Enhancing sustainability of the railway infrastructure: trading energy saving and unavailability through efficient switch heating policies

Silvano Chiaradonna[a], Giulio Masetti[a,*], Felicita Di Giandomenico[a], Francesca Righetti[b], Carlo Vallati[b,*]

[a] *Consiglio Nazionale delle Ricerche, Via Moruzzi 1, Pisa, Italy*
[b] *University of Pisa, Via Diotisalvi 2, Pisa, Italy*

Abstract

Railway is currently envisioned as the most promising transportation system for both people and freight to reduce atmospheric emission and combat climate change. In this context, ensuring the energy efficiency of the railway systems is paramount in order to sustain their future expandability with minimum carbon footprint. Recent advancements in computing and communication technologies are expected to play a significant role to enable novel integrated control and management strategies in which heterogeneous data is exploited to noticeably increase energy efficiency. In this paper we focus on exploiting the convergence of heterogeneous information to improve energy efficiency of railway systems, in particular on the heating system for the railroad switches, one of the major energy intensive components. To this aim, we define new policies to efficiently manage the heating of these switches exploiting also external information such as weather and forecast data. In order to assess the performance of each strategy, a stochastic model representing the structure and operation of the railroad switch heating system and environmental conditions (both weather profiles and specific failure events) has been developed and exercised in a variety of representative scenarios. The obtained results allow to understand both strengths and limitations of each energy management policy, and serves as a useful support to make the choice of the best technique to employ to save on energy consumption, given the system conditions at hand.

Keywords: Energy efficiency; railway infrastructure; heating control policies; stochastic model-based analysis; unavailability; data integration

1. Introduction

The replacement of road transport with rail transport is commonly recognized as the best strategy to significantly reduce atmospheric emissions of transportation of both people and freight [1]. A large-scale transition to railway, however, requires the further improvement of its energy efficiency, in order to ensure the long term sustainability and the necessary reduction of carbon emissions.

Recent computing and communication technologies are expected to play a crucial role in improving the energy efficiency of railways systems. The Internet of Things (IoT) in particular is expected to integrate different sub-systems of the railway into an all-in-one management infrastructure in which heterogeneous data can be collected. This computing paradigm is expected to enable the implementation of novel control strategies based on the availability of heterogeneous data and the fusion with data collected from external systems. In particular, improvements are expected in the energy efficiency of the whole system [2], along with other features like the fulfillment of dependability properties.

This paper tackles the challenge of energy efficiency in the railway transportation system, with focus on the heating system for railroad switches. This is a critical subsystem, responsible for keeping the switches free from snow and ice, necessary to guarantee the correct operation of the switches and so the correct train routing. Depending on the climate conditions of the place where the railway system operates, the energy consumed by this heating system can be very relevant. To provide

*Corresponding author
Email address: `giulio.masetti@isti.cnr.it`, `carlo.vallati@iet.unipi.it` (Carlo Vallati)

concrete examples, in [3] it is reported that the cost for heating the 6800 switches and crosses in Sweden can amount to $10-15$ MEuro/year. In Germany, Deustche Bahn (DB) alone has 64,000 switches heated with electrical resistance and gas heaters, a combined power of 900 MW which consume up to 230 GWh/year [4]. However, while energy saving is the primary target of our study, we are aware that the criticality of the targeted system imposes that dependability requirements be not endangered. Therefore, this work also addresses system unavailability, since the recommended solution lies in a satisfactory trade-off between energy consumption and unavailability.

Reducing energy consumption in the addressed context is pursuable through two major directions: either enhancing the switch heating and protection technology, or improving the policy dedicated to switch on/off the heating. This paper works in line with the second direction and in particular aims at exploiting the accessibility of heterogeneous data from different sources, both internal and external to the railway system, to improve energy efficiency, but keeping the dependability level as required.

The definition of a basic threshold-based policy that exploits external data, like the weather forecast data, has been already presented in a previous work by a subset of the authors of this paper [5]. In this work we significantly advance the study by proposing: i) novel switch heating management policies, targeting efficiency in terms of power consumption, but without impairing the availability requirement; ii) a refined framework for the model-based analysis of such policies, in order to quantify indicators representative of their goodness and appropriate trade-offs thereof, as a support to the selection of the most adequate one to employ in a specific real context; and iii) an evaluation campaign on a representative case study, by exercising the developed analysis framework, to compare the proposed energy management policies in a number of climate profiles. It is worth to mention that the proposed policies can be implemented in any IoT platform that manages the railways infrastructure. The policies only rely on basic communication capabilities between a controller and the sensors and actuators that manage the railway infrastructure.

The rest of the paper is organized as follows. Section 3 discusses the logical structure of the addressed system, the weather model and the working assumptions. Section 4 presents the three policies proposed to control the heating of the railroad switches. In Section 5 a detailed description of the developed stochastic modeling framework, through which the energy efficiency abilities of the switch heating policies are assessed and compared, is provided. Section 6 illustrates the case study, the measures of interest and discusses the results obtained by simulating the model. Section 2 overviews related work, and Section 7 draws conclusions, together with the identification of interesting future research lines.

2. Related work

The railway sector is increasingly in expansion, to cope with the growing global demand for transport, with estimations of more than doubling the overall activity by 2050 (as in the report [6] by the International Energy Agency, where the future of the rail system is discussed). Dependence on electricity is also increasing, spanning four major areas: train movement, auxiliary systems of the train, auxiliary systems of the infrastructure (including switch heaters, addressed in this paper), stations and other related uses [4]. Reduction of energy consumption in this sector is therefore a challenging issue for both manufacturers producing physical technologies and for providers of cyber solutions to energy supply policies. The directions to promote energy efficiency in railway systems is wide, as e.g. discussed in [7].

Up-to-now, only a few works exploit Information and Communication Technologies (ICT) to improve the energy efficiency of railway monitoring and management. For instance in [8, 9] the authors propose an architecture to collect and analyze data from the railway infrastructure to highlight possible actions for improving the energy efficiency. Both the works, however, focus on the definition of the architecture and methodology, taking for granted data availability. In [10], instead, the authors focus on the collection of data, by developing an energy-efficient MAC protocol for Wireless Sensor Networks (WSNs) that can be deployed on the railway infrastructure for data collection. However none of those works specifically focus on integrating and exploiting data from external sources (i.e, weather forecast stations) to efficiently manage specific railways components, such as the switch heaters.

Several works, instead, focused on defining policies and models to improve the energy efficiency of railways components. On the specific topic of railroad switch heating system, as addressed in this paper, in [3] the authors propose to improve the physical technology constituting the heater components. Several other studies, including those performed by a subset of the authors of this paper, focused on the analysis of energy consumption induced by supply policies in smart cities contexts (such as the railway domain, but also home/buildings), contributing an evaluation framework to assess the energy consumption under specific assumptions on the system behavior, fault model and environment conditions [11, 12, 13, 14, 15]. In general, the problem of trade-off between energy consumption and reliability/survivability of the system has been addressed, assuming an on/off strategy based on temperature thresholds. Different modeling formalisms (such as hybrid automata [16], hybrid Petri nets [17], Stochastic Activity Networks (SANs) [18]), and tools for building, evaluation and verification purposes of the switch heating control system models (such as Möbius [19] and Uppaal [20]) have been adopted in these works. Among the most recent contributions in this category, in [5] the authors advance previous studies on the analysis of the rail road switch heating control system, by developing a more accurate evaluation model that takes into account humidity and dew points as parameters affecting the probability of network communication failure. This is the modeling framework also adopted in this paper, which however conducts a more comprehensive study encompassing both architectural and analysis objectives. On the former, the paper proposes variants of the energy supply control policies, aiming at improving in energy consumption. On the latter, the modeling framework is enriched with the representation of the new proposed control policies and comparison among them is performed.

3. Logical representation of the addressed railroad system

Rail road switches are mechanical installations enabling railway trains to be guided from one track to another. They play a critical role in the operation of the railway system, since correct routing of trains strongly depends on the correct operation of such switches. In fact, in presence of malfunctions, train derailments or train collisions could occur, with expected catastrophic consequences for passengers. Major causes for incorrect rail road switch operation are cold conditions, snow and ice. In the past, it was common to have people employed by railway companies to keep the switches clear by sweeping the snow away. More recently, switches have had heaters installed in their vicinity, automatically operated to keep the temperature around the switches above freezing. The heaters may be powered by gas, water circulation, steam or electricity. The most commonly used in railroads worldwide is the electric heating, which is the one we concentrate on in this paper. In addition to the physical composition of the heater and of its placement around the switch, it is the policy adopted to switch on/off the heater that impacts on the energy consumption.

3.1. Components of the railroad switch heating control system

Having as reference the current practice in managing the switch heaters in some European countries, such as Italy, the logical view of the heating control system under study consists in a two-level organization. At higher level, there is a set of coordinators C_h (where $h \geq 1$), with each one supervising a set of controllers L_{hi} (where $i \geq 1$) deployed locally to the group of switch heaters under the management of C_h. More specifically, for each physical switch heater SH_{hi}, there is a controller L_{hi}, coordinated by C_h, close to it and devoted to switch on/off its electrical power supply. Such controllers can be implemented in practice by employing any IoT technology that enables the communication between a control logic running in the controller and the sensors and actuators installed on the switches and the heating control system. It is worth to mention that in this work we only assume that the platform provides basic communication capabilities. The size and relevance of a specific railway station in the railway backbone of a Country determine the number of components of switches and consequently the number of C_h and of its L_{hi}: in general, the bigger is the size and the higher are these numbers. For the purpose of our study, the cardinalities of the set of C_h and of the set of L_{hi} coordinated by each C_h are parameters that assume values once the specific configuration of a railway station is chosen. Each L_{hi} is connected to the proper C_h through a communication channel CH_{hi}, according to some defined topology. From a practical perspective, different communication technologies could be used. A rather common one is Power-Line Communication (PLC), which exploits the same power lines powering the heaters. PLC is the network technology adopted in this study. Note that, although in general rather well balanced, the number of local controllers is not exactly the same for each coordinator. In particular, the switches topology and configuration of the communication network can induce some differences. Although in the following we assume a balanced topology, different topologies and network differences can be easily included and handled.

The joint activities of each C_h and their respective L_{hi}, connected through the PLC communication network, perform the management of the heaters, responsible for keeping the railroad switches free from ice and snow. Details on the energy management algorithms are in Section 4.

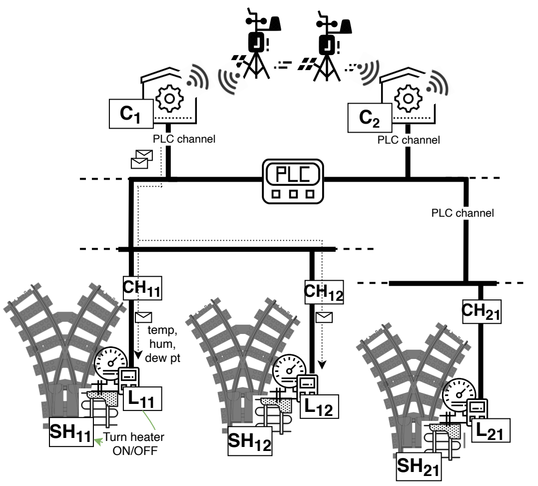

Figure 1: High level organization of the considered railroad switch heating system.

Fig. 1 depicts the described system organization.

3.2. Weather conditions

To decide when it is necessary to switch on the heating system, so to guarantee the safe operation of the switches, environmental data are considered. In particular, air temperature and humidity related values have strong impact on the formation of ice and snow and are therefore exploited by the heating logic. Typically, sensors devoted to periodically measure the air temperature are positioned close to each switch. In addition to temperature values, it is also possible to exploit humidity data. Since the installation of humidity sensors is not common nowadays in railways systems, humidity values can be retrieved by exploiting the integration with external systems, such as weather forecast stations that sample such values in larger geographical area (meaning just one value for the geographical area where the railway station is positioned).

In order to make our study of interest to railway stations where minimal weather information are available, we restrict energy management policies to exploit temperature data obtained from both weather forecast station and sensors at each switch location, while humidity related data only come from weather forecast stations. Concerning these last, more than pure humidity, it is the *dew point* that plays an impactful role in our context. It is the temperature at which some of the water vapor condense into liquid water. The higher the dew point rises, the greater is the amount of moisture in the air and, depending on the temperature value, there is the risk of ice formation. Dew point values are calculated as a function of temperature and humidity, and are usually made available by meteorological stations, so we consider it a meteo information immediately exploitable in our analysis framework.

From a notation point of view, we denote with: t^w the time interval between two consecutive observations of the weather data; $T_0(t_k)$ and $T^{dew}(t_k)$ the temperature and dew point values, respectively, provided by the pertinent weather station at time t_k, where $k \in \mathbb{N}$, $t_k = t^w \cdot k$. Moreover, $T_{hi}(t_k)$ and $H_{hi}(t_k)$ indicate temperature and humidity values, respectively, relative to the position of the heater SH_{hi}. While the former are provided by local sensors, the latter can be determined on the basis of the local temperature and the dew point from the meteorological station, according to the formulation (from [21, 22])

$$H_{hi}(t_k) = e^{\left(\frac{aT^{dew}(t_k)}{b+T^{dew}(t_k)} - \frac{aT_{hi}(t_k)}{b+T_{hi}(t_k)}\right)}, \qquad (1)$$

where $a = 17.27, b = 237.7^\circ$ C.

Although not used by the switch heating control algorithms, the values of $H_{hi}(t_k)$ are relevant for determining the failure rate of the PLC communication network, as detailed in the following when discussing the assumptions model.

Note that the values of $T_{hi}(t_k)$ and $H_{hi}(t_k)$ can be different from those provided by the meteorological station, due to specific conditions at the position of the controlled railroad switch (e.g., better/worse exposition to sun, or the presence of shadow, or others). This is the reason why switch freezing can occur when $T_{hi}(t_k) \leq T^{dew}(t_k)$ and $T_{hi}(t_k) \leq 0^oC$, if the heater is not turned on.

Weather conditions are represented at each instant of time t_k by a stochastic process composed by a $(2n+2)$-tuple of random variables:

$$\begin{aligned}\big(T_0(t_k), T^{dew}(t_k), T_{11}(t_k), \ldots, T_{1n_1}(t_k),\\ \ldots, T_{m1}(t_k), \ldots, T_{mn_m}(t_k),\\ \ldots, H_{11}(t_k), \ldots, H_{1n_1}(t_k),\\ \ldots, H_{m1}(t_k), \ldots, H_{mn_m}(t_k)\big),\end{aligned}$$

where n is the number of all L_{hi}, with $n = n_1 + \cdots + n_m$ and n_h is the number of L_{hi} associated to C_h. These random variables are piece-wise constant over time and change value every t^w units of time. Thus, for example, $T_{hi}(t) = T_{hi}(t_k)$, for each t such that $t_k \leq t < t_{k+1}$, and similarly for the other random variables.

With respect to [5], where temperature and humidity changes are synchronized with control actions to assure no switch freezing, here the interest is in addressing a wide and general variety of weather profiles. To this purpose, the following four key parameters are introduced:

$\bar{T}_0$ It indicates the temperature assigned at time 0 to the variable $T_0(0)$, which represents the expected temperature in the geographical area where the railway station is positioned over the time $[0, t]$. The value assumed by $\bar{T}_0$ determines the distribution and the mean value of $T_0(t_k)$ over the interval of time $[0, t]$.

T_0^a It is a non negative value and represents the amount of variation of $T_0(t_k)$ around its expected value. An increase of T_0^a determines an increase of the maximum and minimum values that $T_0(t_k)$ can assume, thus impacting on the variance of $T_0(t_k)$ in the interval of time $[0, t]$.

ΔT_0^{dew} It is a non negative value and indicates how much the expected value of $T^{dew}(t_k)$ is close to or far from the expected value of $T_0(t_k)$. At increasing the value of ΔT_0^{dew}, the expected value of $T^{dew}(t_k)$ decreases, being $T^{dew}(t_k) \leq T_0(t_k)$.

T_a^{dew} It is a non negative value and represents the same effect that the parameter T_0^a has on $T_0(t_k)$, but this time on $T^{dew}(t_k)$. Similarly, by increasing T_a^{dew} also the maximum and minimum values that $T^{dew}(t_k)$ can assume increase, thus impacting on the resulting variance.

This formulation is a major advancement on the representation of the weather conditions with respect to [5]. In fact, by assigning different combinations of values to the 4-tuple $(\bar{T}_0, T_0^a, \Delta T_0^{dew}, T_a^{dew})$, a variety of weather profiles can be represented: medium, low and very low temperatures ($\bar{T}_0$), low and high temperature variations around the expected value for the temperature (T_0^a), dew point values close to, or far from, the temperature (ΔT_0^{dew}), and low and high variation of the dew point values around the expected value (T_a^{dew}).

With these parameters the stochastic processes $T_0(t_k)$, $T^{dew}(t_k)$ and $T_{hi}(t_k)$ are defined by Eqs. (2), (3) and (5), respectively.

$$T_0(t_k) = \begin{cases} \frac{1}{1-\sigma_f} T_0^a f'(t_k)(t_k - t_{k-1}) + T_0(t_{k-1}) \\ \qquad\qquad \text{if } k \geq 1 \text{ and } Y = 0, \\ \frac{1-2\sigma_f}{1-\sigma_f} T_0^a f'(t_k)(t_k - t_{k-1}) + T_0(t_{k-1}) \\ \qquad\qquad \text{if } k \geq 1 \text{ and } Y = 1, \\ \bar{T}_0 \qquad\qquad \text{if } k = 0. \end{cases} \tag{2}$$

$$T^{dew}(t_k) = \begin{cases} T_v^{dew}(t_k) & \text{if } T_v^{dew}(t_k) < T_0(t_k), \\ T_0(t_k) & \text{otherwise,} \end{cases} \tag{3}$$

where

$$T_v^{dew}(t_k) = \begin{cases} \frac{1}{1-\sigma_f} T_a^{dew} f'(t_k)(t_k - t_{k-1}) \\ + T^{dew}(t_{k-1}) \quad \text{if } k \geq 1 \text{ and } Y = 0, \\ \frac{1-2\sigma_f}{1-\sigma_f} T_a^{dew} f'(t_k)(t_k - t_{k-1}) \\ + T^{dew}(t_{k-1}) \quad \text{if } k \geq 1 \text{ and } Y = 1, \\ \bar{T}_0 - \Delta T_0^{dew} \quad \text{if } k = 0. \end{cases} \tag{4}$$

$$T_{hi}(t_k) = \begin{cases} \frac{1}{1-\sigma_f} T_{hi}^a (T_0(t_k) - T_0(t_{k-1})) \\ + T_{hi}(t_{k-1}) \quad \text{if } k \geq 1 \text{ and } Y = 0, \\ \frac{1-2\sigma_f}{1-\sigma_f} T_{hi}^a (T_0(t_k) - T_0(t_{k-1})) \\ + T_{hi}(t_{k-1}) \quad \text{if } k \geq 1 \text{ and } Y = 1, \\ \bar{T}_0 + \Delta T_{hi} \quad \text{if } k = 0. \end{cases} \tag{5}$$

In Eqs. (2) and (4), $f'(t)$ is the derivative of the differentiable function f(t), that represents the qualitative trend of the temperature over one day. The distance of $T_0(t_k)$ from f(t) depends on the value of σ_f, in particular $T_0(t_k) = \bar{T}_0 f(t)$, for $\sigma_f = 0$. Y is a (discrete) random variable that can take the values 1 and 0 with probabilities $p_Y(1)$ and $1 - p_Y(1)$, respectively. Y represents the event that $T_0(t_k)$ changes direction ($Y = 1$), flipping the slope, with respect to $f(t_k)$, or not ($Y = 0$). If the expected value of Y is 0.5 then the expected value of $T_0(t_k)$ and $T_v^{dew}(t_k)$ is governed by f(t).

In Eq. (5), $T_{hi}(t_k)$ follows the same trend of the actual $T_0(t_k)$, although it can assume different values. Assigning small values to ΔT_{hi} and T_{hi}^a (e.g., 3 and 1, respectively), small variations of $T_{hi}(t_k)$ respect to $T_0(t_k)$ can be represented.

It can be observed that in Eqs. (3) and (4), as the dew point is always lower than or equal to the temperature of the air, i.e., $T_v^{dew}(t_k) \leq T_0(t_k)$, a decrease in the values of ΔT_0^{dew} or T_a^{dew} causes $T_v^{dew}(t_k)$ to tend towards $T_0(t_k)$. This increases the risk of switch freezing, that occurs when $T_{hi}(t_k) \leq T^{dew}(t_k)$, $T_{hi}(t_k) \leq 0^oC$ and the heater is switched off. Intuitively, T_a^{dew} and T_{hi}^a represent the amplitude of the displacement of $T^{dew}(t_k)$ and $T_{hi}(t_k)$, respectively, whereas ΔT_0^{dew} and ΔT_{hi} represent a shift in the displacement of $T^{dew}(t_k)$ and $T_{hi}(t_k)$, respectively. Notice that in Eq. (5) $f(t)$ is replaced by $T_0(t_k)$ so that $T_{hi}(t_k)$ depends on the actual air temperature. Examples for $T_0(t_k)$, $T^{dew}(t_k)$ and $T_{hi}(t_k)$ are depicted in Fig. 5.

3.3. Working assumptions

The development of the heaters management algorithms, and the evaluation framework set up to assess their efficacy, are based on the following assumptions:

- The PLC communication network can experience failures, while all the other components always work correctly. The failure on a communication channel CH_{hi} between the coordinator C_h and its local controller L_{hi} can have impact also on other local controllers, depending on the topology of the network (e.g., in the case the coordinator and its local controllers are connected in series through the same power line). A failed communication channel prevents the exchange of information between the source and the destination; once the channel is repaired, the communication is correctly restored. Communication channel failures depend on weather conditions (both temperature and humidity) and other characteristics of the PLC. A repair operation brings back the failed communication channel to operate correctly. Introducing failure events gradually is helpful to understand the impact of individual phenomena on both the complexity and the efficiency of the control system. Accounting for the failure of other components is certainly relevant and postponed as future work.
- It is assumed that each coordinator C_h has periodic access to weather information (temperature and dew point) from the meteo station and the temperatures computed by the sensors local to switches under their control (and conveyed to them by the pertinent L_{hi} through the PLC network). Each L_{hi} has directly available the temperature values, periodically sampled by the local sensor. Depending on the energy management policy, other weather information can be gathered from C_h (specifically, in one of the policies developed in this paper, L_{hi} receives the dew point from the coordinator).
- Communications along working PLC channels occur instantaneously. This assumption is useful to simplify the modeling and analysis of the energy management policies, since accounting for a communication is not expected to have an impact on the energy consumption that is the goal of the analysis.
- If $T_{hi}(t_k) \leq T^{dew}(t_k)$ and $T_{hi}(t_k) \leq 0$, i.e., there are chances of ice formation, then turning SH_{hi} off causes an instantaneous freezing of the switch (pessimistic assumption). Similarly, if the switch is frozen, turning on SH_{hi} eliminates instantaneously the ice from the switch.
- A heater consumes electrical energy only when it is switched on.

4. Energy management policies

Efficient energy management policies should leverage heterogeneous information like environmental data in order to optimize the process of turning on and off the heaters, so they are activated only when it is necessary. As also pointed out when presenting the working assumptions, each L_{hi} can access only the data collected by a local temperature sensor. The C_h components, instead, are integrated with other systems and can benefit from the availability of additional environmental data (e.g., the value of $T^{dew}(t_k)$), which could be downloaded from an external system through the Internet. Such knowledge can be profitably exploited, in addition to temperature data coming from the local sensors, by the coordinators C_h towards a more energy efficient and resilient management of the switch heating system.

In view of such wider knowledge, energy efficiency is primarily boosted by the C_h components, and the logic of the heating policy they exploit, then actuated through the local heater controls, is presented in Section 4.1. However, its feasibility requires correct operation of the communication between C_h and L_{hi}. In this paper, communication channels are based on the PLC technology, for which the probability of experiencing failures cannot be considered negligible, thus leaving the involved L_{hi} without knowing when to switch on or off the controlled heater(s). So, it is necessary to equip local controllers with a policy to manage the switch heater under their control, in absence of a well working connection with the coordinator.

This leads to the following overall organization of the switch heating system. Whenever the communication channel works properly, each local controller implements the command issued by the coordinator to turn the heater on or off. Instead, for the period in which the communication channel is interrupted, the local controller, disconnected from the coordinator, implements an autonomous policy. The policy that assumes a well working communication channel between L_{hi} and its coordinator C_h is first presented in Section 4.1. Then, three policies for the case of interrupted communication channel are introduced: a basic one, which reflects the current practice, and two variants, developed to improve the energy consumption.

4.1. Policy under functional communication channels

The policy developed for the case where the channel works properly is based on a reference temperature threshold T_{thr} (in line with [4]). It is derived from a previous work of a subset of the authors of this paper [5], and extends current practice by exploiting external information about weather conditions that include also the dew point, in addition to the data collected by the local sensors.

At every period t_k, each L_{hi} transmits to its associated C_h the value of the temperature obtained from its local sensor, and receives from C_h the command to turn on or off the controlled heater. In particular, using the temperature and dew point values received from the meteo station, and the local temperatures received from the associated L_{hi}, each C_h determines the on or off command by checking the following conditions:

- if $T_{hi}(t_k) \leq T^{dew}(t_k) + T_{\text{thr}}$ and $T_{hi}(t_k) \leq T_{\text{thr}}$, the command to L_{hj} is to turn on the heater;
- if $T_{hi}(t_k) > T^{dew}(t_k) + T_{\text{thr}}$ or $T_{hi}(t_k) > T_{\text{thr}}$, the command L_{hj} is to turn off the heater.

The value of T_{thr} is provided by the railway operator on the basis of the geographical area of the location, with possible adjustments based on the analysis of historical weather data, in order to minimize the probability of a temperature drop from T_{thr} to zero in less than the time interval between two control actions t^w.

4.2. Policies under interrupted communication channel

In case of channel failure, each L_{hi} that experiences isolation from its coordinator C_h autonomously takes decision to turn on or off the heater.

In the following we describe three different, locally operated, energy management policies. The definition of sophisticated logics to compensate for unavailability of the rich weather information, as possessed by the coordinator components, is a major novel contribution of this paper.

4.2.1. Basic policy

According to a first basic policy, named P_{bas}, each L_{hi} isolated from the respective C_h, uses the temperature threshold T_{thr} to decide to turn on or off the heater. The operated policy is similar to that of the coordinator, but the check is only limited to the temperature obtained from the local sensor, since there is no availability of the dew point value.

Specifically, the logic of this basic policy is: L_{hi} decides to turn on the heater SH_{hi} if $T_{hi}(t_k) \leq T_{\text{thr}}$; otherwise, if $T_{hi}(t_k) > T_{\text{thr}}$, then SH_{hi} is turned off.

4.2.2. Memory-based policy

A more complex policy with respect to P_{bas} is P_{mem}, a policy that requires every L_{hi} to be able to store some historical environmental data on the local memory, to exploit when the communication between C_h and L_{hi} is interrupted.

When adopting the P_{mem} policy, L_{hi} stores in its local memory the value of T^{dew} that it receives periodically from its C_h. In case of channel failure,

L_{hi} uses the last stored value of T^{dew} for at most $\Delta m \cdot t^w$ time units. After $\Delta m \cdot t^w$ time units, L_{hi} switches to the P_{bas} policy, since the T^{dew} value could be outdated and inaccurate with respect to the current weather conditions.

In detail, P_{mem} performs the following steps.

Let t' be the channel failure time and T^{dew} (t_m) be the last value of T^{dew} before the failure, i.e., $t_m = m \cdot t^w$ such that $t_m \le t' < t_{m+1}$. During the channel failure period, at the time t_k, with $t_k \ge t'$, L_{hi} performs the following decisions. If $t_m \le t_k \le t_{m+\Delta m}$, then L_{hi} adopts the same policy used when the channel is working as described in Section 4.1, with $T^{dew}(t_k) = T^{dew}(t_m)$. Instead, as soon as $t_k > t_{m+\Delta m}$, L_{hi} switches to the policy P_{bas}. Formally:

- if $t_m \le t_k \le t_{m+\Delta m}$ then:
  - turn on if $T_{hi}(t_k) \le T^{dew}(t_m) + T_{\text{thr}}$ and $T_{hi}(t_k) \le T_{\text{thr}}$;
  - turn off if $T_{hi}(t_k) > T^{dew}(t_m) + T_{\text{thr}}$;
- if $t_k > t_{m+\Delta m}$ then:
  - turn on if $T_{hi}(t_k) < T_{\text{thr}}$;
  - turn off if $T_{hi}(t_k) \ge T_{\text{thr}}$.

In a geographical location characterized by slowly changing weather conditions, it is expected that this policy, by using the dew point information in the same manner as the coordinator, brings significant benefits in terms of energy consumption.

4.2.3. Prediction-based policy

The idea at the basis of P_{pre} is that L_{hi} maintains in local memory recent values of T_{hi} and uses them, possibly together with other information on seasonal trends, to perform local prediction of future temperature values. Consequently, when the channel fails, L_{hi} exploits its temperature prediction to decide on whether to turn on or off the heater, instead of comparing the actual temperature with a fixed temperature threshold, as the basic policy P_{bas} does. Improvements in energy consumption with respect to P_{bas} are expected when the prediction shows a future temperature that is increasing; this can allow the heater to be kept off, while P_{bas} would have decided to turn it on. Of course, the accuracy degree of the prediction plays a fundamental role in avoiding the switch to freeze.

In particular, calling $\tilde{T}_{hi}(t_{k+1})$ the predicted temperature at time t_{k+1}, L_{hi} checks:

- if $\tilde{T}_{hi}(t_{k+1}) \le \tilde{T}_{\text{thr}}$, then turn on;
- otherwise, turn off.

The value of $\tilde{T}_{\text{thr}}$ can be decided on the basis of the absolute error ϵ_{f} of the forecasted value for $\tilde{T}_{hi}(t_{k+1})$, for which $\tilde{T}_{hi}(t_{k+1}) - \epsilon_{\text{f}} \le T_{hi}(t_{k+1}) \le \tilde{T}_{hi}(t_{k+1}) + \epsilon_{\text{f}}$, with $\epsilon_{\text{f}} \ge 0$.

This means that, considering $\tilde{T}_{\text{thr}} \ge \epsilon_{\text{f}}$ avoids the freezing of the rail switch, because the switch is turned off when $\tilde{T}_{hi}(t_{k+1}) > \tilde{T}_{\text{thr}} \ge \epsilon_{\text{f}}$, i.e., when $\tilde{T}_{hi}(t_{k+1}) - \epsilon_{\text{f}} > 0$. In fact, being $T_{hi}(t_{k+1}) > \tilde{T}_{hi}(t_{k+1}) - \epsilon_{\text{f}}$, then also the actual temperature $T_{hi}(t_{k+1})$ is greater than 0. With $\tilde{T}_{\text{thr}} = \epsilon_{\text{f}}$, the switch is turned on when $\tilde{T}_{hi}(t_{k+1}) \le \tilde{T}_{\text{thr}} = \epsilon_{\text{f}}$, thus leading to unneeded energy supply if $\tilde{T}_{hi}(t_{k+1}) > -\epsilon_{\text{f}}$, such that $T_{hi}(t_{k+1}) > 0$. Unneeded energy supply occurs always for $\tilde{T}_{\text{thr}} > \epsilon_{\text{f}}$ when $\epsilon_{\text{f}} < \tilde{T}_{hi}(t_{k+1}) < \tilde{T}_{\text{thr}}$, because $T_{hi}(t_{k+1}) \ge \tilde{T}_{hi}(t_{k+1}) - \epsilon_{\text{f}} > 0$.

Differently from the previous setting, considering $-\epsilon_{\text{f}} \le \tilde{T}_{\text{thr}} \le \epsilon_{\text{f}}$ can lead to both freezing of the rail switch and unneeded energy supply. Freezing of the rail switch occurs when the switch is turned off in the case that $-\epsilon_{\text{f}} \le \tilde{T}_{\text{thr}} < \tilde{T}_{hi}(t_{k+1}) \le \epsilon_{\text{f}}$ such that $T_{hi}(t_{k+1}) \le 0$. Unneeded energy supply occurs when the switch is turned on in the case that $-\epsilon_{\text{f}} < \tilde{T}_{hi}(t_{k+1}) \le \tilde{T}_{\text{thr}} \le \epsilon_{\text{f}}$ such that $T_{hi}(t_{k+1}) > 0$.

Finally, considering $\tilde{T}_{\text{thr}} < -\epsilon_{\text{f}}$ leads to the freezing of the rail switch each time that $\tilde{T}_{\text{thr}} < \tilde{T}_{hi}(t_{k+1}) \le -\epsilon_{\text{f}}$, because $T_{hi}(t_{k+1}) < \tilde{T}_{hi}(t_{k+1}) + \epsilon_{\text{f}} < 0$. With $\tilde{T}_{\text{thr}} < -\epsilon_{\text{f}}$ unneeded energy supply cannot occur, because when $\tilde{T}_{hi}(t_{k+1}) \le \tilde{T}_{\text{thr}} < -\epsilon_{\text{f}}$ then $T_{hi}(t_{k+1}) \le \tilde{T}_{\text{thr}} + \epsilon_{\text{f}} < 0$.

From these considerations, the best value of $\tilde{T}_{\text{thr}}$ for avoiding the switch freezing and reducing the unneeded energy supply is ϵ_{f}. The value of the error ϵ_{f} can depend on both the adopted forecast strategy and the different weather profiles.

The literature on weather forecast is vast, so we focused on relatively recent developments based on Exponential Smoothing [23]. In particular, temperature data are collected in the time scale of minutes, so a seasonality of one day has been considered. The standard Holt-Winters approach is not adequate because the seasonality originates too many samples (e.g., 144, sampling temperature every 10 minutes for 24 hours), so we rely on a more elaborate prediction algorithm [24], and exploit its implementation in R called `tbats`[1].

[1] `https://www.rdocumentation.org/packages/forecast/versions/8.11/topics/tbats`.

5. Modeling the energy management policies

In order to assess the efficiency of the developed algorithms for managing the heating of railroad switches, a stochastic model-based approach is adopted. The model is expressed through the SAN [18] formalism and evaluates the measures of interest through the simulation engine of the Möbius tool [19]. SAN is a stochastic extension of the Petri Nets formalism, based on the following primitives: plain and extended places (blue and orange circles, respectively), timed and instantaneous activities (hollow and solid vertical bars, respectively), input and output gates (triangles pointing left or right, respectively). The SAN primitives are defined by expressions or statements of the programming language C++. This formalism is widely adopted in dependability and performability indicators and is very suitable to model the heating management system tackled in this work.

5.1. Stochastic process

The energy consumed by each switch heater SH_{hi} is determined by its state at each time instant t, represented by the stochastic process $\{X_{hi}(t) \,|\, t \geq 0\}$ defined by

$$X_{hi}(t) = \begin{cases} 1 & \text{if } SH_{hi} \text{ is } \textit{on} \text{ at time } t, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

The time to the physical fault of a communication channel CH_{hi} is modeled as a random variable exponentially distributed with rate:

$$\lambda_{hi}(t_k) = c \cdot w^{H_{hi}(t_k)} \tag{7}$$

where w represents the weight, i.e. the impact, over time of the humidity $H_{hi}(t_k)$ on the fault rate, and the constant c represents the impact of all the other influencing aspects (including the characteristics of the power lines, e.g. the distance between the local controllers Ls and the coordinators Cs, the air temperature and others) on the fault rate. Depending on the topology of the specific PLC communication network at hand, the failure of CH_{hi} can propagate to other channels connecting other local controllers to C_h. This propagation path is accounted for in the analysis by building a topology of interactions, as performed in [5]. The recovery time of CH_{hi} is a uniform random variable with mean μ_{hi}.

Every t^w units of time two things happen simultaneously, and in this order: all C_h and L_{hi} perform their actions, and the weather model is updated. Thus, controller actions are based on the previous value of temperature and dew point. If SH_{hi} is set to *off* at time t_k then, even if after weather change it would be needed to set it to *on*, the next change to turn it on is at t_{k+1}, possibly leading to freezing of the switch. Exploiting simultaneous events with priorities allows to model switch freezing (an important addition with respect to [5]) without resorting to an over complicated model.

5.2. Overview of the SAN model

The system under analysis is modeled and evaluated following the *DARep* compositional and modular approach, as proposed in [25] and also detailed in [5], by means of the tool Möbius [19].

Two atomic template SAN models *TM_W* and *TM_L*, depicted in Figs. 2 and 3, are defined to represent, respectively: i) the global weather conditions, and ii) the components CH_{hi} and L_{hi}, and the weather conditions local to each switch. The algorithms of the three switch heating policies are represented in the single template model *TM_L* through C++ objects defined in the SAN primitives as different specializations of C++ template classes and instantiated at compilation time of the simulator solver.

The models *TM_W* and *TM_L* are a significant extension and update of the corresponding models proposed in [5]: the structure and all the activities, excluded *CF* and *CR*, i.e., 4 activities over 6 and 12 gates over 12, are new or considerably modified, and 8 new places over 17 are defined (shared places are counted one time only).

The overall system model is obtained generating and composing automatically through the $\mathcal{D}$ operator, supported by the *DARep* approach, one instance TM_W_1 of the template model *TM_W* and n instances of the template model *TM_L*, i.e., one instance TM_L_{hi} for each L_{hi}.

TM_W_1 represents, through the timed activity *TWU* and the linked primitives (places and input and output gates): i) the changes of the global weather conditions at each instant of time t_k, ii) the trigger for the changes of the weather conditions local to each switch, iii) the transmission of the current weather conditions from the weather forecast service to the coordinators.

Each instance TM_L_{hi} represents: i) the failure and repair of CH_{hi}, through the timed activities *CF* and *CR*, and the failure propagation based on the dependency-aware State Variable (SV) NW (the description is omitted being equal to that proposed in [5], where the detailed description can be found), where the instance NW_{hi} represents the number of channels that propagate the effects of their fault to the channel CH_{hi}. ii) the actions of L_{hi} turning on and off SH_{hi} at each instant of time t_k, through the input gate of the instantaneous activity tTU, iii) the updating of the weather conditions local to each switch, i.e., $T_{hi}(t_k)$ and $H_{hi}(t_k)$, through the output gates of the instantaneous activity tTU. The topology associated to NW is based on the PLC communication network topology.

In *TM_W*, the activity TWU is always enabled with deterministic completion time equal to t^w. Let Y^0 and Y^d be the random variables Y associated to $T_0(t_k)$ and $T^{dew}(t_k)$, respectively. At the k-th completion, i.e., at time t_k, TWU chooses one of the output gates *Var00*, *Var01*, *Var10* and *Var11* with the probability associated to one of these random events $(Y^0 = 0, Y^d = 0)$, $(Y^0 = 0, Y^d = 1)$, $(Y^0 = 1, Y^d = 0)$ and $(Y^0 = 1, Y^d = 1)$, respectively.

Next, TWU performs the code of the input gate $toTUtw$ that: i) adds one token to the place k, representing the number k of current weather updates, and ii) triggers the immediate activity tTU of each *TM_L*$_{hi}$, by assigning one token to each instance TU_{hi} of the dependency-aware SV TU. The topology associated to TU is a one-to-many topology, which links the single instance of *TM_W* to the n instances of *TM_L*. Using this topology, *DARep* defines in *TM_W*$_1$ one different instance TU_{hi} for each *TM_L*$_{hi}$, and in each *TM_L*$_{hi}$ it defines the same instance TU_{hi} already defined for *TM_W*$_1$.

Next, TWU performs the code of the chosen output gate that updates the local extended places *T0*, and the extended places *DP*, *T0d* and *DPp* (shared among all the instances of the template models), which represent $T_0(t_k)$, $T^{dew}(t_k)$, $T_0(t_k) - T_0(t_{k-1})$ and $T^{dew}(t_{k-1})$, respectively, as defined in Eqs. (2) and (3).

In each instance *TM_L*$_{hi}$, the local place *HOn* represents the status of the heater, such that if $HOn = 0$ the heater is turned off. When $TU_{hi} = 1$ the immediate activity tTU is enabled and completes. At completion, tTU chooses the output gate *Var0* or *Var1*, based on the random events $Y^{hi} = 0$ and $Y^{hi} = 1$, where Y^{hi} is the random variable associated to $T_{hi}(t_k)$. Next, the control action for the policy selected at compilation time is performed by the code in the output gate $isTU$, based on the current, not yet updated, weather conditions local to each switch. Next, using the places *T0d*, *DP* and *DPp*, tTU performs the code of the selected output gate that updates the values of the local places *T* and *H*, representing $T_{hi}(t_k)$ and $H_{hi}(t_k)$, respectively, as defined in Eq. (5).

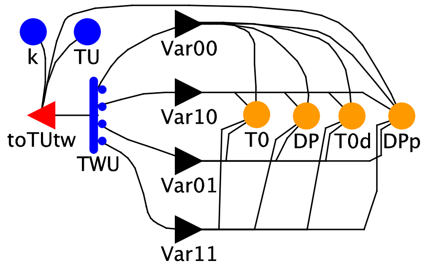

Figure 2: SAN template model *TM_W* for the global weather conditions.

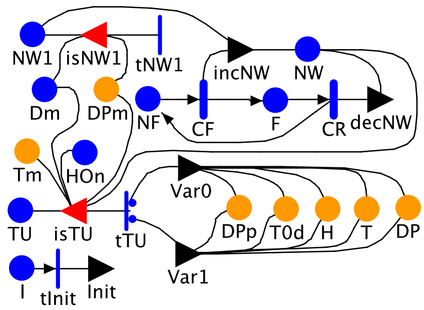

Figure 3: SAN template model *TM_L* for CH_{hi}, L_{hi} and the weather conditions local to each switch.

The extended place *Tm*, initialized by the output gate *Init* and updated by the input gate $isTU$, maintains all the values of $T_0(t_k)$ used by the policy P_{pre} to perform local prediction of future temperature values.

The places *Dm* and *DPm* represent Δm and $T^{dew}(t_m)$, respectively, used and updated in the input gate $isTU$ for the policy P_{mem}. These places are reset in the input gate *isNW1* at each firing of the activity *tNW1*. *tNW1* is enabled at each completion of *CF*, in the same or a different instance of *TM_L*, by the gate *incNW*, which sets to one each *NW1*$_{hi}$ of the dependency-aware SV *NW1* if the current value of NW_{hi} is 0. The topology associated to *NW1* is the same of NW.

To reinforce the correctness of the implementation, the behavior of the model has been extensively tested, adopting parameter settings corresponding to a number of simplified scenarios for which the same outcome of the model can be also manually computed.

Finally, we point out a few basic aspects characterizing this modeling effort:

- the model allows to account for a variety of communication network topologies, since the description of such topology is provided in input to the model. In the scenarios analyzed, both independent and correlated communication channel failures are addressed;

- the model can be easily extended to represent other different switch heating policies, in addition to those presented in Section 4, provided the basic logic of on/off and the assumptions listed in Section 3.3 are maintained.

6. Comparison of the energy management policies

6.1. Case study

The railway station configuration adopted as case study to demonstrate the efficiency of the proposed switch heating policies is that of a typical medium-low size station (actually, inspired by the North Italy Lecco-Maggianico railway station). Its logical structure is shown in Fig. 4.

In particular, it is composed by $n = 19$ railroad switches, partitioned in two groups denoted as *North* and *South* switches, with size 9 and 10, respectively. Each switch has a local controller to manage the associated heater, and there are two coordinators, C_{North} and C_{South} to supervise the operation of the local controllers associated with the *North* and *South* switches group, respectively.

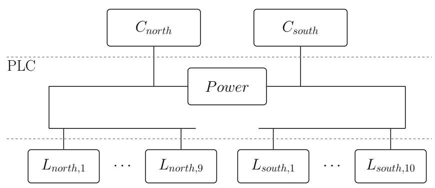

Figure 4: Logical architecture of the PLC of the railway station adopted as case study.

Although rather simple, this case study is sufficiently rich to the purpose of the conducted analysis and to appreciate the benefit of devising more sophisticated energy management policies that exploit available information than just relying on basic actions.

6.2. Assessed indicators

As already said, the efficiency of the developed switch heating policies is addressed in this paper in terms of the energy consumption they result in.

Let P^{SH} be the electrical power required by each switched-on heater SH_{hi}, i.e., when $X_{hi}(t) = 1$. The electrical energy consumed in the interval of time $[0, t]$ by all the heaters is

$$E(t) = \mathbb{E}\Big[P^{SH} \sum_{i=1}^{n} \int_0^t X_{hi}(t)\, dt\Big]. \qquad (8)$$

From the dependability perspective, unavailability has been assessed in the analyzed scenarios, to understand whether, and to which extent, the gain in lower consumption is obtained at the cost of higher unavailability of the railroad switches due to the freezing condition. This is especially the case of the prediction-based policy P_{pre}, which can tune internal parameters so to either favor energy consumption at the expense of switch availability, or vice versa, or to find a satisfactory trade-off between the two.

The unavailability indicator $\mu_U(t)$ has been defined as the expected cumulative time of freezing for each switch belonging to the system configuration that experiences a freezing in the time interval $[0, t]$. Its formulation is:

$$\mu_U(t) = \mathbb{E}\Big[\sum_{i=1}^{n} \int_0^t U_{hi}(t)\, dt\Big], \qquad (9)$$

where the random variable U_{hi} models the presence/absence of ice on SH_{hi}: at a given time instant t, $U_{hi}(t) = 1$ if SH_{hi} is frozen, $U_{hi}(t) = 0$ otherwise.

6.3. Analysis settings

The time window of the analysis is 24 hours, divided in 1440 minutes. Several weather profiles are considered, each characterized by the 4 parameters $\bar{T}_0$, T_0^a, ΔT_0^{dew} and T_a^{dew}, spanning from mild to harsh weather, and taking also into account different climate zones. By varying the combination of values assigned to the four weather-related parameters, 16 weather profile scenarios have been analyzed. The weather profiles are selected on purpose among those more representative of rather cold geographical zones, since in such conditions the heating is more intensive, so representing the best context where to conduct efficiency analysis of the proposed heating policies.

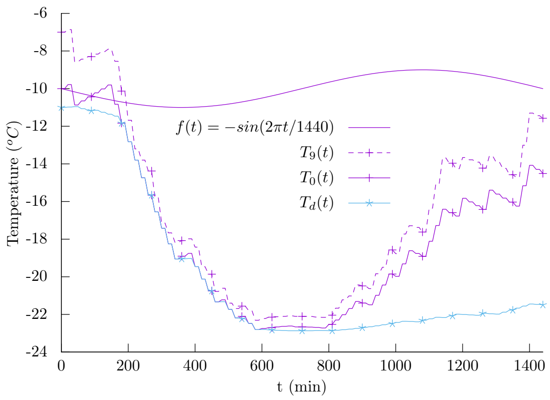

(a) $\Delta T_0^{dew} = 1°, T_a^{dew} = 1°$

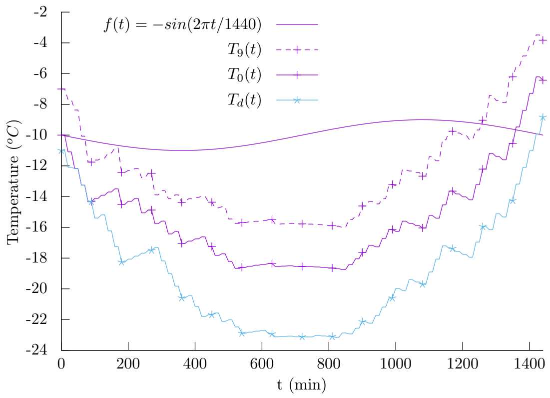

(b) $\Delta T_0^{dew} = 1°, T_a^{dew} = 10°$

Figure 5: Traces of air temperature $T_0(t_k)$, temperature $T_{\text{north},9}(t_k)$, dew point $T^{dew}(t_k)$ and $f(t)$ in two scenarios for $\bar{T}_0 = -10°, T_0^a = 10°$.

More in detail, $\bar{T}_0$, which represents the value of the temperature of the considered geographical area at the time the analysis starts, can assume values in the set $\{0°, -10°\}$. T_0^a, which represents the maximum variability of the temperature with respect to its initial value (both as an increase or a decrease of $\bar{T}_0$), can assume values in the set $\{1°, 10°\}$. Specularly, the other two parameters related to the dew point, ΔT_0^{dew} and T_a^{dew}, can assume values in the set $\{1°, 10°\}$.

The heating control system is activated every $t^w = 10$ minutes. The combination $T_{hi}^a = 0.9$ and $\Delta T_{hi} = -3$ has been selected for L_{hi} under C_{North}, and $T_{hi}^a = 1.1$ and $\Delta T_{hi} = 3$ has been selected for L_{hi} under C_{South}. This parameters setting lets the temperature in the south part of the station having a greater displacement amplitude with respect to the one in the north part of the station, and an upwards shift in the values. In this way, the impact of sun exposition can be taken into account.

The variability is set to $\sigma_f = 0.6$ for $T_0(t_k)$, $\sigma_f = 0.3$ for $T_{hi}(t_k)$ and $p_Y(1) = 0.5$. The function $f(t)$ has been selected equal to $-\sin\left(\frac{2\pi t}{1440 \text{ minutes}}\right)$, so that $T_0(t_k)$ and $T_{hi}(t_k)$ are expected to decrease between midnight and 6 AM, increase between 6 AM and 6 PM, and eventually decrease till midnight. The function $f(t)$ is depicted in Fig. 5 together with a couple of traces of $T_0(t_k)$, $T^{dew}(t_k)$ and $T_{\text{north},9}(t_k)$ obtained in different scenarios.

Concerning the topology of the PLC communication network, in each of the two groups the local controllers are connected in series through a single communication network. This means that the failure of a communication channel between a coordinator C_h and one of its associated local controllers L_{hi} interrupts also the communication with all the subsequent local controllers.

For policy-specific parameters: i) the temperature threshold T_{thr} assumes values 0° and 5° (this last value was used in previous studies [5]); ii) Δm, the number of time intervals P_{mem} holds in memory the dew point, is in the range $[0, 20]$, and iii) the predicted temperature $\tilde{T}_{\text{thr}}$ is in the range $[-3, 3]$.

Regarding the other parameters, introduced when describing the model in Section 5, they assume fixed values: $c = 3.472 \cdot 10^{-4}\,\text{min}^{-1}$, i.e., once every two days on average, $w = 5$ and the mean time to recovery of the communication channel is $\mu_{hi} = 1.66 \cdot 10^{-2}$ minutes^{-1}, i.e., once every hour on average. Although not relative to a specific context, these values seem reasonable ones and in line with the setting assumed in previous studies.

6.4. Results

The conducted analysis spans a rather wide range of scenarios, to help appreciating the behavior of the different switch heating management strategies in a variety of climate conditions. The measures specified in Section 6.2 are evaluated, by simulating the model described in Section 5.2. 10^4 simulation batches are employed to obtain the confidence intervals reported in the figures, with a 95% confidence level.

The first part of the evaluation focuses on the heating management policies individually. The purpose is to understand which are appropriate values to assign to parameters characterizing each individual policy, such that a good trade-off is reached between the two conflicting indicators under analysis (energy consumption and unavailability).

Table 1: Legend for curves in Figs. 6 and 7. In particular, μ_U and μ_E appear in both figures.

μ_U	—+—	—■—	—□—	—△—
μ_E	- + -	- ■ -	- □ -	- △ -
95% confidence				
T_0	-10	-10	-10	-10
T_0^a	1	1	10	10
ΔT_0^{dew}	1	10	1	10
T_a^{dew}	1	1	10	10

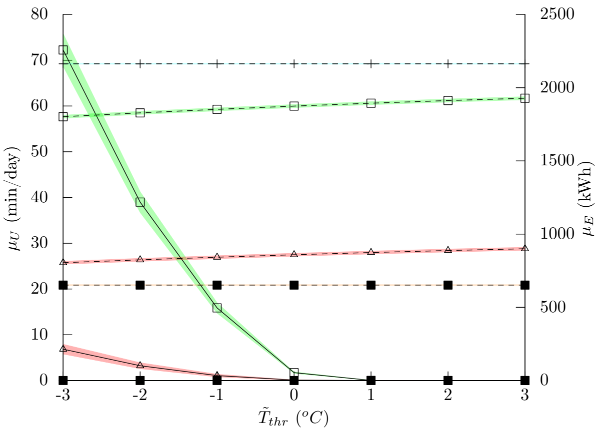

Figure 6: $P_{\rm pre}$: μ_U and μ_E at varying of $\tilde{T}_{\rm thr}$.

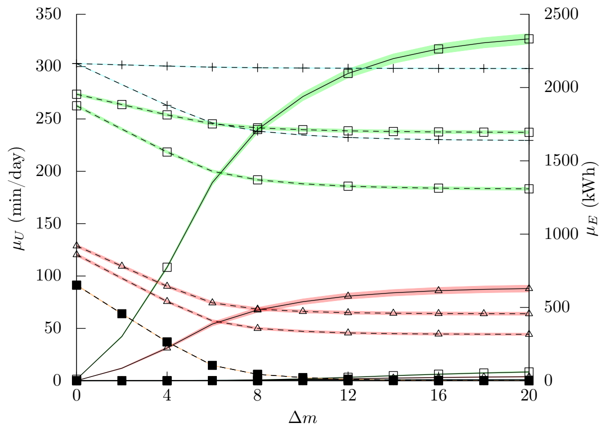

Figure 7: $P_{\rm mem}$: μ_U and μ_E at varying of Δm for $T_{\rm thr} = 0$ (a symbol every 4 points) and $T_{\rm thr} = 5$ (a symbol every 2 points).

Remember that the three policies $P_{\rm bas}$, $P_{\rm mem}$ and $P_{\rm pre}$ are executed by a local controller when the communication channel with its coordinator is interrupted, while in case of well working channel it is the coordinator that decides to set on or off the heating, always according to the logic described in Section 4.1. With the parameter setting adopted in this paper, it turned out that in 24 hours, that is the reference time interval for the analysis, the cumulated time during which at least one communication channel is not working accounted to an average of 8 hours, and this is the time during which one of the three policies $P_{\rm bas}$, $P_{\rm pre}$ or $P_{\rm mem}$ is selected and executed by at least one local controller. To keep the notation simple, in the following we refer to $P_{\rm bas}$, $P_{\rm pre}$ or $P_{\rm mem}$ to indicate the overall railroad switch heating policy that adopts these solutions, respectively, in case of communication channel failure.

Concerning the basic policy $P_{\rm bas}$, its only parameter is the temperature threshold $T_{\rm thr}$, for which we have selected the two extreme values 0 and 5. Clearly, $P_{\rm bas}$ shows the highest energy consumption when $T_{\rm thr} = 5$, and the lowest when $T_{\rm thr} = 0$. The results for $P_{\rm pre}$ and $P_{\rm mem}$ are depicted in Figs. 6 and 7, respectively, when weather conditions are as reported in Table 1. Values and shaded confidence intervals related to both μ_U (left y-axis) and μ_E (right y-axes), as cumulative unavailability time and energy consumption in one day, respectively, are reported.

In Fig. 6, variations in the range $[-3, 3]$ of the policy parameter $\tilde{T}_{\rm thr}$, representing the predicted temperature error, have very slight impact on the plots relative to μ_E. Instead, the trend of the curves related to μ_U suggests to assign values greater than 0 to the parameter $\tilde{T}_{\rm thr}$ to make unavailability of railroad switches negligible. Actually, we face an easy-to-decide situation in this analysis: selecting $\tilde{T}_{\rm thr}$ equal to 0 or 1 saves from experiencing unavailability at almost negligible increase in energy consumption.

Concerning Fig. 7, we recall that the policy $P_{\rm mem}$ is characterized by two parameters: Δm, which is the multiplying factor determining the time interval the policy uses the dew point received from the coordinator, and $T_{\rm thr}$, which is the threshold temperature to adopt when the dew point becomes too old. Δm is the varying parameter on the x-axis, while the same two values of $T_{\rm thr}$ used by $P_{\rm bas}$, i.e., 0°C and 5°C, are adopted. Therefore, when different from zero, two curves are plotted for each weather scenario, represented by a slightly different line style (as indicated in the caption of the figure). As expected, at increasing of $T_{\rm thr}$, the value of μ_E increases and the value of μ_U decreases, both tending to reach an horizontal asymptote. Concerning μ_U, only few curves (relative to $T_{\rm thr} = 0$) appear,

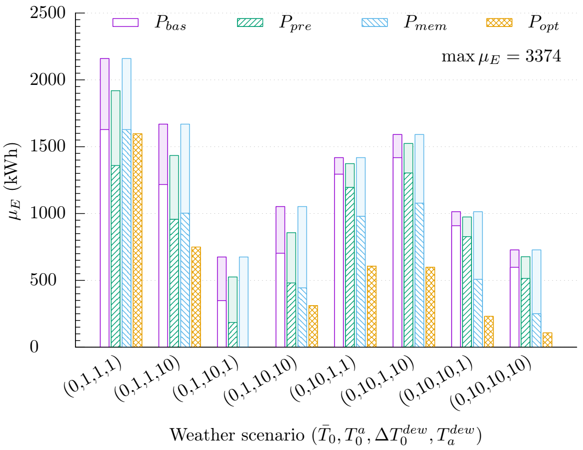

(a) $\bar{T}_0 = 0$

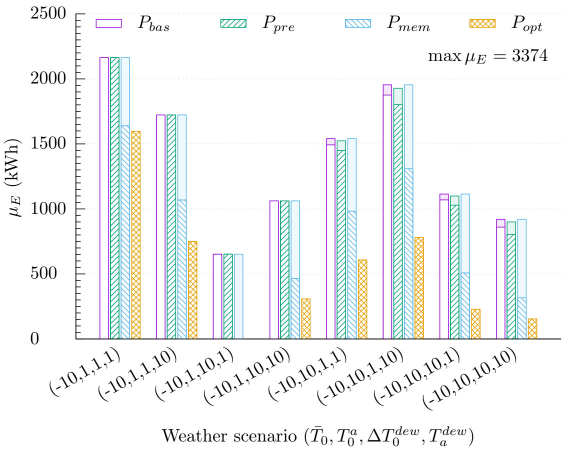

(b) $\bar{T}_0 = -10$

Figure 8: Energy consumption for the different policies when (a) $\bar{T}_0 = 0$ and (b) $\bar{T}_0 = -10$.

since the others have value 0.

An important observation is that, even though μ_{hi} (the mean time to recovery of CH_{hi}) is about 1 hour in the adopted setting, μ_U reaches its asymptote for values of Δm around 8 or 20 – depending on the considered weather profile – and not around 6 as it would be apparently expected. Actually, this is due to the impact of the failure effects induced by the assumed PLC communication topology. Moreover, note that this policy leads to rather higher level of unavailability with respect to P_{pre}: here, μ_U can overcome 300 minutes, while the maximum reported in Fig. 6 is around 70 minutes.

The second part of the analysis focuses on policies comparison. A set of 16 weather scenarios were analyzed and, for the sake of readability, separate figures for energy consumption and unavailability are shown. Results of the energy consumption indicator μ_E are shown in Fig. 8. For each scenario, four bars are included: three of them correspond to the three policies, as indicated in the legend. Analysis similar to those shown in Fig. 6 and Fig. 7 for policies P_{pre} and P_{mem}, respectively, as well as for P_{bas} with T_{thr} in the range $[0°, 5°]$, have been performed for each climate profile.

In Fig. 8, for each scenario, four bars are included: three of them correspond to the three policies, as indicated in the legend. The fourth bar (the most right one in each group) indicates the optimal energy consumption, defined as the minimum energy necessary in each scenario to satisfy the constraint of full availability of all the railroad switches (that is, what a "golden" policy, named in the figure as P_{opt}, would do, by keeping the heaters on only when needed to prevent freezing). Moreover, to keep the comparison fair, for each of the three studied policies, the minimum and maximum energy consumption in each weather scenario are reported (clearly indicated on the same bar).

The maximum amount of energy consumption ($\max \mu_E$), obtained by keeping the switch heaters on for the entire day, is also reported in the figure. The 95% confidence intervals are omitted, being negligible.

As a first general comment, the energy consumption implied by the three policies is significantly lower than the maximum value of μ_E ($\max \mu_E$). Instead, the comparison with the strategy that always makes the optimal choice sees, as expected, all the three policies losers. In fact, working in absence of full knowledge of the real weather conditions (as the optimal strategy would be able to do), they typically incur in some energy waste.

Actually, in Fig. 8a, there is the case of the policy P_{pre} that shows a lower minimum μ_E value in the first considered scenario. However, going deeper in the analysis of this case, the save in energy is payed in terms of unavailability, as can be seen in Fig. 9a, experienced because the switches are not always heated when needed. In fact, an inappropriate heating control action can result either in a pure energy waste, or in lower energy consumption (even lower than the optimal consumption) coupled with unavailability of some switches.

Another comment that applies to all the policies is that the energy values shown in Fig. 8b have less variability than in Fig. 8a. The reason is that the weather profiles considered in the former are more severe than in the other (the starting temperature

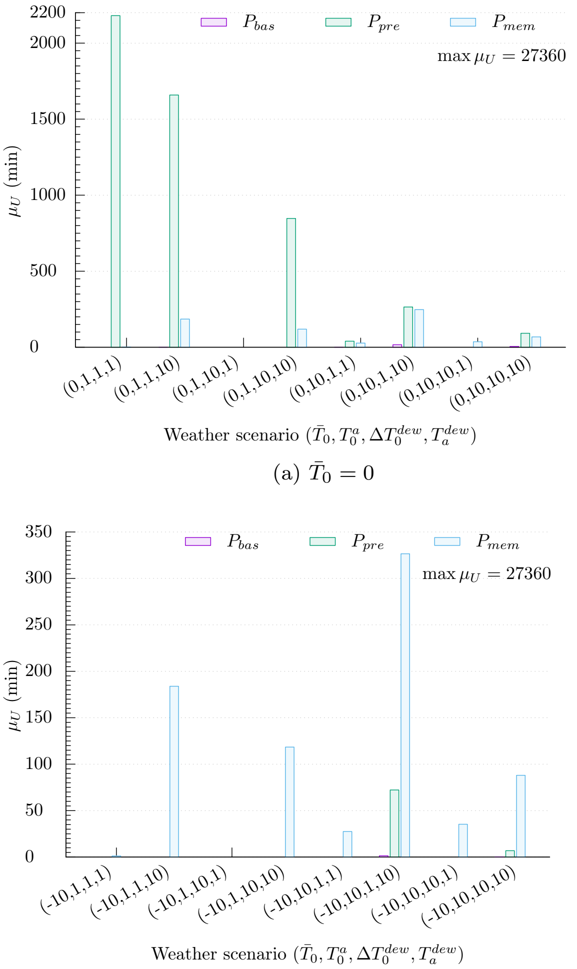

(a) $\bar{T}_0 = 0$

(b) $\bar{T}_0 = -10$

Figure 9: Unavailability for the different policies when (a) $\bar{T}_0 = 0$ and (b) $\bar{T}_0 = -10$.

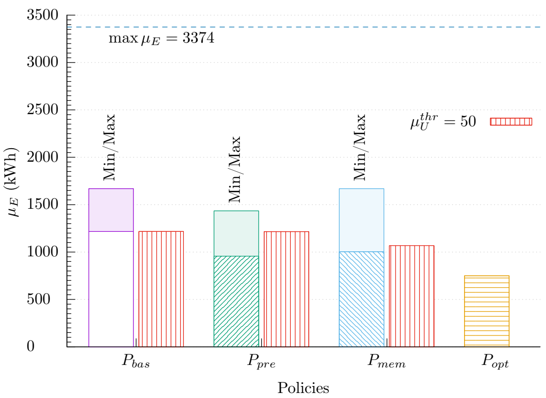

Figure 10: Energy consumption without unavailability constraints (minimum and maximum) and when μ_U^{thr} is 50 min for the different policies and the weather scenario $(0, 1, 1, 10)$.

is $-10°$), thus requiring the heating of the switch for longer time and giving less room for manoeuvre to the policies (especially to P_{pre} and P_{basic}, which have the minimum and maximum energy consumption coincident or very close). Comparing the three policies from the energy consumption only, it can be observed that P_{mem} shows the best behavior in almost all the considered scenarios when the minimum consumption is considered, although the improvements with respect to the other two is variable. Instead, the minimum values of the maximum consumption are shown by P_{pre}. The policy P_{bas} is never better than the others, proving that the variants developed on purpose to enhance this basic policy, really do what is expected. In Fig. 9, for each of the three studied policies, the maximum cumulative unavailability time in each weather scenario is reported. Concerning the minimum cumulative unavailability time, the bars do not appear, since all of them have value 0. From Fig. 9 it can be observed that μ_U behaves differently with respect to μ_E when $\bar{T}_0$ switches from 0 to -10: in Fig. 9a the value of μ_U for P_{pre} is always greater than μ_U for P_{mem}, and only in one case they are comparable, whereas in Fig. 9b the value of μ_U for P_{mem} is always greater than μ_U for P_{pre}. So, for μ_U it is not a matter of reduced variability but a clear distinction exists between mild and cold weathers, the former favoring P_{mem}, the latter P_{pre}. Notice in addition that the worst value of μ_U in Fig. 9b is below 350 minutes, whereas in Fig. 9a μ_U can reach 2200 minutes.

Finally, in Fig. 10 both indicators μ_E and μ_U are considered, showing the minimum and maximum energy consumption induced by the three policies, both when no unavailability constraint is considered and when a maximum unavailability value μ_U^{thr} cannot be exceeded. In the figure, $\mu_U^{\text{thr}} = 50$ minutes and the weather parameters are: $\bar{T}_0 = 0$, $T_0^a = 1$, $\Delta T_0^{dew} = 1$ and $T_a^{dew} = 10$ (one of the scenarios analyzed in Fig. 8). The optimal energy consumption as determined by the golden strategy P_{opt} is also included as the best value to aim to. It can be appreciated how the requirement on availability may lead to an increase in the energy consumption with respect to the minimum possible, especially for the policy P_{pre}. In particular, P_{pre} has a slightly better minimum value of μ_E than P_{mem} when availability is not considered, but the situation reverses, if switches unavailability can be tolerated only for short periods of time (in this scenario, no more than an average of 50 minutes, cumulated along a day by all the 19 railroad switches of the adopted case study).

Of course, both the behaviour of P_{pre} and P_{mem} depend on the values assigned to their internal parameters (which account for the ability to perform correct local prediction of future temperature values for P_{pre}, and for P_{mem} the ability to correctly estimate for how long to use the dewpoint received by the coordinator). In the above analyses, the values assigned to such parameters have been chosen within reasonable ranges (see Section 6.3), but different values could be adopted, if considered better ones by users of such policies, and their impact assessed by exercising the modeling framework with the new setting.

7. Conclusions

This paper presented novel solutions to address energy efficiency of the heating system for railroad switches. This is one of the major energy intensive equipment within the railway infrastructure, but also a critical one from the dependability perspective, since it contributes to the correct operation of the train routing. The proposed energy management strategies exploit availability of heterogeneous data and the fusion with data collected from external systems to achieve energy saving. In particular, weather information made available from meteorological stations are combined with temperature data coming from sensors deployed within the railway infrastructure, according to the different logics at the basis of the policies behavior.

A model-based stochastic framework has been also developed, to perform the assessment of the two proposed strategies and compare them in terms of energy consumption, also with a basic policy which represents current practice. An indicator of unavailability has been also introduced, to understand to which extent the lower energy consumption is obtained by a strategy at the expense of switch availability. The ability to represent generic weather conditions (indicative of different geographical zones, from extremely cold to milder areas) by properly tuning the characterizing four parameters, as well as failure events, is a strength of the assessment framework, since makes it a powerful and widely applicable instrument. It has been exercised on a conveniently chosen case study and, although the considered scenarios and parameters setting are not comprehensive of all the possible studies, they were good enough to show that the newly developed strategies P_{pre} and P_{mem} have the opportunity to be chosen as the most efficient strategy to employ. Of course, the modeling framework can be profitably exploited to analyze a multitude of other case studies, to gather wider knowledge on their behavior, and possibly trigger further refinements.

Further extensions are foreseen in several directions, including: i) to investigate the impact of parameters that have been kept constant in this analysis, especially the duration of the interval of time between two activations of the control system, the communication channels failure rate and the repair time of failed channel; ii) to consider other topology configurations, inducing different dependency among failures of communication channels; iii) further enhance the energy management policies by taking into account also the traffic on the railway tracks close to the switches. The idea would be that freezing of a switch is not dangerous if no train is transiting on the interested lines in a certain time interval, so extra energy could be saved. We believe that the approaches proposed in this paper are an important initial step, with potential for immediate practical exploitation, that helped to learn about the involved dynamics and requirements trade-offs, thus also facilitating the task of adding new aspects in the framework (e.g., the train scheduling).

Acknowledgments. This work is partially supported by the STINGRAY project, co-funded by the Tuscany Region under the POR FESR Toscana 2014-2020 program and by the Italian Ministry of Education and Research (MIUR) in the framework of the CrossLab project (Departments of Excellence).

References

[1] J. T. de Miranda Pinto, O. Mistage, P. Bilotta, E. Helmers, Road-rail intermodal freight transport as a strategy for climate change mitigation, Environmental Development 25 (2018) 100–110.

[2] J. Huang, Y. Deng, Q. Yang, J. Sun, An energy-efficient train control framework for smart railway transportation, IEEE Transactions on Computers 65 (5).

[3] P. Norbbin, J. Lin, A. Parida, Energy efficiency optimization for railway switches & crossings: a case study in Sweden, in: WCRR 2016, 11th World Congress on Railway Research, SPARK knowledge sharing portal, Milan, Italy, 2016.

[4] International Union of Railways, Technologies and Potential Developments for Energy Efficiency and CO2 Reduction in Rail Systems, `https://uic.org/IMG/pdf/_27_technologies_and_potential_developments_`

`for_energy_efficiency_and_co2_reductions_in_rail_systems._uic_in_colaboration.pdf`, online; accessed 15 January 2019.

[5] S. Chiaradonna, F. Di Giandomenico, G. Masetti, D. Basile, A refined framework for model-based assessment of energy consumption in the railway sector, in: M. H. ter Beek, A. Fantechi, L. Semini (Eds.), From Software Engineering to Formal Methods and Tools, and Back, Vol. 11865 of LNCS, Springer, Cham, 2019, pp. 481–501.

[6] International Energy Agency, The Future of Rail - opportunities for energy and the environment, `https://www.iea.org/reports/the-future-of-rail`.

[7] W. Gunselmann, Technologies for increased energy efficiency in railway systems, in: 2005 European Conference on Power Electronics and Applications, 2005.

[8] A. Achilleos, M. Anastasopoulos, A. Tzanakaki, M. Iordache, O. Langlois, J. Pheulpin, D. Simeonidou, Optimal driving profiles in railway systems based on data envelopment analysis, in: The 5th International Conference on Vehicle Technology and Intelligent Transport Systems - Volume 1: VEHITS, INSTICC, SciTePress, 2019, pp. 254–259. doi:10.5220/0007878002540259.

[9] D. Tokody, F. Flammini, The intelligent railway system theory: The european railway research perspective and the development of the european digital railway strategy, International Transportation.

[10] G. M. Shafiullah, S. A. Azad, A. B. M. S. Ali, Energy-efficient wireless MAC protocols for railway monitoring applications, IEEE Transactions on Intelligent Transportation Systems 14 (2) (2013) 649–659.

[11] H. Ghasemieh, B. R. Haverkort, M. R. Jongerden, A. Remke, Energy resilience modelling for smart houses, in: 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2015, 2015, pp. 275–286.

[12] D. Basile, S. Chiaradonna, F. Di Giandomenico, S. Gnesi, A stochastic model-based approach to analyze reliable energy-saving rail road switch heating systems, Journal of Rail Transport Planning & Management, Elsevier 6 (2) (2016) 163–181.

[13] D. Basile, F. Di Giandomenico, S. Gnesi, Model-based evaluation of energy saving systems, in: V. Kharchenko, Y. Kondratenko, J. Kacprzyk (Eds.), Green IT Engineering: Concepts, Models, Complex Systems Architectures, Vol. 74 of Studies in Systems, Decision and Control, Springer, Cham, 2017, pp. 187–208.

[14] D. Basile, F. Di Giandomenico, S. Gnesi, Statistical model checking of an energy-saving cyber-physical system in the railway domain, in: The 32nd ACM Symposium on Applied Computing, SAC 2017, Marrakech, Morocco, 2017, pp. 1356–1363.

[15] D. Basile, F. Di Giandomenico, S. Gnesi, On quantitative assessment of reliability and energy consumption indicators in railway systems, in: V. Kharchenko, Y. Kondratenko, J. Kacprzyk (Eds.), Green IT Engineering: Social, Business and Industrial Applications, Vol. 171, Springer, Cham, 2019, pp. 423–447.

[16] T. A. Henzinger, The theory of hybrid automata, LICS '96, IEEE Computer Society, 1996, pp. 278–292.

[17] R. David, H. Alla, On hybrid Petri nets, Discrete Event Dynamic Systems 11 (1-2) (2001) 9–40.

[18] W. H. Sanders, J. F. Meyer, Stochastic activity networks: Formal definitions and concepts, in: Lectures on Formal Methods and Performance Analysis, First EEF/Euro Summer School on Trends in Computer Science 2000, Revised Lectures, 2000, pp. 315–343.

[19] T. Courtney, S. Gaonkar, K. Keefe, E. W. D. Rozier, W. H. Sanders, Möbius 2.3: An extensible tool for dependability, security, and performance evaluation of large and complex system models, in: 39^{th} Annu. IEEE/IFIP Int. Conf. on Dependable Syst. and Netw. (DSN 2009), Estoril, Lisbon, Portugal, 2009, pp. 353–358.

[20] K. G. Larsen, P. Pettersson, W. Yi, UPPAAL in a nutshell, JSTTT 1.

[21] O. Tetens, Uber einige meteorologische begriffe, Zeitschrift fur Geophysik 6 (1930) 297–309.

[22] M. G. Lawrence, The relationship between relative humidity and the dewpoint temperature in moist air: A simple conversion and applications, Bulletin of the American Meteorological Society 86 (2) (2005) 225–234.

[23] R. Hyndman, A. B. Koehler, J. K. Ord, R. D. Snyder, Forecasting with exponential smoothing: the state space approach, 1st Edition, Springer Series in Statistics, Springer-Verlag Berlin Heidelberg, 2008.

[24] A. M. De Livera, R. J. Hyndman, R. D. Snyder, Forecasting time series with complex seasonal patterns using exponential smoothing, Journal of the American Statistical Association 106 (496) (2011) 1513–1527.

[25] S. Chiaradonna, F. Di Giandomenico, G. Masetti, A stochastic modeling approach for an efficient dependability evaluation of large systems with non-anonymous interconnected components, in: The 28^{th} Int. Symp. on Softw. Reliab. Eng. (ISSRE 2017) - IEEE, Toulouse, France, 2017, pp. 46–55.