

Stochastic Model-Based Analysis of Energy Consumption in a Rail Road Switch Heating System ^{*}

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

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

Abstract. Rail road switches are critical components, since they enable trains to be guided from one track to another, thus potentially resulting in catastrophic consequences (e.g. trains collision) in case of wrong manoeuvring. Among the facilities to help preserve their correct functioning, rail road switches heaters are used to avoid the formation of snow and ice.

Managing the energy consumption of these devices is important in order to reduce the costs and minimise the environmental impact. While doing so, it is important 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 stochastic models based on the Stochastic Activity Networks (SAN) formalism. An on-off policy is considered for heating the switches, with parametric thresholds representing the temperatures activating/deactivating the heating. Initial investigations are carried on to understand the impact of different thresholds on the indicators under analysis (namely, probability of failure and energy consumption).

1 Introduction

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. In this work, we do not consider a failure in the rail road switch heating system as a safety failure (citing [2] "*the possibility of human lives being endangered*"). Indeed interlocking mechanisms guarantees safety in case of failure of switches, and in this work we only consider failures in the heating system.

Unfortunately, during winter, snow and ice can prevent the switches to work properly. Indeed the mechanism which allows a train to be directed can be blocked by an

^{*} This work has been partially supported by the TENACE PRIN Project (n.20103P34XC) and the Italian MIUR project CINA (PRIN 2010LHT4KM) funded by Italian Ministry of Education, University and Research.

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 clear manually by employers who were sweeping the snow away. More recently, heaters are used so that the temperature of the rail road switch can be kept above freezing. The heaters may be powered by gas or electricity.

The managing of the heaters is automatic, and is controlled by a central unit using a Powerline [5]. Powerline communications have the possibility of transmit coded informations through the existing electric lines. A Powerline is a transmission system that uses electric lines, with a very extensive infrastructure in nearly each building. The main advantage of adopting the existing network is the absence of additional costs for the installation of the infrastructure.

Powerlines are used in order to automate and optimize several tasks i.e. switching on and off the lights of the station, checking the status of railway track switches, managing the traffic of trains in the station and the rail road switch heating system.

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 financial terms and in environmental impact. 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.

In this paper, we address reliability and energy consumption of rail road switch heaters. A failure in those switches can lead to major diseases. Indeed, while managing energy optimization we must ensure reliability. In fact, by turning on all the heating system 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. 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 by selecting minimum and maximum temperature thresholds, which guarantees a satisfactory trade-off.

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 preliminary version.

In particular, we exploit *Stochastic Activity Network* (SAN) [16] to model the heating system, and use the Möbius tool [6] to perform experiments.

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.

In [15] Generalized Stochastic Petri Nets [3] are used to solve the dynamic power management problem for systems with complex behaviour. Dynamic power management studies the problem of reducing 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 turn on a component when it is used and turn it off when it is not used for a certain amount of time. Their model is compared with other models based on Markov Decision Processes (MDP). By exploiting GSPN they express a finer model, with synchronizations and conflicts between different modules, that is shown to be more accurate in power saving than MDP models. In our case complex behaviour are modelled with SAN, which are a generalization of GSPN. We also consider a policy of switch 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 compute the physic model of 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.

In [13] the problem of power management in smart grids is handled with Learning Automata (LA), that provide a mechanism of learning from the environment the optimal solution over a period of time. The model of the system is hierarchical: at the root there is a LA-based main power station, that supplies the power to to LA-based transmission system, and adjust the power supply according to requirements based on learning the system. The LA-transmission system calculate the performance of the system. The studied performance metrics are the power utilization and the customer satisfaction, in terms of satisfying the demand of energy. 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 customer satisfaction.

In [8] the dynamic power management problem is interpreted as a hybrid automaton control problem and integrate stochastic control. 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 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 provide 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.

In [14] the applicability of self-organizing systems for different fields of power system control is discussed. 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 manage 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 between 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.

In [9] the authors analyse the survivability of a smart house, 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 [7] are used for modelling this scenario. Different strategies of battery management are considered. In the first 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. In our case the probability of failure is more accurate, as it is derived from the model. We do not consider fixed thresholds, but we analyse how different values for thresholds impact in the energy consumption and reliability. Moreover SAN are shown to be more accurate as they can handle continuous places, stochastic transitions, and C++ code to simulate complex behaviour such as heat transfer by convection.

Concerning the analysis and optimization of a railway station using formal techniques, in [10] Stochastic Activity Network 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 taken 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 [12] an Automatic Train Supervision is designed that prevents the occurrence of deadlocks. A formal model that design 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 [11].

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 impact on possible delays of trains, and deadlocks.

Structure of the paper In Section 2 we describe the considered rail road switch heating system. Stochastic model-based analysis is introduced in Section 3, while in Section 4 we present the models of the rail road switch heating system. Finally, the results of our experiments are discussed in Section 5, while final conclusions are in Section 6.

2 The System Under Analysis

In Figure 3 a rail road switch heater is displayed. The picture is taken from [1]. The heating system 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, and checks if the temperatures are between given thresholds [1]. Based on this general behaviour, we have based the policy employed to activate/deactivate the heating on two threshold temperatures:

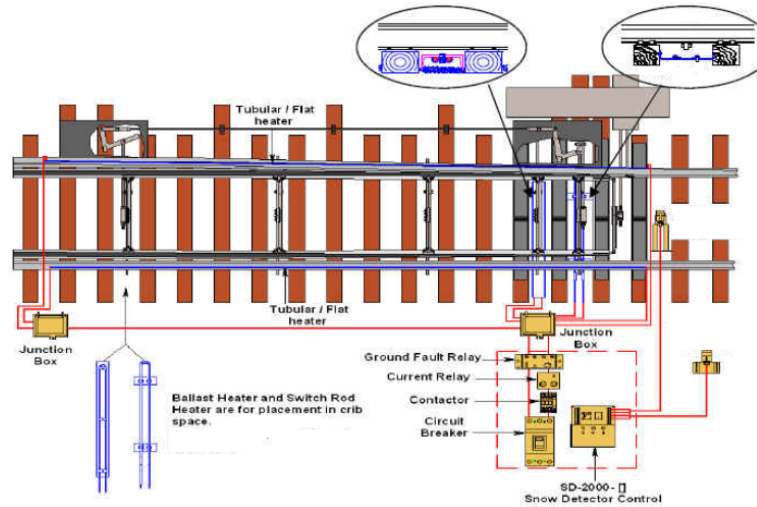


Fig. 1: An example of a rail road switch heater heated by electric induction ©2006 - 2012 Rails Company

- *warning threshold*: this temperature represents the lower temperature that the system should not trespass. If the temperature is lower than the warning threshold, then the risk of ice or snow can lead to a failure of the rail switch and therefore the heating system needs to be activated;
- *working threshold*: 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.

Hence the time during which a single heater is active depends on its location and the weather conditions. Heaters located in a colder geographical area will consume more energy than those under warmer temperatures.

The network of heaters will share information through central computational unit, using the Powerline. The central unit manages the maximum amount of power that can be delivered to the system, in order to prevent possible blackouts.

3 Stochastic Model-based Analysis

Stochastic model-based approaches are useful to support the development of new systems, in all the phases of their life cycle. In the early design phases it is important to validate a model of a system in order to avoid waste of time and resources in the development phase. This can be done by pointing out the properties and the requirements of the system, building a model that represents its behaviour and checking that the properties are satisfied by the model. It is possible to choose between 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 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 behaviour to plan the maintenance and the upgrading operations for the system.

3.1 Stochastic Activity Network

Several stochastic modelling methodologies have been proposed in literature. Stochastic Activity Networks [16] are a widely adopted formalism for the analysis of systems under performance, dependability and quality of service. The formalism is a generalization of Stochastic Petri Nets [4], and has similarities with Generalised Stochastic Petri Nets [3]. A SAN is composed of the following entities: *places*, *activities*, *arcs*, *input gates* and *output gates*. Places in SAN have the same interpretation as those of Petri Nets. Activities are of two types: *instantaneous* and *timed*. Instantaneous activities are fired once the enabling conditions are satisfied. Timed activities are fired following a temporal stochastic distribution of time. There are different policies of activation and reactivation of timed activities, for a marking based policy of reactivation of timed activity. An enabled activity is aborted when the SAN moves into a new marking in which the enabling conditions of the activity no longer holds. Cases are associated to activity, and are used to represent uncertainty about the action taken upon completion of the activity. A marking is stable if no instantaneous activity is activated. Input gates control the enabling of the activity and the change of marking at completion of the activity. Output gates define the change of marking upon completion of activity. It is possible to use C++ code to change of marking. When an activity completes, the following sequence of events is executed:

- one of the cases of the activity is chosen according to its probability,
- the functions of the connected input gates are executed,
- one token is removed from the places connected by the input arc,
- the function of the output gates connected to the activity are executed,
- one token is added to the places connected to the activity or one of its cases by the output gate.

For evaluating the energy consumption and the probability of failure of the system modelled, we use the Möbius tool [6]. Möbius is a software tool for modelling the behaviour of complex systems, supporting various formalisms such as SAN, PEPA, Fault Tree, etc... Developed models are then solved by using different analytical and simulative solvers. This tool 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.

4 Modelling Framework

In this section we describe the rail road heating system model.




Fig. 2: On the left the SAN model ProfileSelector, on the right the SAN model LocalitySelector.

We developed a SAN model to represent the system which is actually structured as composed by five atomic sub models, properly combined through Rep and Join operators (see Figure 4).

Three of the five atomic models are selectors for the profile, locality and the unique identifier. The remaining two are the atomic model for the queue, shared among the replicas, and the main model for the rail road switch heater.

The main parameters of the SAN model are the lower and working temperature of the device, and the maximum power that the system can provide every instant of time, i.e. the maximum number of heaters that can be turned on at the same time.

Weather forecast. To model the external weather conditions, our model takes in input a table containing profiles of average temperature in winter days. Each night will be divided in intervals of two hours, containing the temperature in that moment. The considered interval of time will be from 6 pm to 6 am. We do not consider a whole day because the probability of having temperatures below zero during the day is lower than considering the night only. Moreover, consider a whole day does not affect the main results.

At starting time, a SAN will select one of these profiles. A probability is assigned to each profile. Less frequent profiles will have a lower probability. The probability in which a particular profile will be selected, depends on the aforementioned value. In Figure 2 left, the SAN model corresponding to the action of selecting the profile is displayed. The probability is selected in the SAN model with a timed activity with different cases, each case corresponds to a profile. The probability in which a profile is selected is higher for average temperatures, while it will be lower for less frequent temperatures. In our study five profiles are available. A C++ function in the output gate will load the selected profile.

We note that an alternative way of deciding the actual temperature would be to select stochastically the temperature depending on the previous temperature and the actual time. However, by adopting this methodology we would have obtained a non realistic zig-zag evolution of the temperature during the night. Indeed after each interval temperatures would have a non-zero probability of increasing or decreasing. Thus there could be the unwanted case where temperatures increase and decrease several times during the night.

Location. Rail road switch heater are located in different zones, with different weather conditions, while we assume that they are remotely controlled by the central management unit. At a higher altitude, the probability to have temperatures lower than the average selected for that profile is greater than the case of rail road switch heaters located at lower altitudes. For modelling this variation of temperature, the algorithm takes as input a vector of locations, which are numerical values representing the gap from the average temperature for that profile. In case there are many rail road switch heaters located at higher altitudes, it will be more probable that the actual temperature will be lower than the average. The SAN that selects the location for the device is depicted in Figure 2 right. The network works similarly to the one for the selection of the profile. A timed activity has different cases, according to the different probabilities of selecting a location. Once the activity is fired, a C++ function is called which assigns the given probability to the rail road switch heating module.

4.1 Rail Road Switch Heater

In Figure 3 the SAN model representing the rail road switch heater is depicted. The SAN takes in input the profile and the location of the device, and a unique identifier. This unique identifier will be computed by a separate SAN, not shown here. There are two subnets present in the network.

Heater subnet representing the status of the heater, which can be switched on or off, or it can be a failure state. An extended place *temperature* represents the internal temperature of the device, while the places *on* and *off* represent the status of the system. A shared place *sharedOn* is useful to know how many heaters are turned on at a given time. According to the heating policy, once the system temperature goes below a predefined *warning threshold*, the heating needs to be activated otherwise a switch failure is experienced (represented by the place *failure*). Then, once the temperature raised and reached the *working threshold*, the heating system can be safely turned off.

The energy consumption of the overall system depends on the value of those thresholds. A smaller gap between the two values represent a frequent activation of the heating system, but for a shorter period of time. Alternatively, by taking a wider gap between the two values, we will obtain a less frequent activation, but the activation will be for longer periods of time. In the next section, we instantiate the model with varying values for these thresholds, to show their impact and benefit on both energy consumption and system reliability.

Clock subnet represents a clock which updates at each unit of time the parameters, i.e. the temperature of the rail road and the temperatures of the air. For this case study, the unit of time is assumed to be one hour. Each time the clock activity is fired, the temperatures will be updated in the output gate using C++ functions. We model the physical behaviour of the rail road in terms of temperature decay and increase, when the heating is switched off and on, respectively. For the temperature of the air, a new value will be picked up by the table for the selected profile and depending on the location of the device, this value will be higher or lower than a given amount of degrees, selected in the locality table.

For the internal temperature of the device, we will model the actual temperature of the system by using an equation of heating exchange through convection, which simulates the data that will come as input from the sensors.

The *energy* place represents the amount of energy consumed by the system. Each time the clock activity is fired, if the heater is switched on, then the energy consumed is augmented. By taking into account the power consumed by the heater, it is possible to calculate the consumption of energy in kilowatt.

Queue In Figure 3, the place *free* represents the queue of active rail road switch heaters. This place is shared among all the instances of the SAN. The capacity of the queue represents the number of heaters that can be turned on at the same time, that we call NH_{Max} . For example, assuming that the power consumed by a single heater per hour is $35KW$, if $NH_{Max} = 2$ then the maximum energy that the system can provide is $70KW$: no more than two heaters can be turned on at the same time. If the capacity of the queue is exceeded, a blackout failure may occur. Hence the thresholds for the activation of the switches must be chosen also taking into account this constraint. With the queue model, it is also possible to implement different priorities in which a particular rail road switch heater must be turned on and off.

4.2 Physical Model for Heat Exchange

As mentioned above, we need to simulate the data read by the sensors in order to estimate the energy consumption of the heaters. To do so, we instantiate a physical model representing 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 an electric cable that passes through the rail road in different points in order to warm up the iron. 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 (6pm - 6am).

Every hour the sensor reads a new data for the internal temperature of the rail road track. Assuming that the value of the temperature of the air and the previous internal temperature are known, we foresee the updated internal temperature of the device using the following equation representing exchange of heat by convection. This equation is derived from a differential equation on the time:

$$T_{fin} = T_a + (T_i - T_a) \cdot e^{\frac{-u \cdot A \cdot t}{m \cdot c}} + \frac{Q}{u \cdot A \cdot L}$$

The coefficient of convective exchange u is calculated as:

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

The parameters of the previous equation are: T_{fin} is the new internal temperature of the heater; u is the the coefficient of convective exchange; c is the heat capacity of iron;

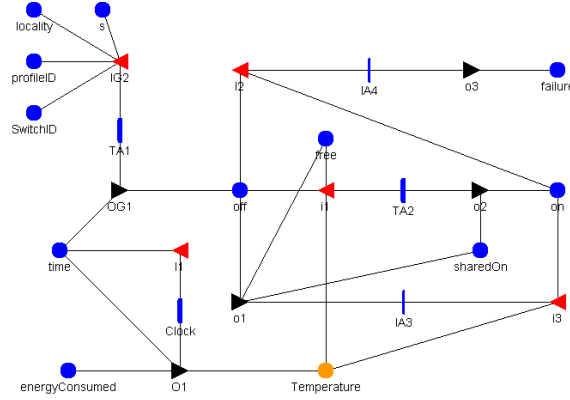


Fig. 3: The SAN model RailRoadSwitchHeater.

A is the surface area exposed to the external temperature; AV is the ratio between area and volume of the iron bar; t is the interval of time, one hour in our case; m is the mass of the iron bar; T_i is the previous internal temperature; T_e is the external temperature of the surrounding area; Q is the power used when the heater is turned on, if the heater is turned off this value will be zero. Moreover L is the length of the electric cable for heating the rail road; g is the gravity acceleration; β is the thermal expansion coefficient; ρ is the density of air; μ is the dynamic viscosity and K is the thermal conductivity of iron.

The function representing the heating exchange is written in C++, and it is called by the output gate $O1$ displayed in Figure 3 to update the temperature of the rail road every hour. This C++ module is the portion of the system representing the real world model that we want to optimize. This module can be easily modified according to model different scenarios of energy optimization. For example, there could be the need to optimize not only the energy consumed by the rail road switch heaters. Indeed, the enlightening of a station could be optimized by turning off the lights when they are not needed. In this case, 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.

4.3 The composed model

The composed model is displayed in Figure 4. The atomic SAN models are *RailSwitch-Heater*, *LocalitySelector*, *ProfileSelector*, *SwitchIDSelector* and *Queue*. The box *Join1* represents the join of all the atomic models except *Queue*. Those atomic models share the places relative to the locality of the device, its weather profile and unique ID. The box *Rep1* represents the network of heaters. A parameter *numRep* identifies the number of devices composing the network. Recall that each device has its own weather forecast profile and locality. Finally, the node *Join2* composes the network of rail road switch

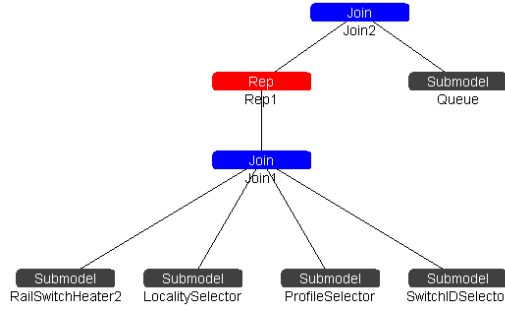


Fig. 4: The composed model

heaters with the model representing the Queue. Indeed all the sub models share the same queue.

5 Analysis results/Experiments

In this section we describe the preliminary experiments we have performed in order to find suitable trade-offs in terms of reliability and energy consumption for different settings of the model parameters. These experiments do not refer to a rail road configuration taken from reality, but consider a restricted plausible configuration adopted as a working example to show the benefits of our analysis framework. Indeed, 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.

5.1 Measures of Interest

We consider two different measures of interest which represent the energy consumption and reliability of the system under analysis:

- 1 $CE(t, l)$: the mean energy consumed by a heater in the interval $[t, t + l]$;
- 2 $PFAIL(t, l)$: the probability of failure of the system at time $t + l$, defined as the mean probability that a switch become frozen during the interval of time.

The interval $[t, t + l]$ goes from 6:00 pm to 6:00 am. $CE(t, l)$ is defined by accumulating in the interval $[t, t + l]$, that is the time that each replica of the SAN model RailSwitchHeater spends in the marking represented by one token in the place *ON*, multiplied for the energy consumed in a unit of time. $PFAIL(t, l)$ is defined as the probability that at time $t + l$ there is one token in the place Failure of the SAN model RailSwitchHeater. $PFAIL(t, l)$ can be used to compute the mean time to a catastrophic failure.

Other measures may help in understanding aspects related with energy consumption and system failure, so to take appropriate actions in improving the system. Some of them are introduced in the following, although their quantitative evaluation is left as future work:

	06 pm	08 pm	10 pm	00 am	02 am	04 am	06 am	Probability
profile 1	2°	0°	-2°	-2.5°	-3°	-3.3°	-3.8°	0.25
profile 2	0°	-1°	-2°	-3°	-4°	-4.5°	-5°	0.2
profile 3	3°	1°	0°	-1°	-2°	-3°	-3.5°	0.2
profile 4	0°	-2°	-3.5°	-4.8°	-6.2°	-6.8°	-6.9°	0.15
profile 5	1°	-1°	-3°	-3.5°	-4°	-4.5°	-5.5°	0.2

Table 1: The profiles of temperatures per night

	Temperature Variation	Probability
locality 1	-0.8 °	0.3
locality 2	-0.3 °	0.3
locality 3	0 °	0.4

Table 2: The temperature variation and the probability of occurrence of such variation for the three considered localities

- the probability of heating activation in an interval of time; it is useful for understanding in which hours the probability that a heater is switched on is greater, i.e. the hours of the night with a greater consumption of energy;
- the probability of activation of a given number of heaters in an interval of time; it is useful for improving planning operations performed by the central computational unit;
- the average number of heaters switched on in an interval of time; again, it is useful for planning purposes.

5.2 Scenarios and Settings

The table of the profiles used for the experiments are displayed in Table 1. We consider average cold winter nights, which are relevant for our study. Starting from the temperature at 6 pm, it decreases and reaches the minimum at 6 am.

We consider three main localities, displayed in Table 2. We note here that it is more probable to select a locality at a lower altitude, where there is no variation between the selected profile temperature. It must be pointed out that these parameters can be easily modified to deal with different conditions.

The parameters considered for calculating the heat exchange between the rail road, the heaters and the external air are showed in Table 3. We consider an iron bar of 6 meters of length, 4 cm of height and 28 cm of width to represent the portion of rail road track that needs to be heated. In the table l, h and w are respectively the length, height and width of the iron bar. The other parameters are described in Section 4, and we use standard value for the parameters of the air and of the iron, i.e. viscosity, density, conductivity, heat capacity etc... We assume that all the rail road switch heaters use the same amount of power, which in this case is 35 KiloWatt.

For this case study, we set the freezing thresholds to zero. If the temperature goes below this value the system reaches a failure state. In our experiments we consider four

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	Q
$17.157 \cdot 10^{-6} Pa \cdot sec$	0.023 $\frac{W}{mK}$	0.0175 m	1.68 m ²	529.13 Kg	450 $\frac{J}{KgK}$	35 KW

Table 3: the parameters for the model of heat exchange by convection

different combinations of thresholds, at different temperatures and with tight or wide gaps between the working threshold and the warning threshold.

The network is composed of four rail road switch heaters, and we study the results when varying of NH_{Max} , representing how many heaters can be turned on at the same time, that in turn represent the maximum throughput of energy of the system in a unit of time. We assume that all the heaters have initial temperature of 5 degrees Celsius.

The four pairs of thresholds are (1,4), (4,7), (5,6), (4,50). We consider different gaps between the thresholds. Moreover we study the behaviour of the heating system for thresholds that are more and more distant from the temperature of failure.

The experiments were taken on the Möbius tool [6], by using simulation with 10000 batches. The results were computed in one minute on a machine with processor Intel Core i3.2328M 2.20 GHz with 8 GB of RAM, showing that the model is well defined.

5.3 Results Discussion

We report the results for two of the aforementioned measures of interest: $CE(t, l)$ and $PFAIL(t, l)$.

The results for the probability of failure are reported in Figure 5. As expected, by minimizing the throughput of energy the probability of failure increases.

Moreover, the probability of failure increases in case the gap between the two thresholds is greater. Indeed the time taken by the heater to warm up the rail road is higher when the gap is wider. Hence the temperature of the other pending heaters may go below the freezing threshold, leading to a failure of the system.

The results for the energy consumption are reported in Figure 6. We note that by increasing NH_{Max} we have more energy available, which in turn results in a major energy consumption (recall Section 4). The same happens when we consider greater value for the thresholds.

Interestingly, a wider gap is worse than a strict gap when NH_{Max} is large enough, while the converse holds in case NH_{Max} is limited.

The threshold (4,50) is used for checking the consumption of energy in case the heater once activated, will never be turned off during the whole night. Indeed at 35 KW of power, the system never reaches the working temperature. Of course this threshold is worst both for the probability of failure, and for the energy consumption. Note that if the starting temperature is not below the warning threshold, the heater will not be activated for the whole night. Indeed the temperature of the heater needs some hours to fall below the warning threshold. For example, in Figure 6 for thresholds (4,50) and $NH_{max} = 1$ we have $CE(t, l) = 92.8 \approx \frac{10.5 \cdot 35}{4}$, so after two or three hours the temperature of the track will go below 4 degrees Celsius.

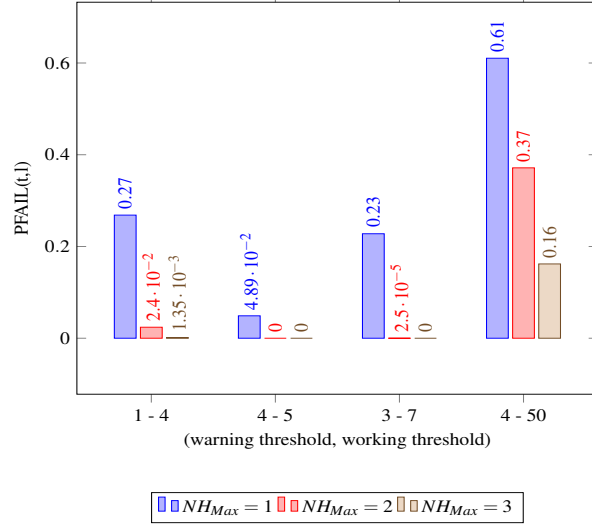


Fig. 5: The graph for the probability of failure

Our experiments show that it is better to take a tight gap between the minimum and the maximum thresholds. Also keeping those thresholds not too close to the freezing threshold improve reliability, as well as guaranteeing a sufficient amount of heaters to be activated at the same time. Indeed this is necessary for guaranteeing the reliability of the system and minimizing the energy consumption.

Hence, the best scenario in this example is represented by the pair of thresholds (1,4) for $NH_{Max} = 3$ and (4,5) for $NH_{Max} = 2$.

6 Conclusion and Future Work

We have presented the result of a preliminary 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 probability of failure.

The system reads in real time both the temperature of the external air, the one of the rail road track, and according to given thresholds decides when to turn on and off the heaters. We evaluated the probability of failure of the system and the energy consumption at the varying of the thresholds and the maximum number of heaters that can be turned on at the same time.

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 [6]. At the moment we have instantiated the model to representative (but realistic) parameters, to obtain a first indication of the behaviour of our system. 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 two heaters working simultaneously;

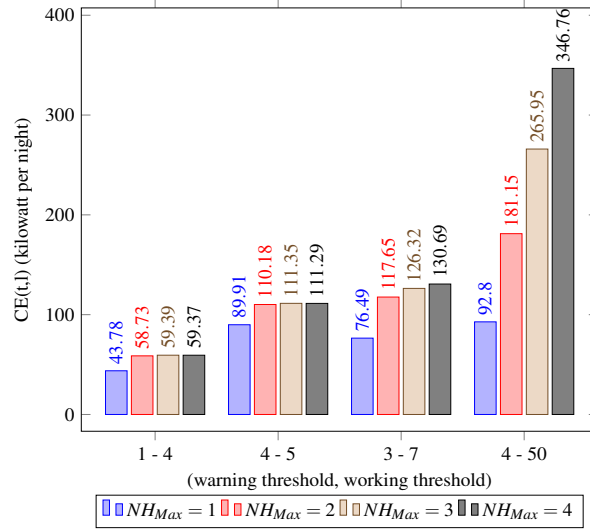


Fig. 6: The graph for the energy consumption

- keep the minimum temperature distant from the temperature of failure;
- use a small gap between the minimum and the maximum temperatures, in order to better distribute the time in which each heater is turned on.

This work represents a first step in the design of a rail road switch heating system with attention to energy consumption and reliability, which often are opposed requirements. We plan to instantiate our model to case study from real world, i.e. an existing rail road switch heating system, to check if our results match the expectations.

The measures of interest unspecified in Section 5 are left as future work. By studying the energy consumption at the varying of those measures we feel confident to obtain more insight on the behaviour of the system and its energy consumption.

Moreover we plan 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 of which the heaters are composed, 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.

Moreover, by implementing a priority queue, it is possible to develop strategies for activating and deactivating those heaters which are in a critical situation, i.e. when the temperature is closer to the failure point. At the moment, in case of conflicts between different heaters that need to be turned on, the choice is made stochastically by the system.

We plan to extend the model to other different case study for energy optimization in the context of a railway station, exploiting the modularity of the proposed approach.

Acknowledgements We are thankful to Gianpaolo Roina for the useful comments.

References

1. http://www.railsco.com/~electric_switch_heater_controls.htm
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. Bause, F., Kritzinger, P.S.: Stochastic petri nets: An introduction to the theory. *SIGMETRICS Perform. Eval. Rev.* 26(2), 2–3 (Aug 1998)
5. Berger, L.T., Schwager, A., Escudero-Garzás, J.J.: Power line communications for smart grid applications. *JECE* 2013, 3:3–3:3 (Jan 2013)
6. 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
7. David, R., Alla, H.: On hybrid petri nets. *Discrete Event Dynamic Systems* 11(1-2), 9–40 (Jan 2001)
8. 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)
9. Ghasemieh, H., Boudewijn, R., Haverkort, M.R.J., Remke, A.: Energy resilience modeling for smart houses. In: *DSN (2015)*, to appear
10. Giandomenico, F.D., Fantechi, A., Gnesi, S., Itria, M.L.: Stochastic model-based analysis of railway 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)
11. 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)
12. 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)
13. 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)
14. 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)
15. Qiu, Q., Wu, Q., Pedram, M.: Dynamic power management of complex systems using generalized stochastic petri nets. In: *DAC*. pp. 352–356 (2000)
16. 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)