties in *critical energy-saving cyber-physical systems*, are discussed. The targeted systems are characterised by physical entities that are controlled by computational units, where an energy management policy is in place to supply the physical components as needed to reliably perform their work. Comparison among different energy management strategies is possible, in terms of indicators representative of the energy consumption and dependability-related indicators, to select the most effective one.

Second, a review of widely adopted formalisms and tools for stochastic model-based evaluation is presented, to provide a basis of existing instruments to implement the proposed analysis framework.

Third, a realistic case study is proposed to exploit our modelling framework, which is taken from the railway transport domain, a critical domain where failures can lead to catastrophic consequences, as for example the derailment of trains. In particular, we evaluate policies of energy consumption for a system of rail road switch heaters. By adopting the proposed method, evaluation of the measures of interest is used for tuning the parameters of the considered energy-saving policies in order to reduce energy consumption while satisfying reliability constraints.

## Related work

Model-based analysis of energy-saving reliable systems is recently gaining interest from the research community, hence the literature concerning this subject is relatively heterogeneous; in the following we briefly review recent works that analyse and optimise the energy consumption in several application domains. Generalized Stochastic Petri Nets [3] are used to solve the dynamic power management problem for systems with complex behaviour in [30]. Dynamic power management addresses reduction of power dissipation in embedded systems, with a selective shut-off or

slow-down of system components that are idle or underutilized. A time-out policy is used for power saving, which turns on a component when it is used and turns it off when it is not used for a certain amount of time. Comparisons are also performed with other models based on Markov Decision Processes (MDP) [29]. The dynamic power management problem is interpreted as a hybrid automaton control problem and integrated stochastic control in [18]. Hybrid automata mixed both a discrete state, representing the power mode of the system, and a continuous one, representing the consumed power. Two strategies are compared: on demand wake-up of a component (that was previously turned off) and pre-emptive wake-up. The former provides better results for conservation of energy and prevention of latency. The applicability of self-organizing systems for different fields of power system control is discussed in [26]. Agent-based decentralized power flow control is compared with current practice based on central decision making. The authors study how to balance the voltage and frequencies stability of the network to meet the demand of energy. These parameters are linked to reliability and safety of the system. It is shown how a decentralized control can improve reliability, safety and efficiency by providing a real-time adaptivity to changes in the network (failure of a node, blackout). The survivability of a smart house is analysed in [21], that is the probability that a house with locally generated energy (photovoltaic) and a battery storage can continuously be powered in case of a grid failure. Hybrid Petri Nets [16] are used for modelling this scenario. Different strategies of battery management are considered. In the first one, all the battery is consumed when needed, in the second there is a minimum threshold of energy saved in case of grid failure. In the third case the battery is also charged to a maximum threshold when the grid is operating. It is shown how the third strategy is better both for the local usage of energy and for the survivability of the smart house. The authors consider a randomly chosen probability of failure and fixed thresholds. The trade-off between energy saving and reliability is studied in [40], by managing frequencies and voltage of the delivered energy. In particular, lower frequencies result in higher reliability while for voltage scaling the reliability decreases dramatically. In our approach the energy consumption is managed by changing the power consumed by the system.

Power management in smart grids is discussed in [25], where it is shown that, by adjusting the power supplied to the different clients, it is possible to obtain a good trade-off between power utilization and customers satisfaction. Services negotiation of energy and reliability requirements is the selected case study in [39], where an energy provider, an energy consumer and a mediator try to find an agreement on the amount of energy delivered, its reliability and price. In Section 2.3 we will review some of the previously discussed models for energy assessment.

Structure of the chapter.

Section 2 discusses general guidelines for supporting the design and evaluation of energy-saving systems and provides a brief survey on the modelling methodologies and tools for evaluating dependability measures and energy consumption of sys-

tems. A case study is proposed in Section 3 where the proposed method is applied for evaluating reliability and energy consumption indicators for the system under analysis. Conclusions are discussed in Section 4.

## 2 Supporting the Design of Energy Saving Systems

In this section, we introduce general guidelines toward the building of a modelling framework to support the design of critical energy saving systems. To provide the context useful to understand the developments, background on model-based analysis is firstly introduced. Then, a general overview of the proposed modelling framework is presented, followed by an overview of the pertinent formalisms and tools from the literature. To concretely demonstrate the potentialities of the proposed approach, instantiation of the general framework in a use case from the railway domain is then addressed, as extensively described in Section 3.

### *2.1 Model-based Analysis*

Quantitative assessment of non-functional properties, especially dependability and performance related ones, is an important activity when developing systems to be employed in critical domains, where strict requirements on these properties need to be met. Such evaluation typically consists in probabilistically estimate the occurrence of faults and their impact on the ability of the system to operate correctly, in presence of possible measures to enhance system resilience (such as fault tolerance mechanisms). Several approaches are available in literature to perform system assessment [2], mainly testing, fault injection and model-based evaluation. For quantitative evaluation of dependability indicators, stochastic model-based analysis [35] has been proven to be particularly useful, versatile and cost-effective for manufacturers [8, 17]. To keep the model manageable, the system needs to be represented at a properly identified abstraction level. Indeed, depending on the properties to be analysed the emphasis on the system representation is focused on those aspects that are relevant for the analysis purposes, while irrelevant aspects are neglected. Therefore, a wide variety of models are used in practice, to tailor the right abstraction level for the system under analysis, in accordance with the properties to be assessed, the desired degree of accuracy and available resources to manage models development and solution.

Stochastic model-based approaches [35] are useful to support the development of systems, in all the phases of their life cycle. In the early design phases, they can be helpful for avoiding waste of time and resources in the development phase. This can be done by pointing out the properties and the requirements that the system must satisfy, as for example the energy consumed, and build a model that captures the relevant structural and behavioural aspects of the system under development. Such

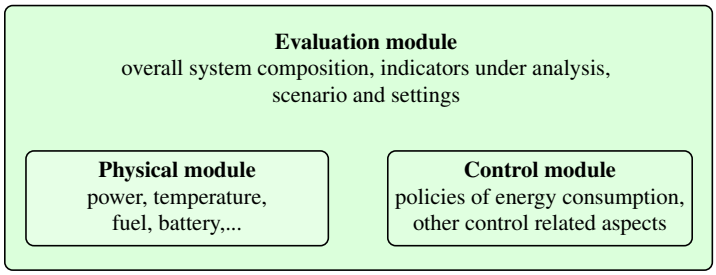


Fig. 1: The proposed analysis framework for energy saving CPS

early modelling phase is therefore exploited to highlight problems in the design of the system, such as bottlenecks, to perform comparisons of different alternative architectural solutions and selection of the most suitable one. For an already existing system, an a-posteriori analysis of properties such as dependability or performance is useful to understand and learn about specific aspects, to detect possible design weak points, to perform a late validation of the dependability requirements, and it is especially useful to improve the system in its future releases. Moreover, with a model-based analysis it is possible to predict future behaviours to plan the maintenance and the upgrading of the system.

## *2.2 Guidelines to Model Critical Cyber-physical Energy Saving Systems*

By resorting to a model-based approach, we propose general guidelines to develop a modelling framework for supporting the design of critical energy saving systems. We target energy consumption in cyber-physical systems (CPS) [23], where physical entities are controlled by computational units. Among the functionalities under the supervision of the cyber-control, there is the strategy for supplying energy to components of the physical system, necessary to keep them effective and reliable in the service they accomplish. Our interest is in assessing measures that are representative of the energy consumption, to be combined with other dependability-related properties dictated by the critical domain the CPS is employed in. The results from this kind of analysis are very helpful for: i) gaining insights on benefits from different energy supply strategies and ii) properly tuning the parameters of these strategies toward most rewarding configurations.

A modular and compositional framework is pursued, to promote wide applicability to a variety of scenarios relevant to cyber-physical systems, as well as ability to undergo future refinements and extensions. A diagram of the proposed framework is depicted in Figure 1. The proposed analysis framework is built around three major modules, as outlined in the following.

- *Physical-aspects module*: this module focuses on the physical components of the system and on their characterisation in terms of relevant aspects from the energy viewpoint. It includes models representing phenomena related with energy supply, which depend on the fabric of the supplied components and environmental conditions impacting on the energy consumption. Examples are: i) internal and external temperatures, properly modelled taking into account their evolution in time, given the different means involved (such as iron or copper for the physical components, winter or summer days for the external air); ii) fuel consumption, represented by properly considering the engine parameters, aerodynamic drag, weight, and other relevant parameters; iii) supplied power, represented by properly considering the laws regulating the involved real process; iv) battery charging and others.
- *Control-aspects module*: this module deals with the policies that dynamically regulate the energy consumption of the physical components. To manage potential complexity while assuring adequate accuracy of the analysis, the representation of such policies is abstracted at the level of their impact on the energy parameters of the controlled physical components. As already mentioned, the primary objective of the proposed model-based approach is to assist the system designer in identifying the best policy to employ among several alternatives, in accordance with pre-established dependability and cost requirements. Trade-offs between energy consumption and dependability requirements are mandatory in critical domains, where, e.g., the energy should not be reduced in safety critical situations. A typical family of energy supply algorithms is of kind on-off, where energy is supplied or turned off on the basis of values assumed by parameters that depend on physical conditions of the system and of the environment.
- *Evaluation module:* Finally, the third module deals with the composition of the several models from the previous two modules, to end up with the overall evaluation framework. Exercising the overall composed model, energy supply policies trading energy saving and dependability properties can be quantitatively evaluated and compared in terms of properly defined indicators.

The generality of the above outlined approach allows assessing a variety of measures of interest to final customers, service providers and operators, in accordance with the specific application domain where the CPS system under analysis is utilised. Given the aim of trading energy consumption with dependability, typical indicators are energy supplied to individual system components or to the overall system in a certain time interval, as well as failure probability of an individual component or of the overall system due to lack of supplied energy.

In Section 3 we will provide an example of an energy saving system designed with this method [6, 5].

## 2.3 Methodologies and Tools for Model-based System Analysis

Assessing the ability of a developed system to fulfil its requirements (both functional and non-functional) is a fundamental phase in the system engineering development process. As already mentioned, model-based evaluation is a versatile support since the early phases of system development and a large variety of modelling formalisms and tools have been developed over the past decades. The interest towards energy consumption analysis gained great attention in very recent years, so there is a relatively limited number of works addressing this aspect. The review below offers a reference to existing modelling formalisms and tools which can be fruitfully employed also to this purpose.

In particular, we briefly describe a subset of *state-space oriented* modelling formalisms, that are used for quantifying non-functional properties (typically, dependability and performance indicators) at different stages of a system development process.

Both *deterministic* and *stochastic* modelling formalisms have been defined, to properly address the different deterministic or probabilistic behaviours of systems to be analysed. Concerning the evaluation of energy saving systems and dependability attributes, stochastic models are more adequate. Indeed, they are capable of expressing the uncertainty about the occurrence of events as failures, and for modelling external factors affecting the energy consumption, as for example weather conditions.

According to the underlying stochastic process, stochastic models can be further classified into *Markovian* and *non-Markovian* [22]. Markovian models are those satisfying the *Markov property*, that is the conditional probability distribution of future states (conditional on both past and present values) depends only upon the present state. Indeed, in Markovian models the occurrence of events always has an exponential distribution of time. For modelling systems without this property non-Markovian models are used. For the discussed modelling methodologies, several tools are available for building and solving them [27, 38, 24].

## 2.4 Markovian Models

Markovian models are stochastic models where it is assumed that future states depend only on the present state. Among the others we discuss *Markov chains, Stochastic Petri Nets, Generalised Stochastic Petri Nets*.

### Markov Chains

Markov chains [22] are capable of modelling several types of systems and they are equipped with efficient solution algorithms implemented in many automated tools. With this formalism the states of a system are modelled with a random variable $X(t)$

that changes through time. Thanks to the Markov property, the distribution for this variable depends only on the distribution of the previous state.

When only the elapsed time is important and not the actual instance of time, then the Markov chain is said to be homogeneous; otherwise, it is said to be non-homogeneous. If the parameter $t$ that indexes the Markov chain is continuous, we have a continuous-time Markov chain; otherwise, we have a discrete-time Markov chain.

A graphical representation of a Markov chain is a directed graph called *state-transition diagram*, where nodes represent states and arcs represent state transitions. Because of the memoryless property, each transition occurs in an exponentially distributed time, and the rate is exactly the inverse of the expected time to the transition, that is, the rate of the corresponding exponential distribution. This is the most severe constraint that limits the applicability of Markov chains. Markov chains can be solved in terms of transient or steady-state analyses. Once the state occupation probabilities of the Markov chain are obtained, the values of the most important dependability measures can be evaluated [8]. A reward structure can also be defined, which assigns reward values to the states (or to the transitions) of the Markov chain model, obtaining the Markov reward models [32], which among the others can be used to evaluate the energy consumption of systems.

### Stochastic Petri Nets

Petri Nets are widely used as modelling paradigm thanks to the efficiency in describing the qualitative and quantitative aspects of complex systems and because of their intuitive and appealing graphical representation. Originally Petri Nets were developed for representing in a compact way concurrency among processes [28].

Stochastic Petri Nets [7] are a popular timed extension of place-transition Petri Nets [28], where places are represented by circles and timed exponential transitions by white or black rectangular boxes connected by directed arcs. Arcs have positive integer weights associated and are of two types: input arcs connect input places to transitions while output arcs connect transitions to output places. Places contain tokens and the whole state of the net, called *marking*, is represented by the number of tokens in each place. Each transition has associated a random firing delay with negative exponential distribution time (Markov property).

When the marking of the input places matches the weights on the input arcs then the transition is enabled, and a timer starts counting the sampled time. If the transition remains continuously enabled the whole time, it will fire by removing the tokens from the input places and adding them to the output places according to the weights of the arcs. Once the marking changes, some transitions can be enabled and others aborted. The set of all possible reachable markings is called the *reachability set*. The *reachability graph* describes all the possible evolutions of the network starting from the initial marking, where nodes are possible markings and arcs are transitions causing the change of marking.

The reachability graph of a network is isomorphic to the correspondent continuous time Markov chain, where states and transitions are in one-to-one correspondence with those of the reachability graph, and the rates of the transitions are equal to the firing rates of the transitions of the network which cause the change of markings. Hence, in order to solve a stochastic Petri Nets it suffices to analyse the associated Markov chain, and the corresponding reward structure [32].

#### Generalized Stochastic Petri Nets

Generalized Stochastic Petri Nets [3] enrich stochastic Petri Net with instantaneous transitions and inhibitor arcs. The first, graphically represented as thin bars, once enabled fire in zero time. Conflicts among instantaneous transitions are solved by associating priorities and weights to transitions with the same priority. Inhibitor arcs, graphically depicted as arcs terminating with a circle, inhibit the firing of the connected transition when the input place of the inhibitor arc has one or more tokens. The reachability graph of this type of networks is not in correspondence with Markov chains because of the instantaneous transitions. Indeed, now the network spend zero time in some markings, which are called *vanishing*. Nevertheless it is possible to eliminate those vanishing markings by performing some approximations, in order to obtain a reduced reachability graph which is isomorphic to a Markov chain. The reward structure can be associated to the underlying Markov chain (Markov reward models) [32].

## *2.5 Non-Markovian Models*

In Generalised Stochastic Petri Nets, when transitions with non exponential duration are represented (i.e. instantaneous activity) then an approximation is introduced in order to solve them. This constraint (Markov property) may be too strict for certain classes of systems, and for them several classes of non-Markovian models have been introduced in literature [9], for which it is possible to represent generally distributed activities and that can be solved by simulation, or analytically; by approximating non-exponential random variables with exponential ones [35, 33, 9]. Here we discuss *Semi-Markov Stochastic Petri Nets* and *Stochastic Activity Network*.

#### Semi-Markov Stochastic Petri Nets

Semi-Markov Stochastic Petri Nets are an extension of Generalised Stochastic Petri Nets for which the firing time of a transition is generally distributed (non-Markovian property), and with the *resampling memory policy* for transitions, where each time a transition fires the remaining firing time of the other transitions is resampled. This nets are in correspondence with Semi-Markov Stochastic Processes [20, 36], where

the holding time of states is generally distributed and the instants of time a timed transition fires are regeneration points. Indeed, after a regeneration point the future of the marking process becomes a replica of the process at time zero. The evolution of the Semi-Markov Stochastic Process is studied for transient and steady-state analyses [36].

Stochastic Activity Network

Stochastic Activity Networks [34] are a formalism widely used for performance, dependability and performability evaluation of complex systems. The SAN formalism is a variant of Stochastic Petri Nets [7], and has similarities with Generalised Stochastic Petri Nets [3]. A SAN is composed of the following primitives: *places, activities, input gates* and *output gates*. Places and activities have the same interpretation as places and transitions of Petri Nets. Input gates control the enabling conditions of an activity and define the change of marking when an activity completes. Output gates define the change of marking upon completion of the activity. Each enabled activity may complete. Activities are of two types: *instantaneous* and *timed*. Instantaneous activities complete once the enabling conditions are satisfied. Timed activities take an amount of time to complete following a general temporal stochastic distribution function. An enabled activity is aborted, i.e. it cannot complete, when the SAN moves into a new marking in which the enabling conditions of the activity no longer hold. Cases are associated to activities, and are used to represent probabilistic uncertainty about the action taken upon completion of the activity. When an activity completes, the following steps are executed:

- one of the cases of the activity is chosen according to its marking-depending probability;
- the function of each input gate of the activity is executed;
- the function of each output gate linked to the case selected at first step is executed.

The primitives of the SAN models are defined using C++ code.

## *2.6 List of Tools*

We now overview a small subset of the vast body of tools for automatising the description and evaluation of stochastic models.

- SHARPE (Symbolic Hierarchical Automated Reliability and Performance Evaluator, `http://sharpe.pratt.duke.edu/node/1`) [33] is a toolkit that provides a specification language and solution methods for models described as, among the others, Markov chains and Generalized Stochastic Petri Nets. Steady-state, transient and interval measures can be computed. For most of the model types, SHARPE provides more than one analysis algorithm from which the user

can choose. SHARPE allows the user to combine models in a hierarchical fashion. The model types include state-space ones such as Markov and semi-Markov reward models and Stochastic Petri Nets;
- SPNP (Stochastic Petri Net Package) [14] supports models as Stochastic Rewards Petri Nets and Generalised Stochastic Petri Nets. It supports transient and steady-state analytical evaluations for Stochastic Rewards Petri Nets models and discrete-event simulation for non-Markovian models. More informations at: `http://people.ee.duke.edu/~kst/chirel/IRISA/spnp.html`;
- SURF-2 (`http://homepages.laas.fr/surf4tst/what-uk.html`) is a dependability evaluation tool for hardware and software systems, based on strict construction, validation and numerical resolution of Markov models. The available modelling formalisms are Markov chains and Generalized Stochastic Petri Nets. Reward structures are used to obtain measures of dependability, performance or cost. It supports transient and steady-state analysis;
- GreatSPN (`http://www.di.unito.it/~greatspn/index.html`) [13] supports formalisms as Generalized Stochastic Petri Nets and Stochastic Well-formed Nets, and it allows the composition of different models. An algorithm for the fast computation of performance bounds based on linear programming techniques is available; together with algorithms for the analysis of stochastic well-formed nets. Steady state and simulation solvers are available;
- Oris (`http://www.oris-tool.org/`) [11] is a tool which supports a variety of timed extensions of Petri Nets, and it implements a symbolic state space analysis which enables an integrated approach to qualitative verification and quantitative evaluation;
- Möbius (`https://www.mobius.illinois.edu/`) [15] is a tool that supports various formalisms such as SAN, PEPA (a process algebra), Fault Tree, and different analytical and simulative solvers. Möbius can be used for studying the reliability, availability, and performability of systems. It follows a modular modelling approach, with proper operators *Rep* and *Join* to compose atomic models into an overall composed model.

## 3 A Case Study of Energy Saving System

In this section we tailor the general method to build evaluation frameworks discussed in Section 2.2 to a scenario of energy saving in the railway domain. Following the spirit of the proposed method, a modular and parametric approach is followed, to assure usability of the developed analysis framework in a variety of system configurations, as well as to promote extension and refinements of the model itself to account for further involved aspects/phenomena and so enhance its adherence to sophisticated and realistic implementations with respect to the current version.

Specifically, the considered case study is a rail road switch heating system. A rail road switch is a mechanism enabling trains to be guided from one track to another. It works with a pair of linked tapering rails, known as points. These points can be

moved laterally into different positions, in order to direct a train into the straight path or the diverging path. Such switches are therefore critical components in the railway domain, since reliability of the railway transportation system highly depends on their correct operation, in absence of which potentially catastrophic consequences may be generated.

Unfortunately, during winter, snow and ice can prevent the switches to work properly, exposing the railway network to potential failures. In the past, the switches were kept clean manually by employees who were sweeping the snow away. Nowadays, heaters are used so that the temperature of the rail road switches can be kept above freezing. The heaters may be powered by gas or electricity; for simplicity, in the next description we refer to electrical supply.

Different policies may be adopted to power the heaters, as for example to heat a selection of switches for a given amount of time or to heat all the switches together. Of course, different policies imply different energy consumption and different probabilities to expose the switches to a malfunction because of lack of sufficient heating.

The proposed analysis contributes to gain insight on the interplay between energy consumption and reliability in order to select an appropriate policy for the heating of the switches, which guarantees a satisfactory trade-off. In particular, the (electric) energy consumption of the system is the amount of power consumed by the system in a unit of time, that is measured in Watt per hours, while reliability is defined as the continuity of correct service, i.e. the ability of the system to avoid service failures that are more frequent or more severe than is acceptable [2].

In the following, we describe the steps to set up and exercise the modelling framework for the evaluation of reliability and energy consumption indicators for a system of (remotely controlled) rail road switch heaters. An on-off policy is considered for heating the switches, with parametric thresholds representing the temperatures triggering the activation/deactivation of the heating. The management of the heaters is automatic, and is remotely controlled by a central computational unit. In a railway station there are tracks which are less important than others, for example the side tracks. In case of extremely cold conditions, the total amount of energy available could not be sufficient to heat the overall system, hence it is important to duly choose the heaters that must be primarily heated and those that may be heated later on. Those rail road switch heaters whose temperature cannot be kept above the freezing thresholds will experience a failure.

Since the analysed system is characterized by stochastic phenomena that cannot be accurately represented by the exponential distribution, (e.g. the time regulating changes of the temperature of the switches), we have resorted to non-Markovian models, in particular the ones defined by Stochastic Activity Networks (SAN) [34], and the Möbius tool (see Section 2.3) for evaluating the measures of interest. We remark that the proposed framework allows to model and analyse the system with different formalisms; however due to lack of space we only resort to SAN models.

### *3.1 Physical-aspects Module*

We briefly describe the real-world devices that constitute the physical components of the system under analysis. The heating system for a single heater consists in a series of tubular flat heaters along the rail road track, which warm up the rail road by induction heating. To accomplish its task, the rail road switch heater system reads through sensors the temperatures of the air and of the rail road [31].

The physical behaviour of the rail road is modelled in terms of temperature decay and increase, when the heating is switched off and on, respectively. For the temperature of the air, statistic data about cold winter nights are used, while for modelling the internal temperature of the device and to estimate the energy consumption of the heater, the exchange of heat through convection is modelled. Indeed the rail road gets cooled by the external temperature and warmed by the heaters.

To make the needed calculation we consider the portion of the rail road track to be heated, which for simplicity is represented by an iron bar. We assume that the bar is exposed to the external temperature both from the top and the bottom.

The heater is represented by a resistance that passes through the rail road in different points in order to warm up the iron. The set-up for the heating device is based on patents of heating switches [10], which also contain data about the power consumed by a single heater and the increment of the temperature of the track in cold winter nights. We assume that the power used by the heater is constant, in order to estimate the kilowatt per hours consumed during the time interval that we consider.

Every hour new data for the internal temperature of the rail road track are computed. Assuming that the values of the external temperature of the surrounding area $T_e$ and the previous internal temperature $T_i$ are known, we foresee the updated internal temperature $T$ of the rail road by solving the differential equation (balance of energy) representing the exchange of heat by convection [12]:

$$mc\frac{\partial T}{\partial t} = -uA(T - T_e) + \dot{Q},$$

where $u$ is the coefficient of convective exchange, $c$ the heat capacity of iron; $A$ the surface area exposed to the external temperature; $t$ the interval of time, one hour in our case; $m$ the mass of the iron bar; $\dot{Q}$ the power used when the heater is turned on, if the heater is turned off this value will be zero.

We remark that in the physical phenomenon we have studied (i.e. heating exchange), the unknown variable is the internal temperature at instant of time $(k+1)*t$, while the external temperature and the internal temperature at time $k*t$ are known.

### *3.2 Control-aspects Module*

We identify two main logical components in the control subsystem: the *heater* and the *central coordinator*. The network of heaters is realised by replicating the heater component, and the activation/deactivation of each heater is controlled by the central coordinator. We now present these two main components.

- *Heater*: we based the policy employed to activate/deactivate the heating on two threshold temperatures:
  - *warning threshold* ($T_{wa}$): this temperature represents the lower temperature that the track should not trespass. If the temperature is lower than $T_{wa}$, then the risk of ice or snow can lead to a failure of the rail road switch and therefore the heating system needs to be activated;
  - *working threshold* ($T_{wo}$): this is the working temperature of the heating system. Once this temperature is reached, the heating system can be safely turned off in order to avoid an excessive waste of energy.

  The energy consumption of the overall system depends on the value of $T_{wa}$ and $T_{wo}$. Reducing the gap between these thresholds will result in a more frequent activation of the heating system, but for a shorter period of time.
- *Coordinator*: the coordinator collects the requests of activation from the pending heaters, and it manages the energy supply according to a FIFO prioritized order. Indeed, the first heater which asks to be turned on will be the first to be activated. We assign priorities to switches based on their criticality on the track; the purpose of considering priorities is to guarantee higher reliability to those switches that are vital for the correct functioning of the overall station. The percentage of heaters that can be turned on at the same time is called $NH_{max}$. This value represents the maximum amount of energy deliverable by the system, and cannot be exceeded. If there is no energy available, each request will be enqueued in the queue of pending heaters.

We now give details about the exchange of messages between the network of heaters and the central coordinator:

- *Heater*: at starting time each heater $h_i$ is switched off and its internal temperature is above $T_{wa}$. Once the internal temperature goes below $T_{wa}$, $h_i$ asks the coordinator to be turned on and waits. When notified, $h_i$ is turned on. After that, two events can happen:
  - the heater $h_i$ reaches an internal temperature above $T_{wo}$, communicates to the central coordinator the termination of the heating phase and is switched off;
  - a second component $h_j$ with a higher priority asks to be turned on. The energy delivered to $h_i$ is turned off, even though it has not yet reached an internal temperature above $T_{wo}$. If the temperature is below $T_{wa}$, $h_i$ will issue a new request of activation to the coordinator.
- *Coordinator*: at starting time the central coordinator is waiting for a message from one of the heaters $h_i$ in the network. Two messages can be received:

- $h_i$ asks to be activated. This request is inserted in the queue of pending requests in case there is no energy available and the priority of $h_i$ is not higher than that of the already activated switches. Otherwise, the request is accepted and the coordinator switches to a busy state. In this state, two events are possible. In case there is energy available, $h_i$ will be activated by issuing a notification. Otherwise, if no energy is available but $h_i$ has a priority higher than one of the activated heaters, firstly the heater with lower priority will be turned off with a negative notification, and then the activation is notified to $h_i$;
- $h_i$ asks to be deactivated. After the deactivation, if there are no heaters that are waiting for being activated then no action is performed. Otherwise, one of the pending heaters $h_j$ (the first in the queue of pending heaters with higher priority) is activated by issuing a notification to it.

### SAN models

The overall model is obtained by the composition of the atomic models, using the *Join* and *Rep* operators (see Section 2.6), as shown in Figure 2.

The atomic model *Coordinator* represents the central coordinator. In addition to the models representing the physical and heating policy of the railway switch heating system, there are other models that accomplish supporting operations. The atomic SAN models *ProfileSelector*, *LocalitySelector*, and *SwitchIDSelector* represent, respectively, the selector for the weather profile, the location of the switch and the unique identifier of each switch. The SAN model *RailRoadSwitchHeater* represents the rail road switch heater. The submodel *HeaterModuleM* represents an instance of a single heater module, obtained by the composition, using the join operator, of the four atomic SAN models. Those atomic models share the places relative to the locality of the device, its weather profile and the unique ID. The submodel *HeatersNetM*, obtained by replicating *numRep* times the model *HeatherModuleM*, represents the network of heaters, where the parameter *numRep* identifies the number of devices composing the network. Finally, the model *SwitchHeatingSysM*, obtained using the join operator, represents the overall system. Indeed all the submodels share the same coordinator.

All these SAN models interact through *shared places*, a feature available in the Möbius tool [15] for joining different SAN models thus allowing composability.

### Rail Road Switch Heater SAN model

In Figure 3 the SAN model representing the rail road switch heater is depicted. We identify three logical components inside this SAN model: the *init sub-net*, the *clock sub-net* and the *heater sub-net*.

*Init sub-net.* The *init* subnet initialises the C++ data structures based on the values of the places *locality*, *profileID*, *SwitchID* shared with the SAN selector models

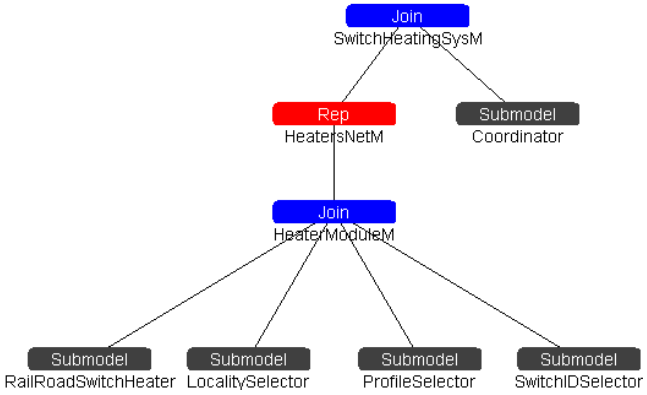


Fig. 2: The composed model.

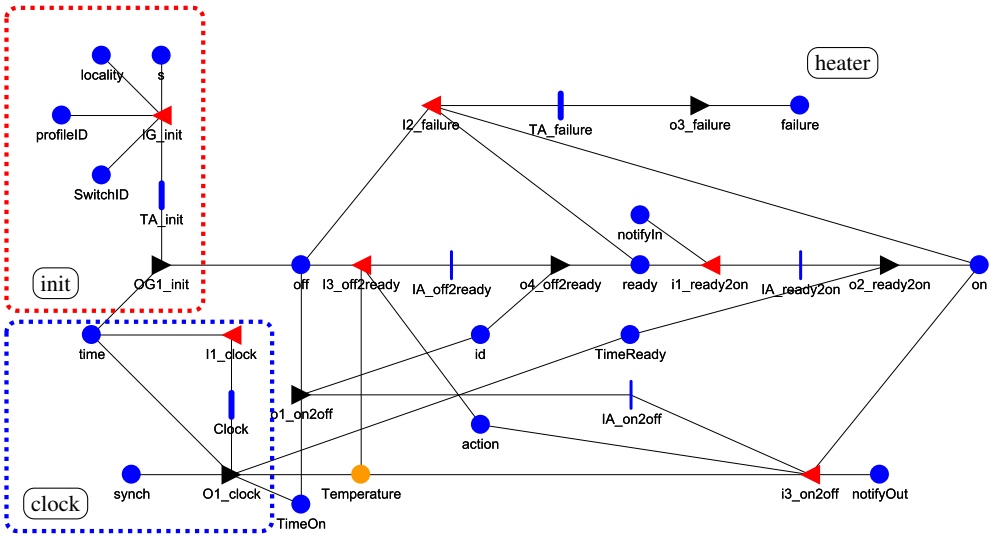


Fig. 3: The SAN model RailRoadSwitchHeater, logically divided into three sub-nets: the init sub-net, the clock sub-net and the heater sub-net.

(there is a bijection between the marking of the network and the data structures used in C++), and assigns a priority to the rail road switch heater.

*Clock sub-net.* The clock sub-net models the evolution of time (during one day in our analyses), and it is used to load the external temperature and the temperature of the rail road track. We consider as unit of time one hour. The activity *clock* has a deterministic distribution of time (non-Markovian) and completes each hour and updates (among the others) the places *TimeOn* and *TimeReady* of the heater sub-

net which are respectively counting the time that a heater is activated and the time that a heater is waiting. Moreover, when *clock* completes, the place *Temperature* is updated: if the heater is turned on then the temperature increases, otherwise the temperature will be updated according to the temperature of the environment. The function representing the heating exchange (see Physical-aspects Module) is defined in C++, and it is called by the output gate $O1_{clock}$ of this SAN model to update the temperature of the rail road track each interval of time $t$. We note that in this particular case study there is no neat distinction between the SAN models concerning the physical module and the one for the control module. Indeed the discussed SAN model captures aspects concerning both modules (i.e. energy policies, temperature).

*Heater sub-net.* The heater sub-net represents the status of the rail road switch heater. The heater can be activated (one token in the place *on*), waiting for being activated (one token in the place *ready*), turned off (one token in the place $off$), or failed (one token in the place $failure$). Indeed, according to the heating policy, once the system temperature goes below a pre-defined warning threshold ($T_{wa}$), the heating needs to be activated, otherwise the associated switch fails. Then, once the temperature raises and reaches the working threshold ($T_{wo}$), the heating system can be safely turned off.

The heater sub-net interacts with the Coordinator SAN model through places shared among all the replicas of the heater model and the Coordinator model implementing the logic described above. For example, if the heater is on state ready, in order to be turned on, the input gate $i1_{ready2on}$ checks if the marking of the shared place $notifyIn$ is equal to the marking of the place *SwitchID*, which means that the coordinator has notified the heater to be turned on.

The activity $TA_{failure}$ models the failure of a component. It has an exponential distribution of time based on the temperature of the rail road track: the more the temperature is below the freezing threshold the more is probable that the activity will fire (the activity is not activated if the temperature is positive).

For a comprehensive description of all the SAN models we refer the interested reader to [6, 5].

## *3.3 Evaluation Module*

In the evaluation module the overall composed model is instantiated with varying values of $T_{wa}$,$T_{wo}$ and $NH_{max}$, in order to find the best compromise between reliability and energy consumption. We analyse two different measures of interest. The first concerns the energy consumption while the second addresses the reliability of the system under analysis.

1 $CE(t,l)$: the time (number of hours) an heater is activated in the time interval $[t,t+l]$. This measure is defined by accumulating the marking of the place *on* of the *RailRoadSwitchHeater* net in the interval $[t,t+l]$, that is the time that each replica of the SAN model *RailRoadSwitchHeater* spends in the marking represented by one token in the place *ON*. Hence $CE(t,l)$ is the hours that an

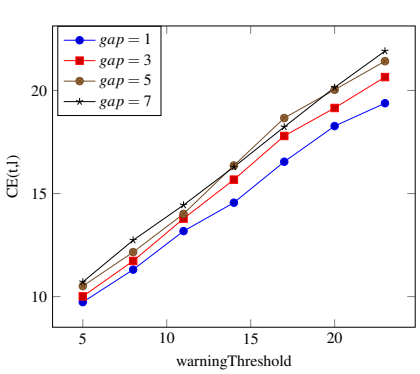


(a) The amount of energy consumed, where the different lines correspond to different gaps between the thresholds $T_{wa}$ and $T_{wo}$.

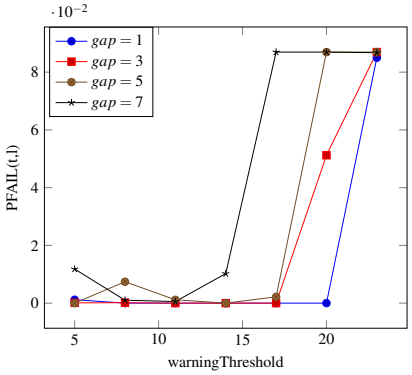


(b) The probability of failure, where the different lines correspond to different gaps between the thresholds $T_{wa}$ and $T_{wo}$.

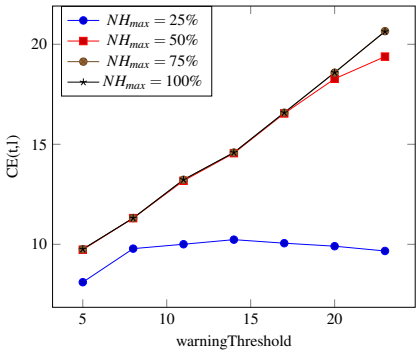


(c) The amount of energy consumed, where the different lines correspond to different levels of energy available to the whole system ($NH_{max}$).

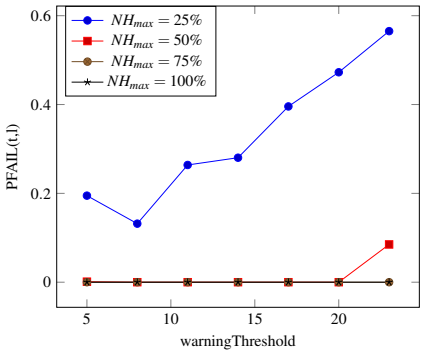


(d) The probability of failure, where the different lines correspond to different levels of energy available to the whole system ($NH_{max}$).

Fig. 4: The graphs for the measures of interest $CE(t,l)$ and $PFAIL(t,l)$.

heater is active. By multiplying $CE(t,l)$ for the power consumed (kilowatt per hour) it is possible to derive the energy consumed by the system;

2 $PFAIL(t,l)$: the probability that at least a switch fails (becomes frozen) at time $t+l$, given that at time $t$ is not failed. This measure is defined as the probability that at time $t+l$ there is one token in the place $failure$ of the SAN model $RailRoadSwitchHeater$.

We remark that reliability is computed as the probability that no failure occurs in the interval of time under analysis [37], that is $1-PFAIL(t,l)$. Simulation-based evaluations have been performed using Möbius tool [15], considering from a minimum of 1000 batches to a maximum of 10000 batches. The measures of interest were estimated within an interval of confidence of 0.95, and the model is instantiated to real-world data, available at [6, 5].

### *3.4 Results presentation*

Figure 4b and Figure 4a depicts respectively the measures of interest $PFAIL(t,l)$ and $CE(t,l)$ at the varying of the gap between $T_{wa}$ and $T_{wo}$, while Figure 4d and Figure 4c depicts respectively the measures of interest $PFAIL(t,l)$ and $CE(t,l)$ at the varying of the available energy $NH_{max}$.

From Figure 4 it is possible to observe the results of our evaluation. For values of the thresholds $T_{wa} = 7°\text{C}$, $T_{wo} = 8°\text{C}$, and energy available $NH_{max} = 50\%$ we obtain the optimal trade-off between energy consumption and reliability. Indeed, for values of $T_{wa}$ lower than 7 °C, as a result of the environment temperature, the internal temperature of the heater will fall quickly under zero, hence increasing the probability of failure. Instead, for values greater than 7 °C, the active heaters jeopardize all the energy available, leaving the other pending heaters waiting for too long, hence increasing the probability of a failure. Since the kilowatts consumed by each heater are constant, an increment in the gap between the thresholds will result in a longer period of activation of the heater. Finally, in Figure 4 we only have analysed the high-priority switches, which are predominant for this case study, and in particular $NH_{max} = 50\%$ is enough for heating all the high-priority switches; while for lower values of $NH_{max}$ we have an increment in the probability of failure. More details and measures of interest for this case study are available at [6, 5].

## 4 Conclusions

We have discussed techniques for modelling and evaluating energy saving reliable CPS, where the management of the energy resources is handled by specific computational units that control physical entities. Policies of energy saving are generally adopted for implementing energy saving techniques. While applying these policies, computational units must also guarantee the overall reliability of the system and other dependability related requirements. Energy saving and reliability aspects are often opposite requirements, and it is important to find a good trade-off between these measures.

Hence, we have proposed a model-based approach to the analysis of CPS, which is very useful for evaluating the measures of interest in order to find a good set-up for the parameters of the policies of energy saving, so to satisfy both energy and reliability requirements. By adopting a model-based approach it is possible to evidence, far before they become obvious, weakness points in the system under analysis, saving both in time and economic terms.

We have presented widely adopted modelling formalisms for the evaluation of energy-saving aspects, focusing on Markovian models, non-Markovian models and extensions of Petri Nets capable of expressing probabilistic aspects of the systems under analysis. A list of available tools for supporting the design and evaluation of models of energy saving systems has been presented, with pointers to the related literature.

To illustrate and corroborate our proposal, we have illustrated a case study taken from the railway transport domain. In particular, we have tailored our modelling framework to the evaluation of reliability and energy consumption indicators for a system of rail road switch heaters, by using stochastic activity networks and the Möbius tool. The adopted policy of energy saving is based on an on-off mechanisms for the heating of the switches, with parametric thresholds representing the temperatures triggering the activation/deactivation of the heating.

We address some lines of future research concerning energy saving CPS. Firstly, due to their physic nature and their interactions with the environment, the behaviour of these systems is in general unpredictable. New control techniques are necessary, for restricting their possible behaviours in order to predict and avoid possible failures, improve and verify their dependability [1]. Moreover, defining the energy requirements that a CPS must satisfy, verifying that the proposed model satisfy them or proving that such requirements are not satisfiable is an hot research topic [4].

Likewise, energy saving solutions for CPS are expected to be impacted by these foreseen evolutions. In turn, analysis support to help defining and selecting appropriate energy saving policies which satisfy required dependability properties needs to keep the pace with this future prospect. Actually, the framework outlined in this paper is amenable to undergo extensions and adaptation. Of course, the implementation aspects will be highly impacted and may raise significant challenges, especially in terms of scalability and efficiency. However, research advancements in model/based solution techniques in relation with modern and future complex systems are expected to alleviate such potential weaknesses of model/based analysis.

Concerning the proposed case study, several extensions have been identified. It would be interesting to study how the energy consumption is modified by changing parameters of the underlying physical model. Indeed, the obtained results may suggest that, by changing the material the heaters are composed of, its length or the power consumed, a better trade-off between reliability and energy optimization can be obtained. It would also be interesting to let the power consumed by the system vary at different weather conditions. This may help to improve the reliability of the system. Indeed in case of emergency a major throughput may prevent a failure. Adapting the thresholds to the different class of priorities may increase the reliability of the system. Finally, modelling and evaluating the case study with different formalisms and tools would help the validation of the obtained results.

## References

1. Antsaklis, P.: Goals and challenges in cyber-physical systems research editorial of the editor in chief. IEEE Transactions on Automatic Control 59(12), 3117–3119 (Dec 2014)
2. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. Dependable and Secure Computing, IEEE Transactions on 1(1), 11–33 (Jan 2004)
3. Balbo, G.: Introduction to generalized stochastic petri nets. In: Bernardo, M., Hillston, J. (eds.) Formal Methods for Performance Evaluation, LNCS, vol. 4486, pp. 83–131. Springer Berlin

Heidelberg (2007)
4. Banerjee, A., Venkatasubramanian, K.K., Mukherjee, T., Gupta, S.K.S.: Ensuring safety, security, and sustainability of mission-critical cyber-physical systems. Proceedings of the IEEE 100(1), 283–299 (Jan 2012)
5. Basile, D., Chiaradonna, S., Di Giandomenico, F., Gnesi, S.: Stochastic model-based evaluation of reliable energy-saving rail road switch heating systems. Tech. rep., Istituto di Scienze e Tecnologie dell'Informazione, Consiglio Nazionale delle Ricerche, Pisa (2016), http://puma.isti.cnr.it/dfdownload.php?ident=/cnr.isti/2016-TR-009
6. Basile, D., Chiaradonna, S., Di Giandomenico, F., Gnesi, S., Mazzanti, F.: Stochastic model-based analysis of energy consumption in a rail road switch heating system. In: Software Engineering for Resilient Systems - 7th International Workshop, SERENE 2015, Paris, France, September 7-8, 2015. Proceedings. pp. 82–98 (2015)
7. Bause, F., Kritzinger, P.S.: Stochastic petri nets: An introduction to the theory. SIGMETRICS Perform. Eval. Rev. 26(2), 2–3 (Aug 1998)
8. Bernardi, S., Merseguer, J., Petriu, D.C.: Model-Driven Dependability Assessment of Software Systems. Springer (2013)
9. Boguñá, M., Lafuerza, L.F., Toral, R., Serrano, M.A.: Simulating non-markovian stochastic processes. Phys. Rev. E 90, 042108 (Oct 2014)
10. Brodowski, D., Komosa, K.: A railroad switch and a method of melting snow and ice in railroad switches (2013), https://data.epo.org/publication-server/rest/v1.0/publication-dates/20131225/patents/EP2677079NWA1/document.html
11. Bucci, G., Carnevali, L., Ridi, L., Vicario, E.: Oris: a tool for modeling, verification and evaluation of real-time systems. STTT 12(5), 391–403 (2010), http://dx.doi.org/10.1007/s10009-010-0156-8
12. Cannon, J.R.: The One-Dimensional Heat Equation. Cambridge University Press (1984), cambridge Books Online
13. Chiola, G., Franceschinis, G., Gaeta, R., Ribaudo, M.: Greatspn 1.7: Graphical editor and analyzer for timed and stochastic petri nets 24, 47–68 (1995)
14. Ciardo, G., Muppala, J., Trivedi, K.: Spnp: Stochastic petri net package (1989)
15. Clark, G., Courtney, T., Daly, D., Deavours, D., Derisavi, S., Doyle, J.M., Sanders, W.H., Webster, P.: The möbius modeling tool. In: Proceedings of the 9th International Workshop on Petri Nets and Performance Models. pp. 241–250 (2001)
16. David, R., Alla, H.: On hybrid petri nets. Discrete Event Dynamic Systems 11(1-2), 9–40 (Jan 2001)
17. Diab, Hassan B.; Zomaya, A.Y.: Dependable computing systems: paradigms, performance issues and applications. John Wiley & Sons (2005)
18. Erbes, T., Shukla, S.K., Kachroo, P.: Stochastic learning feedback hybrid automata for dynamic power management in embedded systems. In: SMCia/05, IEEE Mid-Summer Workshop on Soft Computing in Industrial Applications, June 2005 (2005)
19. Friedler, F.: Process integration, modelling and optimisation for energy saving and pollution reduction. Applied Thermal Engineering 30(16), 2270 – 2280 (2010), http://www.sciencedirect.com/science/article/pii/S1359431110001936, selected Papers from the 12th Conference on Process Integration, Modelling and Optimisation for Energy Saving and Pollution Reduction
20. German, R.: Performance Analysis of Communication Systems with Non-Markovian Stochastic Petri Nets. John Wiley & Sons, Inc., New York, NY, USA (2000)
21. Ghasemieh, H., Boudewijn, R. Haverkort, M.R.J., Remke, A.: Energy resilience modeling for smart houses. In: DSN (2015), to appear
22. Haverkort, B.R.: Lectures on formal methods and performance analysis. chap. Markovian Models for Performance and Dependability Evaluation, pp. 38–83. Springer-Verlag New York, Inc., New York, NY, USA (2002), http://dl.acm.org/citation.cfm?id=567305.567307
23. Lee, E.A.: Cyber physical systems: Design challenges. In: Proceedings of the 2008 11th IEEE Symposium on Object Oriented Real-Time Distributed Computing. pp. 363–369. ISORC

'08, IEEE Computer Society, Washington, DC, USA (2008), http://dx.doi.org/10.1109/ISORC.2008.25
24. Lyu, M.R. (ed.): Handbook of Software Reliability Engineering. McGraw-Hill, Inc., Hightstown, NJ, USA (1996)
25. Misra, S., Krishna, P.V., Saritha, V., Obaidat, M.S.: Learning automata as a utility for power management in smart grids. IEEE Communications Magazine 51(1), 98–104 (2013)
26. Müller, S.C., Häger, U., Rehtanz, C., Wedde, H.F.: Application of self-organizing systems in power systems control. In: Dieste, O., Jedlitschka, A., Juzgado, N.J. (eds.) PROFES 2012 Proceedings. LNCS, vol. 7343, pp. 320–334. Springer (2012)
27. O'Connor, P.P., Kleyner, A.: Practical Reliability Engineering. Wiley Publishing, 5th edn. (2012)
28. Peterson, J.L.: Petri nets. ACM Comput. Surv. 9(3), 223–252 (Sep 1977), http://doi.acm.org/10.1145/356698.356702
29. Qiu, Q., Wu, Q., Pedram, M.: Stochastic modeling of a power-managed system: construction and optimization. In: Proceedings of the 1999 International Symposium on Low Power Electronics and Design, 1999, San Diego, California, USA, August 16-17, 1999. pp. 194–199 (1999)
30. Qiu, Q., Wu, Q., Pedram, M.: Dynamic power management of complex systems using generalized stochastic petri nets. In: DAC. pp. 352–356 (2000)
31. http://www.railsco.com/\~electric_switch_heater_controls.htm
32. Reibman, A., Smith, R., Trivedi, K.: Markov and markov reward model transient analysis: An overview of numerical approaches. European Journal of Operational Research 40(2), 257 – 267 (1989), http://www.sciencedirect.com/science/article/pii/0377221789903354
33. Sahner, R.A., Trivedi, K., Puliafito, A.: Performance and Reliability Analysis of Computer Systems: An Example-Based Approach Using the SHARPE Software Package. Springer Publishing Company, Incorporated (2012)
34. Sanders, W.H., Meyer, J.F.: Stochastic activity networks: Formal definitions and concepts. In: Brinksma, E., Hermanns, H., Katoen, J. (eds.) Lectures on Formal Methods and Performance Analysis, First EEF/Euro Summer School on Trends in Computer Science 2000, Revised Lectures. LNCS, vol. 2090, pp. 315–343. Springer (2000)
35. Front matter. In: Karlin, H.M.T. (ed.) An Introduction to Stochastic Modeling (Revised Edition), pp. iii –. Academic Press, revised edition edn. (1994), http://www.sciencedirect.com/science/article/pii/B978012684885450001X
36. Front matter. In: Grabski, F. (ed.) Semi-Markov Processes: Applications in System Reliability and Maintenance, pp. i – ii. Elsevier (2015), http://www.sciencedirect.com/science/article/pii/B9780128005187099889
37. Trivedi, K.S.: Probability & statistics with reliability, queuing and computer science applications. John Wiley & Sons (2008)
38. Trivedi, K.S., Malhotra, M.: Messung, Modellierung und Bewertung von Rechen- und Kommunikationssystemen, chap. Reliability and Performability Techniques and Tools: A Survey, pp. 27–48. Springer Berlin Heidelberg, Berlin, Heidelberg (1993)
39. Čaušević, A., Seceleanu, C., Pettersson, P.: Distributed energy management case study: A formal approach to analyzing utility functions. In: Margaria, T., Steffen, B. (eds.) Leveraging Applications of Formal Methods, Verification and Validation. Specialized Techniques and Applications, Lecture Notes in Computer Science, vol. 8803, pp. 74–87. Springer Berlin Heidelberg (2014)
40. Zhu, D., Melhem, R., Mossè, D.: The effects of energy management on reliability in real-time embedded systems. In: Computer Aided Design, 2004. ICCAD-2004. IEEE/ACM International Conference on. pp. 35–40 (Nov 2004)