A Stochastic Model-Based Approach to Analyze Reliable Energy-Saving Rail Road Switch Heating Systems

Davide Basile, Silvano Chiaradonna, Felicita Di Giandomenico, and Stefania Gnesi

Istituto di Scienza e Tecnologia dell'Informazione "A. Faedo", Consiglio Nazionale delle Ricerche, ISTI-CNR, Pisa, Italy

Abstract. Rail road switch heaters are used to avoid the formation of snow and ice on top of rail road switches during the cold season, in order to guarantee their correct functioning.

Effective management of the energy consumption of these devices is important to reduce the costs and minimise the environmental impact. While doing so, it is critical to guarantee the reliability of the system.

In this work we analyse reliability and energy consumption indicators for a system of (remotely controlled) rail road switch heaters by developing and solving a stochastic model-based approach based on the Stochastic Activity Networks (SAN) formalism. An on-off policy is considered for heating the switches, with parametric thresholds of activation/deactivation of the heaters and considering different classes of priority.

A case study has been developed inspired by a real rail road station, to practically demonstrate the application of the proposed approach to understand the impact of different thresholds and priorities on the indicators under analysis (probability of failure and energy consumption).

1 Introduction

Nowadays, there is a great attention towards cautious usage of energy sources to be employed in disparate application domains, including the transportation sector, to save both in economical terms and in environmental impact [16, 33, 24, 31]. Therefore, studies devoted to analyse and predict energy consumption are more and more gaining importance, especially in combination with other non functional properties, such as reliability, safety and availability. The (electric) energy consumption of a system is the amount of power consumed by a 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 a system to avoid service failures that are more frequent or more severe than is acceptable [2]. While several works devoted at analysing reliability of systems or their energy consumption are available in literature [16] [25] [13] [31], there is relatively less work concerning the interplay of these two measures. In this paper, we address reliability and energy consumption through stochastic model-based analysis considering a realistic case study, a *rail road switch heating system* [26, 8].

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 and the mechanisms which allow a train to be directed can be blocked by an excessive amount of snow or ice. To overcome this issue, the rail road switches need to be cleaned from possible snow or ice forming on top of it. 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 [8]. A failure in those switches can lead to major malfunctions.

In case electricity is used to heat the switches it is important to optimise the energy consumption and at the same time to ensure reliability. Different policies may be adopted as for example to heat a selection of switches for a given amount of time. A different approach would be to heat all the switches together. Note that, by turning on all the heating systems at the same time, an overhead of energy consumption can lead to a blackout. Alternatively, an excessively parsimonious policy to save on energy can cause the failure of some rail road switch heaters [4].

In this paper, we adopt a stochastic model-based approach to analyse the behaviour of the rail road switch heating system under different circumstances. 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. In particular, we exploit *Stochastic Activity Network* [27] to model the heating system, and use the Möbius tool [11] to perform evaluations.

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. A prioritized approach has been considered, where the heaters are categorized according to their importance inside the analysed railway station. Note that a failure of the heating system is accounted for by other components of the railway system, namely interlocking mechanisms which guarantee safety; however we do not include them in our analysis. This paper extends a previous work [4] in different directions: we provide a major update of the model by introducing a first in, first out (FIFO) prioritized approach, the analysis results are improved with new measure of interests and we consider real world data for our evaluations. Finally, we show that the reliability of the system augments when a prioritized approach is considered.

Structure of the paper Stochastic model-based analysis is introduced in Section 2. In Section 3 we describe the considered rail road switch heating system and its behaviour. We present the models of the rail road switch heating system in Section 4, which are then instantiated to a real scenario in Section 5. The results of our experiments are discussed in Section 6, while related work and conclusions are in Sections 7 and Section 8.

2 Stochastic Model-based Analysis

Given the stochastic phenomena involved in our analysis, namely the failure occurrence and weather forecast, we have adopted a stochastic model-based approach [7, 14]. Stochastic model-based methods 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 analysed system must satisfy, which can be non-functional properties such as dependability, performability and others. Then a model of the system under analysis is built, that represents its behaviour, in order to guarantee that the requirements are met. It is possible to compare different alternatives for the same system, and select the one that better suits the requirements. An early modelling phase is also useful to highlight problems in the design of the system.

When the design phase is completed, a model allows for predicting the overall behaviour of the system, fostering an analysis for the fulfilment of constraints in the design phase and the acceptance cases. For an already existing system, an a-posteriori analysis of properties such as dependability or performance is useful in order 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 [29].

2.1 Stochastic Activity Network

Several stochastic modelling methods have been proposed in literature [3, 27, 1, 5, 12].

According to the underlying stochastic process, they are classified into *Markovian* and *non-Markovian* [18]. Markovian models are those satisfying the *Markov property* (memoryless), that is the conditional probability distribution of future states (conditional on both past and present values) depends only upon the present state. Non-Markovian models do not satisfy the Markov property and thus allow for different probability distributions. Since our 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) [27].

SAN is a formalism widely used for performance, dependability and performability evaluation of complex systems, given its high expressiveness and the powerful tools for modelling and evaluating them [11]. The SAN formalism is a variant of Stochastic Petri Nets [5], 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:

FIFO	First In, First Out
SAN	Stochastic Activity Networks
T_{wa}	Warning Threshold
T_{wo}	Working Threshold
NH_{max}	Maximum Amount of Deliverable Energy
$CE(t,l)$	Energy Consumed in the interval (t,l)
$AHAT(t)$	Activated Heaters At Time t
$PFAIL(t,l)$	Probability of Failure in the interval (t,l)

Table 1: Table of frequently used symbols and acronyms.

instantaneous and *timed*. Instantaneous activities complete once the enabling conditions are satisfied. Timed activities take an amount of time to complete following a 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. SAN models are defined and solved by using the multi-formalism multi-solvers tool Möbius [11]. Möbius is a tool, that supports various formalisms such as SAN, PEPA, 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, where atomic models are building blocks which can be composed with proper operators *Rep* and *Join* to generate an overall composed model.

3 The System Under Analysis

We briefly describe the real-world devices that we want to model and the underlying logical system we have built for analysis purposes. In Table 1 the frequently used acronyms and symbols are displayed.

3.1 Description of the network of rail road switch heaters

In Figure 1, a rail road switch heater is displayed. The picture is taken from [26].

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 [26].

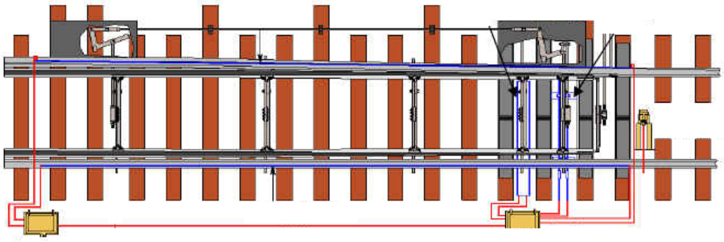

Fig. 1: An example of a rail road switch heater heated by electric induction ©2006 - 2012 Rails Company.

The management of the heaters is automatic, and is remotely controlled by a central computational unit using a powerline [6]. Powerline communications have the possibility to transmit coded information through the existing electric lines. The central unit manages the maximum amount of power that can be delivered to the system, in order to prevent possible blackouts.

Rail road switch heaters can be located in different zones (according to the topology of the network of switches) with different weather conditions. At a higher altitude, the probability to have temperatures lower than the average is greater than the case of rail road switch heaters located at lower altitudes. Moreover, different heaters could have different levels of exposure to the sun light, hence the temperature can be slightly different.

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.

3.2 Logical structure of the system

We identify two main logical components that are the *heater* and the *central coordinator*. The network of heaters is realised by replicating the heater component, and their activation/deactivation is controlled by the central coordinator. We now discuss the 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}. A smaller gap between these thresholds will result in a frequent activation of the heating system, but for a shorter period of time. Alternatively, a wider gap between the thresholds will result in a less frequent activation, but it will be for longer periods of time. Additionally, the time during which a single heater is active depends on the topology of the network of switches and the weather conditions. Heaters located in a colder geographical area will consume more energy than those under warmer temperatures.
- *Coordinator*: the coordinator will collect the requests of activation from the pending heaters, and it will manage 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 will assign priorities to switches based on their criticality on the track and we will exploit the assigned priority in the performed analysis, so 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. Below we describe in details the behaviour of the central coordinator.

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 the 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.

In the next section the SAN based models representing the network of switch heaters and the coordinator are described.

4 Modelling Framework

In this section we describe the SAN based model realising the rail road switch heating system logic and protocol described previously.

4.1 The composed model

The overall model is obtained by the composition of the atomic models, using the join and rep operators, as shown in Figure 2 (where the atomic models are the leaves of the tree while the overall composed model is the root).

The atomic model $Coordinator$ represents the central coordinator, described in Section 4.5. The atomic SAN models $LocalitySelector$, $ProfileSelector$ and $SwitchIDSelector$ represent, respectively, the selector for the weather profile, the location of the switch and the unique identifier of each switch. 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 *HeaterModuleM* using the Rep operator, 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 [11] for joining different SAN models thus allowing modularity.

4.2 Parameters of the model

The main parameters of the SAN composed model are the two temperature thresholds T_{wa},T_{wo} and NH_{max} that we recall to be the maximum power that the system can provide every instant of time, expressed in percentage of heaters that can be turned on at the same time. Reliability aspects are modelled by taking into account the possibility of experiencing a failure in a switch. Other aspects that need to be accounted for in the modelling framework are the weather forecast and the topology of the network of switches.

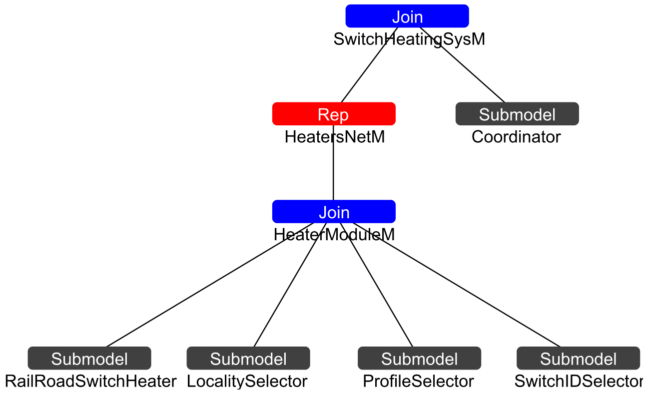

Fig. 2: The composed model.

Weather forecast. To model the external weather conditions, our model takes in input a table containing profiles of average temperatures in those days for which the analysis is relevant (a.g. winter days).

The time window under analysis is divided in intervals to which an average reference temperature is assigned. The current instance of the model concentrates on a whole day, divided in intervals of two hours. However, the model can be easily modified to consider longer (or shorter) periods, as well as different number of intervals.

A probability is assigned to each weather profile, based on weather statistics. At starting time, a SAN model will select one of these profiles according to its probability.

Topology. To model the temperature of the heaters taking into account their displacement, a gap from the average temperature is considered, in order to perturb the table of average temperatures per day. In case we are considering networks of heaters where there are more rail road switch heaters located at higher altitudes or in the shade, it will be more probable that their actual temperature will be lower than the average.

4.3 Selectors

We now describe the SAN models adopted for selecting the appropriate identifier, weather profile and localization of the corresponding heater.

Weather Profile Selector. The SAN model shown on Figure 3a defines the weather profile associated to the corresponding heater, represented by the number of tokens in the place *profileID* (shared with the SAN model *RailRoadSwitchHeater*). Indeed, each weather profile has a unique identifier and if, for example, *profileID* has two tokens, then the second weather profile is assigned to this model.

Each case of the instantaneous activity *selectProfile* represents the probability to select a different weather profile at completion of the activity.

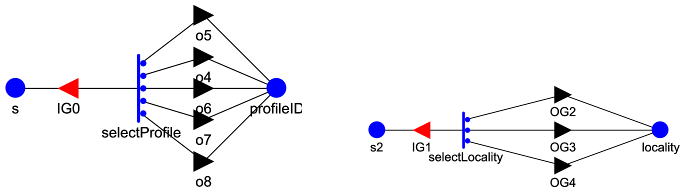

(a) ProfileSelector (b) LocalitySelector

Fig. 3: The SAN models ProfileSelector and LocalitySelector.

Locality Selector. The SAN model *LocalitySelector*, shown in Figure 3b, defines the temperature of the corresponding heater according to its locality, represented by the number of tokens in the place *locality*; codified similarly to the weather profile. At each case of the instantaneous activity *selectLocality* is associated the probability to select a different location at completion of the activity.

Switch ID Selector. In Möbius tool [11] different replicas of SAN models are anonymous, and it is not possible to distinguish between them. However, we need to identify each single heater to assign its parameters. Hence a SAN model is used for assigning a unique ID to each heater module.

The SAN model *SwitchIDSelector* is depicted in Figure 4. It assigns a unique ID, represented by the number of tokens in the place *SwitchID*, to each replica of the submodel *HeaterModuleM*. The place *SwitchID* is shared with the SAN model *RailRoadSwitchHeater*. The uniqueness of the ID is guaranteed by the instantaneous activity $IA1$ that assigns to each replica a different ID according to the place $p1$ which is shared among the different replicas of the heater module and it is used as a counter.

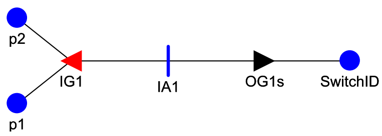

Fig. 4: The SAN model SwitchIDSelector.

4.4 Rail Road Switch Heater

In Figure 5 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*.

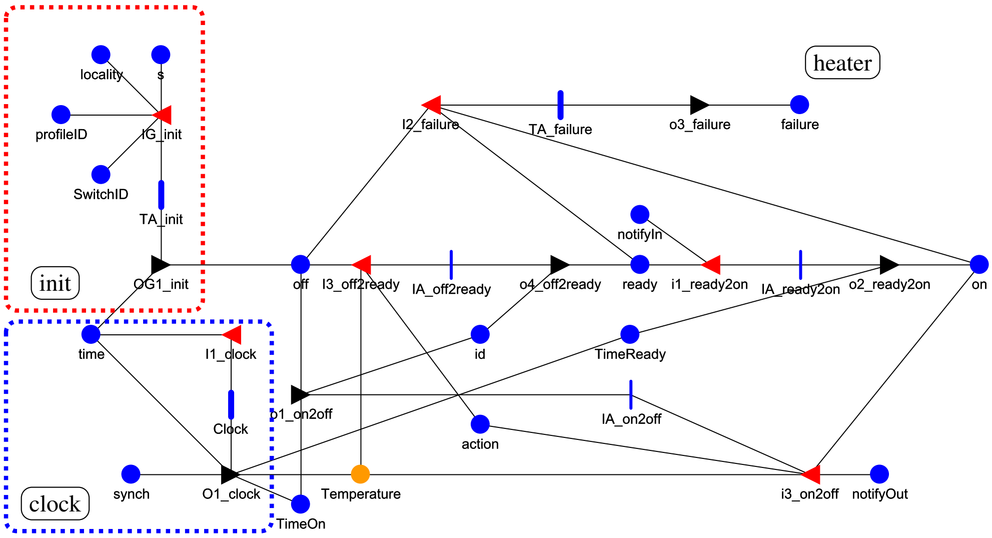

Fig. 5: The SAN model RailRoadSwitchHeater, logically divided into three sub-nets: the init sub-net, the clock sub-net and the heater sub-net.

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 (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 update the environment temperature and the temperature of the rail road track. In this paper we have considered as unit of time one hour. The activity *clock* completes each hour, indeed it has a deterministic distribution of time (non Markovian). It 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.

Heater sub-net. The heater sub-net represents the status of the rail road switch heater. As discussed in Section 3 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.

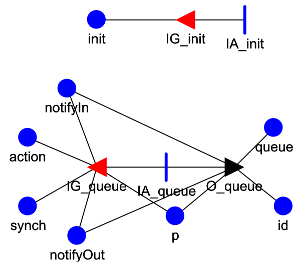

Fig. 6: The SAN model Coordinator.

The heater sub-net interacts with the Coordinator SAN model discussed in Section 4.5, through places shared among all the replicas of the heater model and the Coordinator model, implementing the protocol described in Section 3. 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).

4.5 Coordinator model

The SAN model Coordinator, depicted in Figure 6, represents the central management unit and it interacts with all the heaters in the network by activating, deactivating or moving them in a waiting state. The places $notifyIn, notifyOut, action, id$ and *synch* are shared with the heaters and represent the prioritized FIFO queue.

At completion of an instantaneous activity, different actions are performed depending on the required action:

- *insert*: if there is enough energy available, the requesting heater is switched on by assigning to the marking of the place $notifyIn$ the marking of the place id, which corresponds to the place *SwitchID* of the selected switch. Otherwise, if no energy is available but the pending heater has a higher priority than one of the activated heaters, then the marking of the place $notifyOut$ is set to the id of the heater to be turned off and the marking of the place $notifyIn$ is set to the marking of the place id. If no energy is available and the priority is not higher than those activated switches, the request will be ignored.
- *remove*: the heater h_i is turned off. If there are no other heaters in state ready, one token on the place *queue* will be removed. Otherwise, if another heater h_j in the network is in state ready, then the coordinator will notify to h_j its activation by setting the marking of the place $notifyIn$ to the id of h_j. Note that, among the

heaters in state ready, h_j is the one which has been waiting the activation for longer time (FIFO order).

4.6 Physical Model for Heat Exchange

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, each interval of time a new value will be picked up according to the table for the selected weather profile and the locality table. To model the internal temperature of the device and to estimate the energy consumption of the heater, we model the exchange of heat through convection. 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 an iron bar representing the rail road track. 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 [8], which also contain the power consumed by a single heater and data about 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 [9]:

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

where u is the coefficient of convective exchange, calculated as:

$$u = \left(\frac{g \cdot \beta \cdot (T_i - T_e) \cdot \rho^2 \cdot AV^3}{\mu^2}\right)^{\frac{1}{4}} \cdot \frac{0.54 + 0.26}{2} \cdot \frac{K}{AV},$$

The other parameters are: c, the heat capacity of iron; A, the surface area exposed to the external temperature; AV, the ratio between area and volume of the iron bar; 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. Moreover g, the gravity acceleration; β, the thermal expansion coefficient; ρ, the density of air; μ, the dynamic viscosity and K the thermal conductivity of iron.

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.

The function representing the heating exchange is defined in C++, and it is called by the output gate $O1_{clock}$ shown in Figure 5 to update the temperature of the rail road each interval of time t.

In the next section, we instantiate the composed model with varying values of T_{wa},T_{wo} and NH_{max}, to show their impact and benefits on both energy consumption and system reliability.

5 Scenarios and Settings

In this section we describe the scenarios and settings we have considered for our analysis. In order to corroborate the proposed methodology, real world data are used. In particular, to keep the presentation simple, an average medium-size railway station of a northern city has been chosen. The cold winter days are considered for the temperatures and the parameters of the physical model are based on data available from real devices [8] [26].

5.1 The railway station

We have considered a generic medium-size railway station, based on data about real world railway stations [10]. In our analysis we considered 41 switches that are divided into three different classes of priority depending on the criticality of the service offered by the switches:

- *low priority* (12 heaters): the switches that have access to the siding tracks. A failure of these switches does not affect immediately the traffic inside the station, as the trains can transit using the main tracks while the access to the siding tracks may be changed;
- *medium priority* (6 heaters): the switches that are placed on the bay tracks. Indeed, a failure in these switches can be considered less harming than one on the main tracks;
- *high priority* (23 heaters): the switches placed on the main tracks of the station: a failure in one of those switches may block the transit of trains inside the station, causing serious delays and loss of money for the affected companies.

5.2 Weather data

In the analyses reported in this paper, the environment temperatures are based on those of mid-January 2015 in the city of Montreal [32] (displayed in Figure 7), which were the coldest of that winter. In the proposed model it is possible to consider several days, by iterating the loop which reads the temperature at different intervals of the day, and by expanding the table of the weather profiles.

The data structure representing the external temperatures for the profile 1, corresponding to the values of Figure 7, is shown in Table 2. We assume that the data of each different profile are selected based on uniform distribution.

Table 3 shows the gap from the average temperatures and the probability of occurrence of it for three different localities. The data we have used are representative, to illustrate the possibility of having a fine grained model of the temperatures of rail road switch heaters. It must be pointed out that these parameters can be easily modified to deal with different conditions.

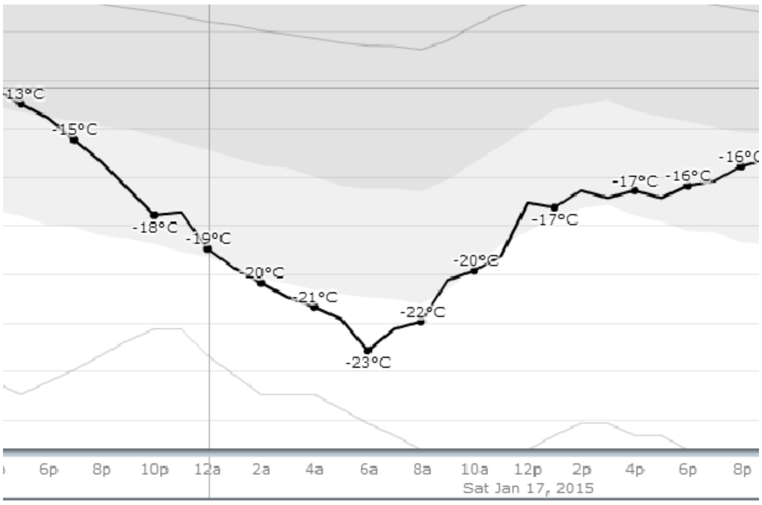

Fig. 7: The temperatures for the 16th and 17th of January 2015 in Montreal, retrieved from [32].

	06 pm	**08 pm**	**10 pm**	**12 pm**	**02 am**	**04 am**	**06 am**	**08 am**	**10 am**	**12 am**	**02 pm**	**04 pm**
profile 1	-14 °C	-15 °C	-18 °C	-19 °C	-20 °C	-21 °C	-23 °C	-22 °C	-20 °C	-17 °C	-17 °C	-15 °C

Table 2: The row of the weather profiles corresponding to Figure 7.

5.3 Parameters of the physical model

The parameters and the default values considered for calculating the heat exchange between the rail road, the heaters and the external air are shown in Table 4. These data were retrieved from different real devices and patents [8] [26].

The parameters l, h and w are respectively the length, height and width of the portion of iron representing the rail road track that needs to be heated. Standard values are used for the parameters of the air and of the iron, described in Section 4. We assume that all the rail road switch heaters use the same amount of power, which in this case is 330 Watts per meter [8].

	Gap from Average Temperature	Probability
locality 1	-0.8 °C	0.3
locality 2	-0.3 °C	0.3
locality 3	0 °C	0.4

Table 3: The gap from the average temperature and the probability of occurrence of it for the three considered localities.

l	h	w	T_i	β	ρ
6 m	28 cm	4 cm	5 °	0.0037 C^{-1}	1.30 $\frac{Kg}{m^3}$

μ	K	AV	A	m	c
17.157 $\cdot 10^{-6} Pa \cdot sec$	0.023 $\frac{W}{mK}$	0.0175 m	1.68 m^2	529.13 Kg	450 $\frac{J}{KgK}$

Table 4: the parameters for the model of heat exchange by convection.

Finally, the considered freezing threshold is set up to zero degrees. If the temperature of the track goes below this value the system has an increasing probability to experience a failure (due to formation of snow and ice on top of it).

Remark 1. The parametric nature of our model allows its customization to a wide variety of system configurations in terms of both size and characterization of system components, as well as different weather profiles and rail road track layouts. Indeed, the priorities assigned to the switches can be easily changed. Moreover, this system can be adopted in a network of remotely controlled switches, where temperatures in different zones can be more different than what we have considered.

6 Analysis results

We now describe the results of the evaluations we have performed in order to find suitable trade-offs in terms of reliability and energy consumption for different settings of the parameters of the model.

In our experiments, we consider different ranges of thresholds and of available power represented by NH_{max}. For T_{wa} the range of temperatures goes from 4 °C to 25 °C, the gap between T_{wa} and T_{wo} goes from 1 °C to 8 °C. For the parameter NH_{max}, we have considered four scenarios where the availability is sufficient to supply 25%, 50%, 75% and 100% of all the considered switches, respectively.

In all the considered evaluations, we assume that at starting time the system is in a safe condition, that is the internal temperature of all switches is equal to its working temperature. This assumption is useful for avoiding instantaneous failure. We have considered all the combinations of those parameters (thresholds and NH_{max}) and here we report the most significant results. Moreover we have analysed the trend of heaters that are active during the hours of the day. Finally, we show how a prioritized heating approach augments the reliability of the station, by ensuring a lower probability of failure for the portion of tracks that are vital for the correct functioning of the railway station.

6.1 Measures of Interest

We consider three different measures of interest. The first two concern the energy consumption while the third addresses the reliability of the system under analysis.

1 $CE(t,l)$: the time (in 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 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 $AHAT(t)$: the total number of heaters that are active at the instant of time t. This measure can be useful for planning purposes. It is evaluated by accumulating the marking of the place *on* for all the replicas of the rail road switch heater model in the instant of time considered.

3 $PFAIL(t,l)$: 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 [30], that is $1-PFAIL(t,l)$.

6.2 Results Discussion

We have plotted the results of our experiments in different graphs, at the variation of different parameters, that are T_{wa}, the gap between T_{wa} and T_{wo}, NH_{max} and the prioritization.

Figure 8 concerns $PFAIL(t,l)$ for different values of NH_{max}. Moreover Figures 8a, 8c and 8e consider a network with no priorities, while in Figures 8b, 8d and 8f the priorities are considered. Finally, Figures 8a and 8b consider a gap between thresholds of 1°C, Figures 8c and 8d consider a gap between thresholds of 4°C, Figures 8e and 8f consider a gap between thresholds of 7°C. Figure 9 is similar to Figure 8, but $CE(t,l)$ is considered.

In Figure 10 the measure of interest $PFAIL(t,l)$ is considered. The different curves in each graph represent different gaps between T_{wa} and T_{wo}. Figures 10a and 10c consider a network with priorities, while in Figures 10b and 10d no priorities are considered. Moreover, in Figure 10a and 10b we have $NH_{max}=50\%$, while in Figure 10c and 10d we have $NH_{max}=75\%$. Figure 11 shows different values of $CE(t,l)$, where the graphs are organized as in Figure 10.

Finally, we only have considered the case of switches with higher priority, since they are the majority of heaters in the network (23 over 41), and the benefits of using an approach with priority have more impact on them.

Figure 12 shows the values of $AHAT(t)$, where on the x axis different hours of the day are reported while on the y axis the resulting values of $AHAT(t)$. The amount of energy available to the system is fixed to 50% of active heaters per time. Figure 12a and Figure 12b show a network with priority, where different bars correspond to different priorities. In the two graphs different thresholds are considered, to show how they impact the behaviour of the system. The last graph (Figure 12c) instead shows the case where no prioritization is performed, and the considered heaters are only 23 heaters over 41, again to compare them with the ones with higher priority.

We now discuss the major evidences resulting from the output of our experiments.

Supplied energy. As expected, a greater amount of energy delivered corresponds to a lower probability of failure, as shown in Figure 8, and the number of enabled heaters increases as a result of having more energy. We remark that the maximum amount of delivered energy ($NH_{max} = 100\%$) corresponds to the energy sufficient for heating all the switches in the network. Overproduction of energy is not considered. In particular, for $NH_{max} = 25\%$ there are no acceptable values for $PFAIL(t,l)$ (too high probability of failure), while the converse holds for $NH_{max} = 100\%$. For $NH_{max} = 50\%$ and $NH_{max} = 75\%$ the values of the thresholds are crucial for augmenting the resilience of the system.

Concerning the consumption of energy $CE(t,l)$, in Figure 9 we observe that for lower values of T_{wa}, the energy consumed with different values of NH_{max} (50%,75% and 100%) is similar, while for $NH_{max} = 25\%$ there are more or less 5 hours (of active heating) of difference. By increasing T_{wa}, in the case of no priority, the gap between the four considered curves for NH_{max} increases, while this is not the case if priorities are considered. This trend is related to the probability of failure. Indeed if $PFAIL(t,l)$ increases then $CE(t,l)$ decreases, because when a switch fails it will no longer consume energy. Note that only 23 heaters out of 41 are considered (high priority switches), more or less 50% of heaters in the network. This explains why, by considering values of $NH_{max} > 50\%$ and priorities, the energy delivered is not consumed by high priority switches and their corresponding values of the measures of interest are equal.

Gap between thresholds. The probability of failure has a convex trend with a global minimum. This is true either when NH_{max} is fixed or when the gap between the thresholds is fixed (Figure 10 and Figure 8). In these experiments, $PFAIL(t,l)$ is minimized for a value of T_{wa} around 7 °C. Curiously, an exception is represented in Figure 10 for a gap of 5 °C.

An explanation of this phenomenon follows. 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. Indeed, since the kilowatts consumed by each heater are constant, an increment in T_{wa} will result in a longer period of activation of the heater. This does not happen when the amount of energy available is enough for heating all the devices, trivially because no heaters are stuck waiting their turn.

Concerning the gap between the thresholds, Figure 10 emphasizes how a tight gap between the thresholds is better than a wide gap. Indeed, with a tight gap there is a better distribution of energy between all the heaters that are waiting to be activated, while for a wide gap once again there will be an increment in the probability of failure of the heaters that are waiting to be activated. Hence, by fixing T_{wa} to 7 °C, an optimal value for the gap between the thresholds in these experiments is of 1°C, as shown in Figure 10. However, if we consider higher values for T_{wa}, the better gap is no longer the tighter. For example, in Figure 10, for $NH_{max} = 50\%$ and no priorities, with T_{wa} higher than $23-24$ °C, the tighter gap suddenly becomes the worst case. An optimum difference between T_{wa} and T_{wo} is around 3 °C, while for $NH_{max} = 75\%$ and no priorities 5 °C is the better gap. Indeed, when T_{wa} is higher, even a tight gap requires longer periods of activation, and it is not enough to keep the heater safe when the heating phase has terminated. This is the reason why a wider gap becomes better. We must note however that for higher values of T_{wa} the resulting values of $PFAIL(t,l)$ are unacceptable.

Finally, for the energy consumed the trend is opposite to the probability of failure. Indeed, when the probability of failure is low, the energy consumption $CE(t,l)$ augments linearly with T_{wa}, while this is not the case when the probability of failure increases. Concerning the relations between $CE(t,l)$ and the gap between the thresholds, in Figure 11 we observe that by changing the gap there is a little variation in $CE(t,l)$. For the case with no priorities, for lower value of T_{wa} a tight gap consumes less energy than a wider one. This is not always true if we increase T_{wa}. Once again, the reason is due to the variations of $PFAIL(t,l)$ described above. By fixing T_{wa} to its optimal value for $PFAIL(t,l)$ (7 °C), the tighter gap also minimise the consumption of energy $CE(t,l)$, for both $NH_{max} = 50\%$ and $NH_{max} = 75\%$. Considering the priorities, the chosen gap affects marginally the values of $CE(t,l)$, because the relative values of $PFAIL(t,l)$ are closer to zero.

Priorities. We now discuss the differences between an approach which considers priorities and one which does not. In Figure 8 we observe that by switching to a network with priorities, for $NH_{max} = 25\%$ the probability of failure decreases of more or less 0.2 when T_{wa} is low. By increasing T_{wa}, the probability of failure is similar in both cases (with or without priority). For greater values of NH_{max}, the benefits of an approach with priorities are emphasized. Indeed, the decrement in the probability of failure is more relevant. In particular, $PFAIL(t,l)$ is closer to zero for values of T_{wa} lower than $15-16$ °C, while for greater values $PFAIL(t,l)$ increases. Moreover, in Figure 8 we observe again that there is an important decrement in $PFAIL(t,l)$. Note that when $NH_{max} \geq 75\%$, *all* heaters with high priority can be activated when necessary, indeed they never fail. For the energy consumption, in Figure 9 we note that for $NH_{max} = 75\%$ and $NH_{max} = 100\%$ the values of $CE(t,l)$ are almost equal. Indeed in those cases the probability of failure is similar. Considering $NH_{max} = 50\%$ and $NH_{max} = 75\%$, concerning the case with priority the consumption of energy is similar to the case of $NH_{max} = 100\%$. This is due to the fact that for such amounts of energy available, (almost) all the heaters with high priority have instantaneous access to energy when needed, thus the consumption is similar.

The trend of the energy consumption at the variation of T_{wa} and the gap for the case with priorities is linear in Figure 11. In this case, a tighter gap is always better than a wider one. Indeed, the resulting values of $CE(t,l)$ are not perturbed by $PFAIL(t,l)$.

Heaters activated during the day. Figure 12 depicts the alternation of heaters with different priorities in their heating phase. We recall that 6 heaters have low priority and 12 have medium priority. In the approach without priorities (Figure 12c), during the coldest hours, the amount of activated heaters per hour is constantly around 10 or 12. This is because we are considering 23 heaters out of 41 and $NH_{max} = 50\%$. Indeed, proportionally almost half heaters (10 or 12 over 23) are active while the others are pending.

The situation changes when we consider priorities. Now all the 21 slots of energy are taken by the heaters with higher priority. In Figure 12a, we have chosen the optimal values for T_{wa} and the gap. With these values, a better distribution of energy between all the heaters is obtained, since heaters with medium and low priority are activated when needed.

Instead, if we chose “bad” values for the thresholds, i.e. 17 °C for T_{wa} and 8 °C of gap, in Figure 12b the heaters with higher priority become hungry of energy, leaving the other heaters inactive for too long and thus experiencing a failure. Indeed, after 12 hours (i.e. around 2 am), the heaters with medium and low priority will never be activated any more, and we can safely deduce that they have failed. This is an example of how the values of thresholds affect the reliability of the system.

As final remark on the observed results, we conclude that for the data of our analyses, it is necessary to deliver an amount of energy that must be at least sufficient to heat half of the heater in the network ($NH_{max} = 50\%$), and choose for T_{wa} a value around 7 °C and a value of T_{wo} around 8 °C. With these values, we guarantee the best compromise between consumption of energy and reliability of the system.

Experiments performance Simulation-based evaluations have been performed using Möbius tool [11], 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 using simulations. For computing the results of $PFAIL(t,l)$ and $CE(t,l)$, considering a network with priorities, 835 experiments where performed (only a portion of the overall results has been discussed in this paper) in 225 minutes, with an average of 15 seconds per experiment. Concerning the network without priority, we have calculated the results of 225 experiments in 40 minutes, with an average of 10 seconds per experiment. For computing $AHAT(t)$, 100 experiments were performed in 11 minutes, with an average of 6 seconds per experiment (the data are similar for both the studies with and without priorities). It has been used a machine with CPU Intel Core i5-4570 at 3.20 GHZ with 8 GB of RAM, running 64-bit Windows 10.

7 Related works

Although not tailored to the rail road switch heating system, there are several works in the literature that analyse and optimise the energy consumption in several application domains using formal approaches. A few of them are recalled in the following.

Services negotiation of energy and reliability requirements is the selected case study in [31], 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. The entities and their protocol of communication have been modelled with Remes Hdcl language [28] and rendered as timed automata [1], through which the absence of deadlock has been certified with the model checker Uppaal [19]. Moreover, the value of an optimal utility function (i.e. weighted sum of energy and reliability requirements) has been calculated, together with the time spent by the parties in order to find an agreement. While reliability and energy requirements are given parameters that are negotiated between the parties; in our approach we are interested in compute their optimal values at the varying of prescribed parametric policies.

In [25] Generalized Stochastic Petri Nets [3] are used to solve the dynamic power management problem for systems with complex behaviour. Dynamic power manage-

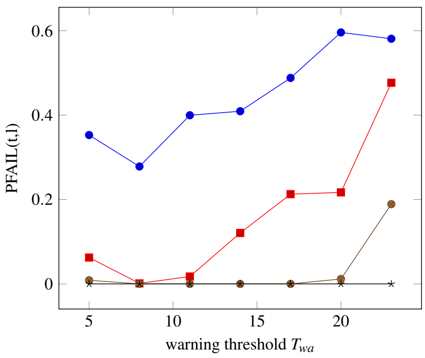

(a) No priorities, gap of 1 °C.

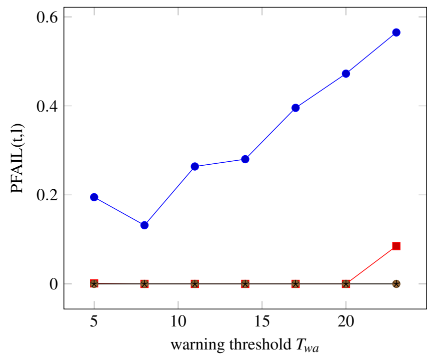

(b) With priorities, gap of 1 °C.

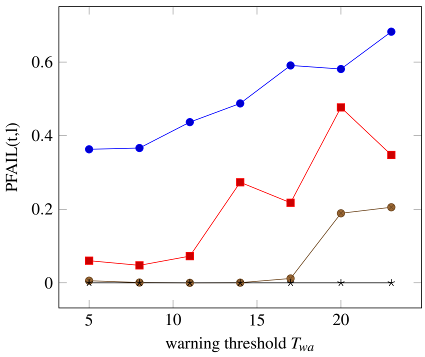

(c) No priorities, gap of 4 °C.

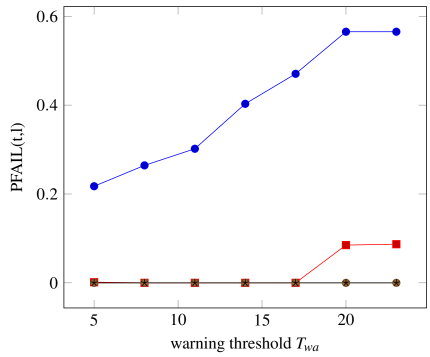

(d) With priorities, gap of 4 °C.

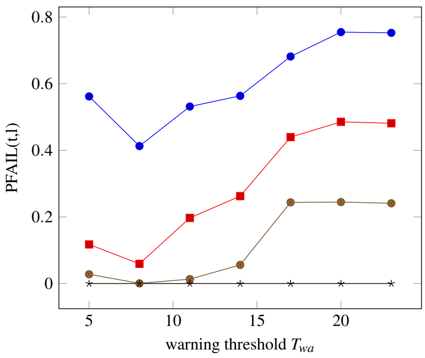

(e) No priorities, gap of 7 °C.

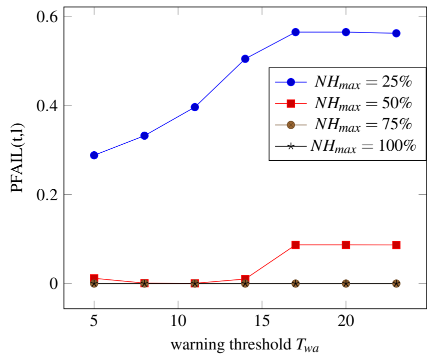

(f) With priorities, gap of 7 °C.

Fig. 8: The probability of failure of high priority switches is considered. The different lines correspond to different level of energy available to the whole system.

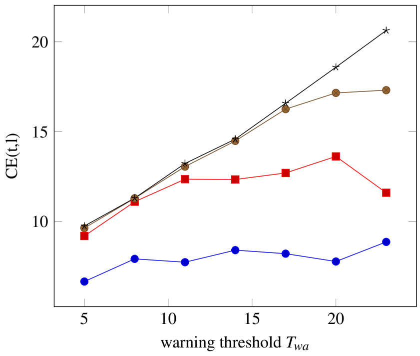

(a) No priorities, gap of 1 °C.

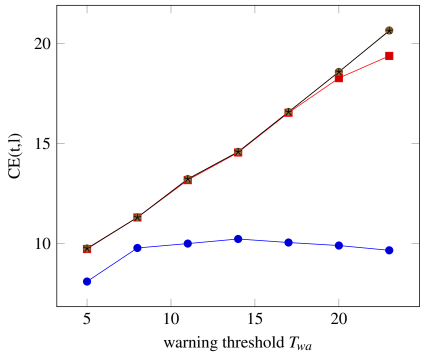

(b) With priorities, gap of 1 °C.

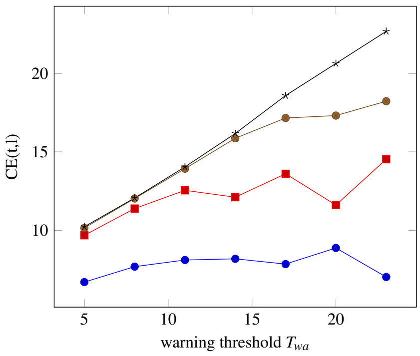

(c) No priorities, gap of 4 °C.

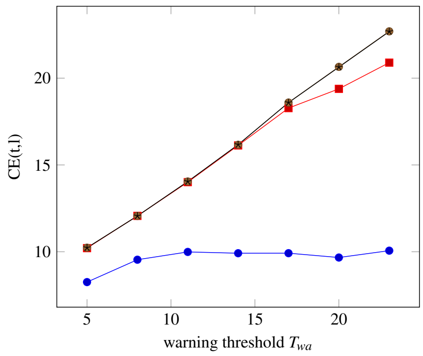

(d) With priorities, gap of 4 °C.

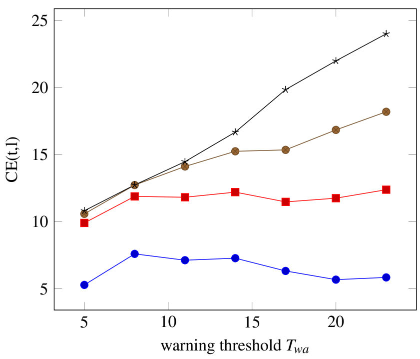

(e) No priorities, gap of 7 °C.

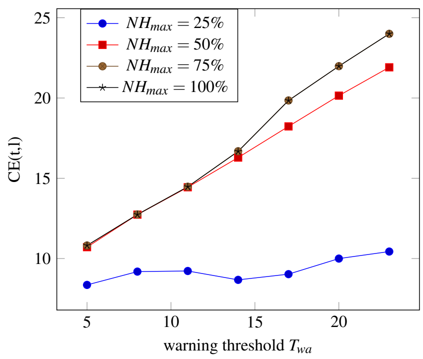

(f) With priorities, gap of 7 °C.

Fig. 9: The energy consumed by the high priority switches of the railway station is considered. The different lines correspond to different level of energy available to the whole system.

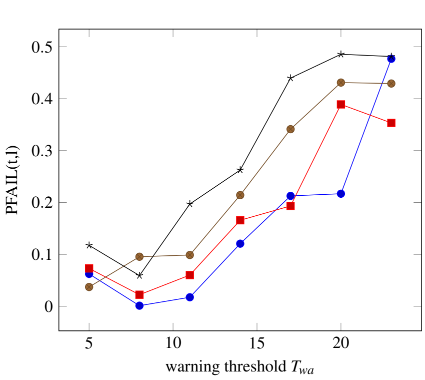

(a) No priorities, $NH_{max} = 50\%$.

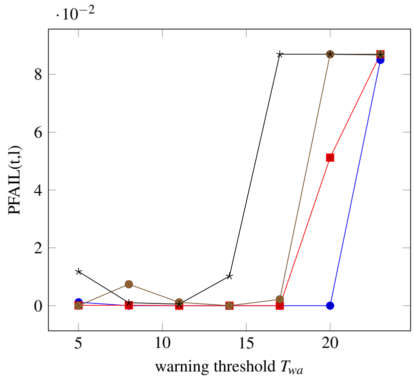

(b) With priorities, $NH_{max} = 50\%$.

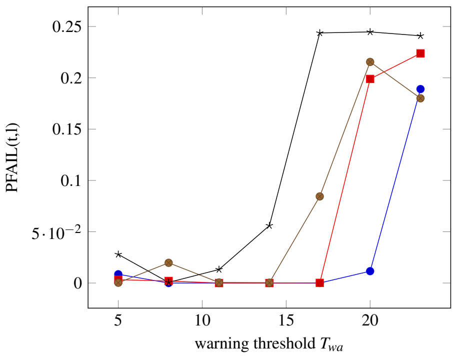

(c) No priorities, $NH_{max} = 75\%$.

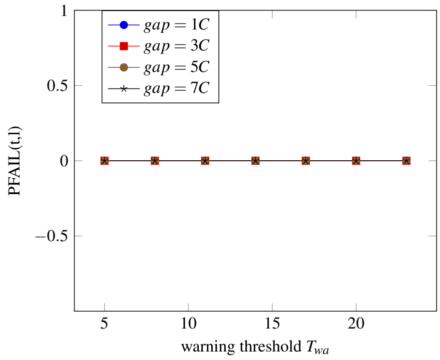

(d) With priorities, $NH_{max} = 75\%$.

Fig. 10: The probability of failure of the high priority switches of the case study is considered. The different lines correspond to different gap between T_{wa} and T_{wo}.

ment 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 [24]. Generalised Stochastic Petri Nets allows to express a finer model, with synchronizations and conflicts between different modules, that is shown to be more accurate in power saving than Markov Decision Processes models. In our case complex behaviours are modelled with SAN, which are a generalization of Generalised Stochastic Petri Nets. We also consider a policy of switching on/off the heater when a given temperature threshold is reached. We express a finer behaviour by using C++ code in the SAN model which computes the physical model of the heat transfer. Indeed the amount of time in which a heater is not used, (i.e. turned off) is derived from the external temperature, the internal temperature of the iron bar, and the heat transfer law.

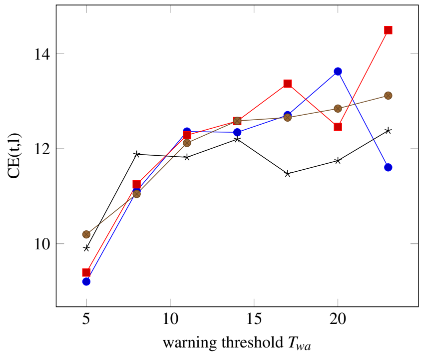

(a) No priorities, $NH_{max} = 50\%$.

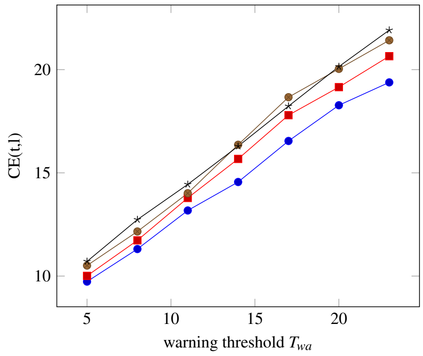

(b) With priorities, $NH_{max} = 50\%$.

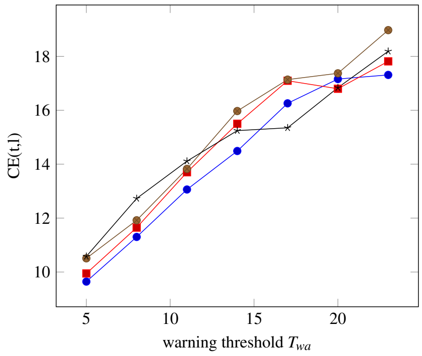

(c) No priorities, $NH_{max} = 75\%$.

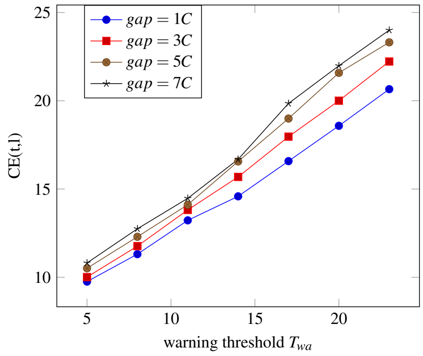

(d) With priorities, $NH_{max} = 75\%$.

Fig. 11: The amount of energy consumed by the high priority switches of the case study is considered. The different lines correspond to different gap between warning and working threshold.

The problem of power management in smart grids is handled with Learning Automata [23] in [21], that provides a mechanism to learn from the environment the optimal solution over a period of time. The model of the system is hierarchical: at the root there is a Learning Automata-based main power station, that supplies the power to Learning Automata-based transmission system, and adjusts the power supply according to requirements based on learning the system. The Learning Automata-transmission system calculates the performance of the system. The studied performance metrics are the power utilization and the customer satisfaction, in terms of satisfied energy demand. 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.

The dynamic power management problem is interpreted as a hybrid automaton control problem and integrated stochastic control in [15]. Hybrid automata mixed both a discrete state, representing the power mode of the system, and a continuous one, representing the consumed power. An integrated stochastic control is synthesised based on

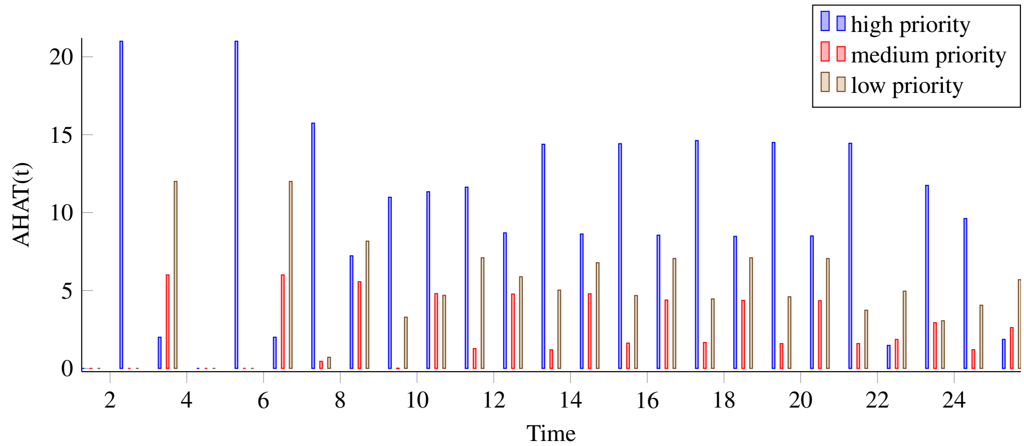

(a) The considered T_{wa} is 7 °C and T_{wo} is 8 °C (1 °C of gap). Priorities are considered.

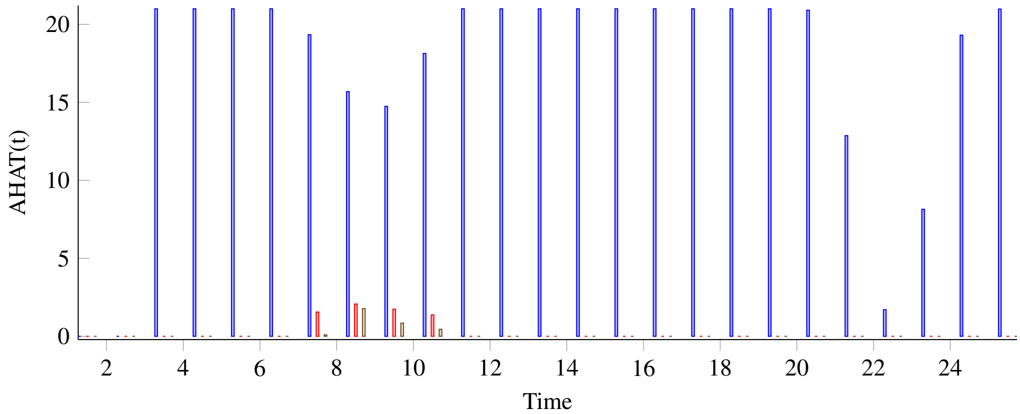

(b) The considered T_{wa} is 17 °C and T_{wo} is 25 °C (8 °C of gap). Priorities are considered.

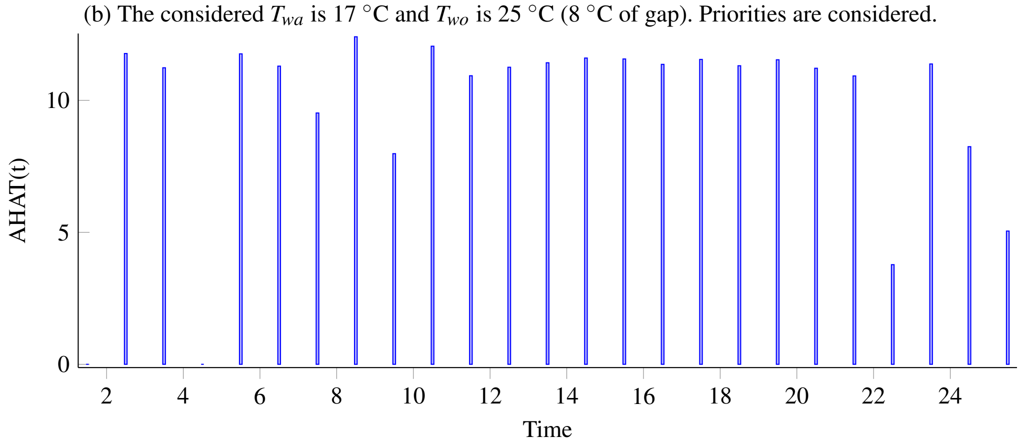

(c) The considered T_{wa} is 7 °C and T_{wo} is 8 °C (1 °C of gap). No priorities are considered.

Fig. 12: The plots show the amount of heaters activated each hour, the amount of power considered is sufficient to activate half of the heaters available in the network (i.e. $NH_{max} = 50\%$).

a learning feedback, and it is used to predict probabilistically the range lengths of the future idle period based on the past history. 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. In our work we do not implement learning mechanisms, and the power supplied is fixed by the number of heaters that can be turned on in a unit of time. It would be interesting to relax this constraint and implement a power adjustment with a prediction mechanism for minimising the power supplied, for example in case of warmer nights.

The applicability of self-organizing systems for different fields of power system control is discussed in [22]. 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). In our case we consider a central unit which manages the different heaters. The demand of energy is adjusted according to the maximum energy that can be delivered by the central unit. In case of failure of a heater, the energy is automatically shared among the remaining active heaters. We show that by managing the temperature thresholds it is possible to improve reliability even in case of low energy demand.

The survivability of a smart house is analysed in [16], 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 [12] 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. Instead, in our case the probability of failure is derived from the model. Moreover, we do not consider fixed thresholds, but we analyse how different values for thresholds affect the energy consumption and reliability.

The trade-off between energy saving and reliability is studied in [33], 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.

Concerning the analysis and optimization of a railway station using formal techniques, in [13] Stochastic Activity Networks are used to improve timetable and delay minimization of the traffic in a station. The model takes in input the railway topology and the required service. Experiments are performed to measure the capacity of the line in terms of number of trains that traverse the line, and the percentage usage of each track segment. In [20] an Automatic Train Supervision is designed that prevents the occurrence of deadlocks. A formal model that designs railway layout and the Automatic Train Supervision behaviour is used to verify such deadlock properties. The verification phase is performed by using the UMC model checking verification framework [17]. It would be interesting to integrate such studies with the possible failure of switches studied here, in order to analyse how a failure in a switch impacts on possible delays of trains, and deadlocks.

8 Conclusion and Future Work

We have presented the result of a research activity in model-based analysis for a rail road switch heating system. We used Stochastic Activity Networks to evaluate both the energy consumption and the reliability.

The developed SAN modelling framework represents the behaviour of the system that reads in real time both the temperature of the external air and the one of the rail road track, and according to given thresholds and priorities decides when to turn on and off the heaters. We evaluated the probability of failure of the system, the energy consumption, and the heaters activated per hours at the varying of several parameters that are the thresholds, the maximum number of heaters that can be turned on at the same time, and the different prioritizations. To represent the heat exchange between the portion of the rail road track, the external air and the heaters, we describe a physical model of heat exchange by convection. Simulations of the model have been taken using the Möbius tool [11].

We have considered a realistic scenario for our case study. The data concerning the layout of the railway station were taken from a real case [10], and the data for the temperatures in extreme cold conditions were retrieved from [32]. The physical model for the heater has been instantiated by taking values available at [8] [26].

We have categorized the heaters in the station in 3 different classes of priority, according to the layout of the station.

The results suggest that, in order to have a good trade-off between probability of failure and energy consumption it is important to guarantee enough energy to allow at least half of the overall heaters in the network to be activated, keep the warning threshold at around 7 °C and use a small gap between the minimum and the maximum temperatures (1 °C), in order to better distribute the time in which each heater is turned on.

Moreover, a prioritized approach has been proved better than one without priorities for both the energy consumption and the reliability of the system.

Several directions for extending this study have been identified. Indeed, the presented model can be tailored to reflect different scenarios of energy optimization. For example, energy optimization could be applied to the enlightening of a station, by turning off the lights when they are not needed. Continuing this example, the temperature should represent the quantity of light during that part of the day, the thresholds should represent when the lights must be turned on and off, while the energy consumption equation should take into account the amount of consumed electricity.

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.

Finally, adapting the thresholds to the different class of priorities may increase the reliability of the system.

Acknowledgements This work has been partially supported by the projects TENACE (PRIN n.20103P34XC) and CINA (PRIN n.2010LHT4KM), both funded by the Italian Ministry of Education, University and Research.

References

1. Alur, R., Dill, D.L.: A theory of timed automata. Theoretical Computer Science 126(2), 183 – 235 (1994)
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. Basile, D., Chiaradonna, S., Giandomenico, F.D., 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), http://dx.doi.org/10.1007/978-3-319-23129-7_7
5. Bause, F., Kritzinger, P.S.: Stochastic petri nets: An introduction to the theory. SIGMETRICS Perform. Eval. Rev. 26(2), 2–3 (Aug 1998)
6. Berger, L.T., Schwager, A., Escudero-Garzás, J.J.: Power line communications for smart grid applications. JECE 2013, 3:3–3:3 (Jan 2013)
7. Bernardi, S., Merseguer, J., Petriu, D.C.: Model-Driven Dependability Assessment of Software Systems. Springer (2013)
8. 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
9. Cannon, J.R.: The One-Dimensional Heat Equation. Cambridge University Press (1984), cambridge Books Online
10. https://en.wikipedia.org/wiki/Carlisle_railway_station
11. 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)
12. David, R., Alla, H.: On hybrid petri nets. Discrete Event Dynamic Systems 11(1-2), 9–40 (Jan 2001)
13. Di Giandomenico, F., Fantechi, A., Gnesi, S., Itria, M.L.: Stochastic model-based analysis of way operation to support traffic planning. In: Gorbenko, A., Romanovsky, A., Kharchenko, V.S. (eds.) Software Engineering for Resilient Systems, 5th International Workshop, SERENE 2013. Proceedings. LNCS, vol. 8166, pp. 184–198. Springer (2013)

14. Diab, Hassan B.; Zomaya, A.Y.: Dependable computing systems: paradigms, performance issues and applications. John Wiley & Sons (2005)
15. 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)
16. Ghasemieh, H., Haverkort, B.R., Jongerden, M.R., Remke, A.: Energy resilience modelling for smart houses. In: 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2015. pp. 275–286. IEEE (2015)
17. Gnesi, S., Mazzanti, F.: An abstract, on the fly framework for the verification of service-oriented systems. In: Wirsing, M., Hölzl, M. (eds.) Rigorous Software Engineering for Service-Oriented Systems, LNCS, vol. 6582, pp. 390–407. Springer Berlin Heidelberg (2011)
18. 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
19. Ivanov, D., Orlic, M., Seceleanu, C., Feljan, A.V.: Remes tool-chain - a set of integrated tools for behavioral modeling and analysis of embedded systems. In: Proceedings of the 25th IEEE/ACM International Conference on Automated Software Engineering (ASE 2010) (September 2010)
20. Mazzanti, F., Spagnolo, G.O., Longa, S.D., Ferrari, A.: Deadlock avoidance in train scheduling: A model checking approach. In: Lang, F., Flammini, F. (eds.) Formal Methods for Industrial Critical Systems - 19th International Conference, FMICS 2014. Proceedings. LNCS, vol. 8718, pp. 109–123. Springer (2014)
21. 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)
22. 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)
23. Narendra, K.S., Thathachar, M.A.L.: Learning Automata: An Introduction. Prentice-Hall, Inc., Upper Saddle River, NJ, USA (1989)
24. 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)
25. Qiu, Q., Wu, Q., Pedram, M.: Dynamic power management of complex systems using generalized stochastic petri nets. In: DAC. pp. 352–356 (2000)
26. http://www.railsco.com/~electric_switch_heater_controls.htm
27. 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)
28. Seceleanu, C., Vulgarakis, A., Pettersson, P.: Remes: A resource model for embedded systems. In: Engineering of Complex Computer Systems, 2009 14th IEEE International Conference on. pp. 84–94 (June 2009)
29. 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
30. Trivedi, K.S.: Probability & statistics with reliability, queuing and computer science applications. John Wiley & Sons (2008)

31. 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)
32. https://weatherspark.com/#!graphs;ws=27985
33. 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)