

# Analysis of a sewage treatment facility using hybrid Petri nets

Hamed Ghasemieh  
University of Twente  
Netherlands  
h.ghasemieh@utwente.nl

Anne Remke  
University of Twente  
Netherlands  
a.k.i.remke@utwente.nl

Boudewijn R. Haverkort  
University of Twente,  
Netherlands  
b.r.h.m.haverkort@utwente.nl

## ABSTRACT

Waste water treatment facilities clean sewage water from households and industry in several cleaning steps. Such facilities are dimensioned to accommodate a maximum intake. However, in the case of very bad weather conditions or failures of system components the system might not suffice to accommodate all waste water. This paper models a real waste water treatment facility, situated in the city of Enschede, The Netherlands, as Hybrid Petri net with a single general one-shot transition (HPnGs) and analyses under which circumstances the existing infrastructure will overflow. This required extending the HPnG formalism with *guard arcs* and *dynamic continuous transitions* to model dependencies both on continuous places and on the rate of continuous transitions. Using recent algorithms for model checking STL properties on HPnGs, the paper computes survivability measures that can be expressed using the path-based until operator. After computing measures for a wide range of parameters, we provide recommendations as to where the system can be improved to reduce the probability of overflow.

## 1. INTRODUCTION

Any water that has been affected in quality either by households or by industries is considered as waste water. It is usually conveyed in the sewerage system of the community to the nearest waste water treatment facility. The treatment process consists of several physical, chemical and biological cleaning steps. The goal of the process is to separate the clean water from the so-called sludge, that can later be safely disposed or used as fertilizer. The cleaned water is usually released to surface water in the area.

In the Netherlands, communities normally have contracts with waste water treatment facilities about the maximum amount of waste water that needs to be taken in by the treatment facility. Hence, these facilities are dimensioned to accommodate the treatment of a maximum amount of sewage, without taking into account the possibility of unforeseen events. However, in the case of heavy rainfall, which is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VALUETOOLS 2013, December 10-12, Torino, Italy  
Copyright © 2013 ICST 978-1-936968-48-0

DOI 10.4108/icst.valuetools.2013.254384



Figure 1: An eagle view picture of the sewage treatment facility in Enschede, the Netherlands. The picture is retrieved using Google Maps.

hard to predict, it may happen that the amount of waste water in the community sewerage exceeds the available storage capacity. In such cases, the sewerage system of the community overflows and waste water is spilled on the streets. Recently, this happened in the city of Enschede, the Netherlands [1, 2, 3] and caused hindrance to citizens and traffic.

This paper investigates under which circumstances flooding occurs and what can possibly be done to reduce the probability of such flooding. For this purpose we have employed data like the capacity of tanks and the average residence time of water in the different cleaning stages from the treatment facility in the city of Enschede. A bird's-eye view of this facility is shown in Figure 1. This information is used to model the operation of the treatment facility as a Hybrid Petri-net with a single General one-shot transition (HPnG).

The modelling formalism of HPnG has recently been introduced for the analysis of fluid critical infrastructures [4], and efficient algorithms have been introduced for their analysis [5]. Note that even though the algorithms are currently restricted to HPnGs with a single general one-shot transition, it is still possible to thoroughly investigate the system evolution using parametrization of different factors, since analysis is extremely quick.

Moreover, in [6] the syntax and semantics of Stochastic Time Logic (STL) has been introduced together with algorithms to efficiently model check STL properties on HPnGs. Especially the analysis of the path-based until formula is suitable to evaluate how well the system performs

in the presence of bad weather conditions or failures at the intake to the treatment facility. Such properties are often referred to as a survivability properties [7], [8], [9], and mostly evaluated for, so-called, Given the Occurrence Of Disaster (GOOD) models. In such models, as the name suggests, the occurrence of a disaster is assumed to happen at a certain time of the day instead of trying to predict the probability of a disaster using risk management. The focus then lies on the effect, the handling and the recovery of the disaster, once it has happened. Note that since this is a case study paper, we do not cite all the related work with respect to our approach. However, pointers to related work can be found in [4], [5] and [6].

In this paper, we extend the HPnG formalism by two new features, namely, **guard arcs** and **dynamic transitions**, since they have been shown to be essential when modelling waste water treatment facilities. Guard arcs combine test and inhibitor arcs, as previously present in the formalism, but additionally allow to control discrete events of the system based on the values of continuous variables. Dynamic transitions are continuous transitions, where the rate depends on the rate of other continuous transitions. As will be explained in this paper, both extensions can be incorporated in the analysis and model checking algorithms without increasing their complexity.

Computing measures of interest for the HPnG modelling of the treatment facility for a wide range of parameters, then allows us investigate how and where the community and the treatment facility could invest best, e.g., through installing larger buffers or more pumping equipment, to reduce the residence time of waste water in the treatment phases in order to decrease the probability of spilling waste water in the streets.

To the best of our knowledge the quantitative evaluation of effects of failures or very bad weather conditions is not usually performed for waste water cleaning facilities in particular or in civil engineering in general. The common way of dimensioning such systems is to use static models and calculations [10]. Risk assessment is generally performed for civil engineering facilities in various ways [11], however, is not able to predict the consequences of failures in a quantitative way. Another approach that is commonly used is simulation [12], which is, however, very time consuming and does not allow to analyse a wide range of parameter settings quickly.

This paper is organized as follows, Section 2 introduces a modified version of the definition and modelling formalism of HPnG with addition of the new features of guard arcs and dynamic transitions. Section 3 revisits region-based analysis and the idea of the partitioning of the state space. Also, a brief definition of STL is provided. Section 4 demonstrates the use of region-based analysis and STL for a simple control example. In Section 5, using the new features of guard arcs and dynamic transitions, a new component for modelling overflow places in real world water treatment facilities is introduced. In Section 6, a case study for modelling and analysis of the sewage treatment facility in the city of Enschede is investigated. Finally, Section 7 concludes the paper.

## 2. MODEL DEFINITION

In the following we extend the HPnG definition from [4] by introducing dynamic transitions and guard arcs connecting fluid places and discrete transitions. Dynamic transi-

tions are a special form of continuous transitions where the outflow may depend on the flow of other continuous transitions, as explained later. This can be used to model, for example, overflow places. Test and inhibitor arcs have been used before to control the enabling of discrete or continuous transitions via the content of a connected discrete place. Guard arcs, as introduced in this paper, combine the functionality of test and inhibitor arcs and additionally allow to control the enabling of a transition by comparing the content of a continuous place with a given boundary condition. This allows to model typical control examples.

### 2.1 Model

As before, an HPnG is defined as a tuple  $(\mathcal{P}, \mathcal{T}, \mathcal{A}, \mathbf{m}_0, \mathbf{x}_0, \Phi)$ , where  $\mathcal{P} = \mathcal{P}^D \cup \mathcal{P}^C$  is a set of *places* that can be divided into two disjoint sets  $\mathcal{P}^D$  and  $\mathcal{P}^C$  for the discrete and continuous places, respectively. The discrete marking  $\mathbf{m}$  is a vector that represents the number of tokens  $m_P \in \mathbf{N}$  for each discrete place  $P \in \mathcal{P}^D$  and the continuous marking  $\mathbf{x}$  is a vector that represents the non-negative level of fluid  $x_P \in \mathbf{R}_0^+$  for each continuous place  $P \in \mathcal{P}$ . The initial marking is given by  $(\mathbf{m}_0, \mathbf{x}_0)$ .

Four types of *transitions* are possible, as follows. The set of immediate transitions, the set of deterministically timed transitions, the set of general transitions, and the set of continuous transitions together form the finite set of transitions  $\mathcal{T} = \mathcal{T}^I \cup \mathcal{T}^D \cup \mathcal{T}^G \cup \mathcal{T}^C$ . Note that, as in previous works, also in this paper the number of general transitions is restricted to  $|\mathcal{T}^G| = 1$ . The set of continuous transitions itself consists of two disjoint sets: static and dynamic transitions, denoted by  $\mathcal{T}^{Dy}$  and  $\mathcal{T}^{St}$ , respectively. Static continuous transitions are the same as the previous continuous transitions, i.e., they have constant nominal firing rates. Note that, due to rate adaptation the nominal rate of a static transition may change and is then called actual rate. In contrast to static transitions, where the nominal rates are always constant, the nominal rate of dynamic transitions may depend on the actual rate of any other static transition in the HPnG at hand.

The set of *arcs*  $\mathcal{A}$  consists of three sets: The set of discrete input and output arcs  $\mathcal{A}^D$ , connects discrete places and discrete transitions and the set of continuous input and output arcs  $\mathcal{A}^C$  connects continuous places and continuous transitions. The set of *guard arcs*  $\mathcal{A}^G$  connects discrete places to all kinds of transitions, and also, continuous places to all but continuous transitions. These arcs ensure that a transition is only enabled in case the number of tokens (in case of a discrete place) or the amount of fluid (in case of a continuous place) fulfills a certain condition that is specified on the guard arc.

The tuple  $\Phi = (\phi_b^P, \phi_p^T, \phi_d^T, \phi_{St}^T, \phi_{Dy}^T, \phi_g, \phi_w^A, \phi_u^A, \phi_s^A, \phi_p^A)$  contains 10 *functions*. Function  $\phi_b^P : \mathcal{P}^C \rightarrow \mathbf{R}^+ \cup \infty$  assigns an upper bound to each continuous place. In contrast to the definition of HPnG in [4] in the following  $\phi_p^T : \mathcal{T}^D \cup \mathcal{T}^I \rightarrow \mathbf{N}$  specifies a *unique priority* to each immediate and deterministic transition to resolve firing conflicts, as in [13]. Deterministic transitions have a constant firing time defined by  $\phi_d^T : \mathcal{T}^D \rightarrow \mathbf{R}^+$  and continuous static transitions have a constant nominal flow rate defined by  $\phi_{St}^T : \mathcal{T}^{St} \rightarrow \mathbf{R}^+$ . Mapping  $\phi_{Dy}^T : \mathcal{T}^{Dy} \rightarrow f$ , assigns a function  $f : \mathbf{R}^{|\mathcal{T}^{St}|} \rightarrow \mathbf{R}^+$  to each dynamic continuous transition, which determines how its nominal flow rate depends on the continuous static transitions rates. The general transition is associated with a random variable  $S$ , representing its firing time, according to a

cumulative distribution function (CDF)  $\phi_g(s)$ , and its probability density function (PDF) is denoted  $g(s)$ . We assign a weight to all discrete input and output arcs:  $\phi_w^A : \mathcal{A}^D \rightarrow \mathbf{N}$  which defines the number of tokens that is taken from or added to connected places upon the firing of the transition.  $\phi_u^A : \mathcal{A}^G \rightarrow \{(\triangleright, \mathbf{R})\}$ , with  $\triangleright = \{\geq, <\}$  assigns a tuple which consists of a comparison operator and a real number to all guard arcs. The functions  $\phi_s^A, \phi_p^A$  specify the share and the priority of a static continuous transition as will be explained later.

## 2.2 Graphical representation

The primitives of the hybrid Petri net formalism with general one-shot transitions are shown in Figure 2. A discrete place is graphically represented by a single circle and a fluid place is represented by two concentric circles. A general transition is drawn as an empty rectangle, a deterministic transition is drawn as a grey rectangle, a continuous static transition is shown as an empty rectangle with double lines, a continuous dynamic transition is shown by a double lined solid rectangle and an immediate transition is a thick black bar. The discrete input and output arcs are drawn as single arrows and fluid input and output arcs are represented with double lines. Guard arcs are drawn with two triangular arrowheads.




Figure 2: Graphical representation of primitives of HPnG.

## 2.3 System evolution

Markings are collected into two vectors, the discrete marking  $\mathbf{m} = (m_1, \dots, m_{|\mathcal{P}^D|})$  and the continuous marking  $\mathbf{x} = (x_1, \dots, x_{|\mathcal{P}^C|})$ . The initial marking is composed of a discrete part  $\mathbf{m}_0$  that describes the initial amount of tokens in all discrete places and a continuous part  $\mathbf{x}_0$  that describes the initial amount of fluid in all continuous places.

The *state* of an HPnG is defined by  $\Gamma = (\mathbf{m}, \mathbf{x}, \mathbf{c}, \mathbf{d}, \mathcal{G})$ , where vector  $\mathbf{c} = (c_1, \dots, c_{|\mathcal{T}^D|})$  contains a clock  $c_i$  for each deterministic transition that represents the time that  $T_i^D$  has been enabled. When a transition is disabled the clocks do not evolve, but the clock value is preserved until the transition is enabled again. Clocks are only reset when the corresponding deterministic transition fires. If the general transition has not fired yet, it can be considered as a deterministic transition, whose firing time is sampled from the corresponding general firing time distribution. This sampling happens only once per model execution, and it occurs when the general transition becomes enabled for the first

time. Vector  $\mathbf{d} = (d_1, \dots, d_{|\mathcal{P}^C|+|\mathcal{T}^D|})$  indicates the drift (speed of change) of all continuous variables. For continuous places it shows the change of fluid per time unit, and for deterministic transitions it is the clock drift which is either one or zero, if the transition is enabled or disabled, respectively. Note that even though the vector  $\mathbf{d}$  is determined uniquely by  $x, m$ , and weight of guard arcs, it is included in the definition of a state for ease of analysis. The general transition is only allowed to fire once, hence, the flag  $\mathcal{G} \in \{0, 1\}$  indicates whether the general transition has already fired ( $\mathcal{G} = 1$ ), or not ( $\mathcal{G} = 0$ ). So, the initial state of the system is  $\Gamma_0 = (\mathbf{m}_0, \mathbf{x}_0, \mathbf{0}, \mathbf{d}_0, 0)$ . A system state can be seen as a snapshot of the system evolution at a specific time, and assumed general transition firing time; this is elaborated in more detail in the next section.

The firing rules of deterministic, general, immediate and fluid transitions differ. Whether a transition is allowed to fire depends (1) on the structure and the current marking of the Petri net (*concession*) and (2) on the type of the transitions [14].

Continuous transitions that have concession are always enabled, and continuously transport fluid along fluid arcs. Conflicts in the distribution of fluid occur when a continuous place reaches one of its boundaries. To prevent overflow, the fluid input has to be reduced to match the output, and to prevent underflow the fluid output has to be reduced to match the input, respectively. The firing rate of both, static and dynamic continuous transitions is then adapted according to the share  $\phi_s^A : \mathcal{A}^C \rightarrow \mathbf{R}^+$  and priority  $\phi_p^A : \mathcal{A}^C \rightarrow \mathbf{N}$  that is assigned to the continuous arcs that connect the transition to the place. This is done by distributing the available fluid over all continuous arcs. Those with highest priority are considered first and if there is enough fluid available, all transitions with the highest priority can still fire at their nominal speed. Otherwise, their actual fluid rates are adapted according to the firing rate of the connected transitions and the share of the arc, according to [14]. The adaptation of fluid rates in these cases results in a piecewise constant fluid derivative per continuous place.

Non-fluid transitions that have concession may be enabled, depending on their type. If an immediate transition has concession the marking is said to be *vanishing* otherwise the marking is said to be *tangible*. Immediate transitions have precedence over deterministic and general transitions. In a vanishing marking deterministic and general transitions are disabled and cannot fire. The clock of each enabled deterministic transition  $T_i^D$  evolves with time at rate  $dc_i/d\tau = 1$  and when a clock reaches its firing time, i.e.,  $c_i = \phi_d^T(T_i)$  transition  $T_i^D$  fires. Similarly, the enabling time of the enabled general transition, that has not fired yet, evolves with time at rate 1. The general transition then fires with probability  $\phi_g(\tau + \Delta\tau) - \phi_g(\tau) = \int_{\tau}^{\tau + \Delta\tau} g(s)ds$  in any time interval  $[\tau, \tau + \Delta\tau]$ .

Whenever a non-fluid transition fires the marking evolves according to a firing rule, depending on the type of the transition. All discrete transition types, i.e., immediate, deterministic and general, change the discrete part of the marking  $\mathbf{m}$  in a similar way. For a more detailed description of HPnGs and their evolution, we refer to [15].

## 3. ANALYSIS

We recall how the underlying state space of a HPnG is partitioned and how this can be used to efficiently analyse the system at hand in Section 3.1. Section 3.2 explains the




Figure 3: Generic presentation of STD.

basic idea behind model checking HPnGs and how to formulate STL properties.

### 3.1 Stochastic time diagram

The Stochastic time diagram (STD) as introduced in [5] provides a genuine way of representing the evolution of a HPnG for a given initial state. The main reasoning behind this is that, for an initial state of an HPnG and a predefined value for the firing time of the general transition, denoted  $s$ , for all future time instances  $t$  we can determine the state of the system. The STD is a two-dimensional diagram with  $t$  and  $s$  as its vertical and horizontal axis, respectively. Each point in this diagram is associated with a unique HPnG state, which is denoted by  $\Gamma(s, t)$ . Note that, for a fixed value of  $s$ , the evolution over time is deterministic and associates with a vertical line in STD. A generic version of this diagram is shown in Figure 3. Two main areas can be distinguished in this figure. The area above the line  $t = s$ , called *stochastic area*, contains all the HPnG states for which the general transition has fired ( $t > s$ ), whereas the area below the line  $t = s$ , called *deterministic area*, includes those states for which the general transition has not yet fire ( $t < s$ ). To compute measures of interest for HPnGs, the STD needs to be deconditioned with the probability density function  $g(s)$ .

The main idea behind the method proposed in [5], is that instead of dealing with infinitely many points in the  $ts$ -plane, we can partition it into several *regions*. These regions exist, because the state of the system does not change until an *event* occurs. In each system state, three types of potential events can occur: (I) a continuous place reaching its lower/upper boundary, (II) a continuous place reaches the weight of the guard arc connected to it, and (III) an enabled transition, either deterministic or general fires. Event type (I) imposes a change in the drift of the continuous place, due to rate adaptation [14], and event type (II) will enable or disable a transition. In case of an immediate transition, it will fire and alter the discrete marking immediately, and if it is a deterministic transition its clock drift will be set to one, therefore changing a continuous variable. Finally, event type (III) alters the discrete marking or the general transition flag. Hence, in overall, an event may cause a change in the discrete marking, a change in drift (either for clocks or fluid levels) or a change in a general transition flag. We define a region as a set of states that while the system remains in them no event occurs, i.e., discrete marking, drift of continuous variables and general transition flag remain

unchanged. Moreover, by occurrence of an event the system enters another region. This leads to the following definition.

DEFINITION 1. A region  $\mathcal{R}$  is a set of  $(s, t)$  points in a given STD, for which we have:

$$\forall (s_1, t_1), (s_2, t_2) \in \mathcal{R} : \begin{cases} \Gamma(s_1, t_1).\mathbf{m} &= \Gamma(s_2, t_2).\mathbf{m}, \\ \Gamma(s_1, t_1).\mathbf{d} &= \Gamma(s_2, t_2).\mathbf{d}, \\ \Gamma(s_1, t_1).\mathcal{G} &= \Gamma(s_2, t_2).\mathcal{G}. \end{cases}$$

In which, while  $\Gamma(s, t)$  is used to refer to the vector of discrete markings,  $\Gamma(s, t).m_P$  is used to refer to the discrete marking of a specific place  $P$ . A similar notation is used for the continuous marking.

Note that this definition is different from the one in [6], since now vector  $\mathbf{d}$ , in addition to drifts of continuous places, also includes clock drifts for deterministic transitions, which are either one or zero. The reason for this is that, by introducing guard arcs, a deterministic transition can be enabled or disabled for the same discrete marking, due to a change in the continuous marking. This contradicts the idea of grouping system states into regions. Because of this fact, clock values of deterministic transitions are explicitly added to the definition.

An example partitioning of the state space into regions is shown in Figure 4 together with probability density function  $g(s)$ . Note that the shape of these regions depends on the structure of the model at hand. In [5] it is shown that, inside a region all continuous variables, i.e., the amount of fluid and the clock valuations, can be represented by simple linear equations in  $s$  and  $t$ . Intuitively, this is because in a region all continuous places are associated with a constant drift and clocks also have a constant drift (of one). Using this we infer that the boundaries between regions, which represent the occurrence of an event, are characterized by linear functions of  $s$  and  $t$ . Hence, each region in the STD is a polygon. Note that, introduction of dynamic fluid transitions does not change this fact, because their nominal rates depend on the actual rates of other static continuous transitions, which are constant, *within* each region. Hence, we can safely treat dynamic transitions as static transitions. To compute the probability to be in a specific system state at time  $\tau$ , it suffices to find all regions intersecting the horizontal line  $t = \tau$  that correspond to the specific system state and integrate  $g(s)$  over the intersection. This idea is illustrated for a given partitioning in Figure 4.




Figure 4: Deconditioning according to the probability density function  $g(s)$ .

Even though reachability computations on the STD are

always performed for a given and finite time bound, there is still a possibility of having an infinite number of regions in the STD before the finite time bound. This happens whenever an infinite sequence of vanishing markings occurs. This problem is well-known for all Petri nets formalism that allow immediate transitions. However, if we require that models have to be bounded, infinite sequences of vanishing markings can only take place in the form of cycles of vanishing markings, which can be detected and removed [16]. This ensures that we always reach a tangible marking in a finite number of steps and the number of regions in the STD before a finite time bound is also finite. Hence, for a bounded model and a finite time bound the algorithm always terminates.

### 3.2 Stochastic Time Logic

Stochastic Time Logic (STL) introduced in [6], allows us to represent and evaluate path-based formula, such as time-bounded reachability, to answer whether a certain property is reached, within a certain time, while another given condition holds. STL is basically an extended version of the state-based logic in [5], with until operator. STL is used to reason about the underlying state space of an HPnG, i.e., it is possible to reason whether an STL formula holds for a certain system state  $\Gamma(s, t)$ . However, note that STL reasons on the conditioned state-space of an HPnG, that is, on the regions of an STD, which does *not* yet take into account the distribution of the general transition.

**DEFINITION 2 (STOCHASTIC TIME LOGIC).** *An STL formula  $\Psi$  is defined as*

$$\Psi := \mathbf{tt} \mid x_P \geq c \mid m_P = a \mid \neg\Psi \mid \Psi \wedge \Psi \mid \Psi \mathcal{U}^{[T_1, T_2]} \Psi,$$

where  $T_1, T_2 \in \mathbb{R}^+$ ,  $x \geq c$  and  $m = a$ , with  $a \in \mathbb{N}, c \in \mathbb{R}^+$ , are called *continuous and discrete atomic properties, respectively*.

Note that, although the above definition allows nested until formula, we have only considered non-nested until formula, so far [6].

In the following two different satisfaction relations are introduced. Firstly,  $\models^{s,t}$  between a single system state  $\Gamma(s, t)$  and an STL formula  $\Psi$ , which is intuitively true if the system state at that certain point satisfies the formula. Secondly,  $\models^t$ , between an interval on the support of the general transition and an STL formula  $\Psi$ . The different indices on the satisfaction relation are used to stress their dependencies on  $s$  and  $t$ .

**DEFINITION 3 (SATISFACTION ON SYSTEM STATES).**

$$\begin{aligned} \Gamma(s, t) \models^{s,t} \mathbf{tt} & \quad \forall t, s, \\ \Gamma(s, t) \models^{s,t} m_P = a & \quad \text{iff } \Gamma(s, t).m_P = a, \\ \Gamma(s, t) \models^{s,t} x_P \geq c & \quad \text{iff } \Gamma(s, t).x_P \geq c, \\ \Gamma(s, t) \models^{s,t} \neg\Psi & \quad \text{iff } \Gamma(s, t) \not\models^{s,t} \Psi, \\ \Gamma(s, t) \models^{s,t} \Psi_1 \wedge \Psi_2 & \quad \text{iff } \Gamma(s, t) \models^{s,t} \Psi_1 \wedge \Gamma(s, t) \models^{s,t} \Psi_2, \\ \Gamma(s, t) \models^{s,t} \Psi_1 \mathcal{U}^{[T_1, T_2]} \Psi_2 & \quad \text{iff } \exists \tau \in [t + T_1, t + T_2] : \\ & \quad \Gamma(s, \tau) \models^{s,t} \Psi_2 \wedge (\forall \tau' \in [t, \tau] : \Gamma(s, \tau') \models^{s,t} \Psi_1). \end{aligned}$$

For the STL until operator  $\Psi_1 \mathcal{U}^{[T_1, T_2]} \Psi_2$  and a system state  $\Gamma(s, t)$ , we have to check for a fixed value of  $s$  and starting time point  $t$ , whether the evolution of the system is such that a time point  $\tau$  exists at which  $\Psi_2$  holds and before which  $\Psi_1$  holds. Recall that, for the system state  $\Gamma(s, t)$  and a fixed sample  $s$ , the evolution over time is deterministic and coincides with a vertical line in the STD, starting at

point  $(s, t)$ . Hence, the analysis of the STL until operator for a given system state is to check, whether this line only intersects with regions where  $\Psi_1$  holds until a region is hit where  $\Psi_2$  holds within the defined time interval. For details, we refer to [6].

We also introduce a satisfaction relation  $\models^t$  for intervals on the support of the distribution of the general transition, denoted  $I_\psi \subseteq \mathbb{R}_{\geq 0}$ , and STL formula  $\Psi$ , at time  $t$ . This allows for a more efficient model checking procedures than checking each system state individually.

**DEFINITION 4 (SATISFACTION ON INTERVALS).**

$$I_\psi \models^t \Psi \quad \text{iff } \forall s \in I_\psi : \Gamma(s, t) \models^{s,t} \Psi.$$

**DEFINITION 5.** *The set of satisfaction intervals  $Sat^t(\Psi)$  is defined as the set of all intervals satisfying  $\Psi$  at time  $t$ , i.e.,  $Sat^t(\Psi) = \{I_\psi : I_\psi \models^t \Psi\}$ .*

While the explicit dependency on  $s$  (or sets of  $s$ -values) is used for the efficient computation of properties, in the end we want to know whether a given STL formula holds at time  $t$  for the HPnG model of interest with a certain probability. Hence, we introduce a probability operator  $\mathbb{P}_{\sim p}(\Psi)$  which is wrapped around an STL formula, where  $p \in [0, 1]$  is a real number and  $\sim \in \{\leq, <, >, \geq\}$  a comparison operator. It abstracts from the possible values of  $s$  by deconditioning with the probability density function  $g(s)$ , as follows.

**DEFINITION 6.** *Let  $\Gamma(t) = \{\Gamma(s, t) \mid s > 0\}$  be the set of possible system states at time  $t$ , then the satisfaction relation for the probability operator  $\mathbb{P}_{\sim p}$  is defined as:*

$$\Gamma(t) \models \mathbb{P}_{\sim p}(\Psi) \quad \text{iff } Prob(\Psi, t) \sim p,$$

where

$$Prob(\Psi, t) = \sum_{I_\psi \in Sat^t(\Psi)} \int_{I_\psi} g(s) ds.$$

Model checking algorithm for STL logic, involves computational geometry, especially polygon clipping algorithms. Detailed description of the algorithms is given in [6].

## 4. CONTROL EXAMPLE

In this section we discuss a well-known control example to demonstrate the use of guard arcs. In the example as shown in Figure 5, a tank denoted  $P_M$  with the capacity of 11 Litres, is connected to two pumps. Either, the producer pump  $T_P$  fills the tank with rate 1 Litres/Minute, or the consumer pump  $T_D$  takes out the fluid with rate 2 Litres/Minute. For control purposes, the amount of fluid in the reservoir needs to be between 1 and 10 avoiding both, underflow and overflow. We also assume that the overall flow from the place  $P_M$ , can not be stopped, i.e., the two pumps  $T_P$  and  $T_D$  can not be off at the same time. Two switches can turn the pumps on and off, both with a delay of 2 Minutes, which is modelled by two deterministically timed transitions. Transition  $T_a$ , with firing time of 2 is connected to the reservoir via a guard arc with condition  $(\geq, 8)$ . Hence, when the amount of fluid is greater or equal to 8, it will be enabled and after 2 minutes it will fire. As a result the pump  $T_P$  becomes disabled and pump  $T_D$  becomes enabled. Also, transition  $T_b$ , with firing time of 2, is connected to the reservoir via a guard arc with condition  $(<, 5)$ , so when the amount of fluid is smaller than 5, the transition will be enabled and fires after two minutes.

Additionally, the producer pump may fail at different points in time  $\alpha$ , which is modelled by the deterministic timed transition  $T_F$ , with firing time  $\alpha$ . Whenever this transition fires, the general transition  $T_R$  becomes enabled. This transition models the repair procedure, which is stochastically distributed according to any arbitrary given probability distribution. Note that the input arc between place  $P_b$  and transition  $T_F$  ensures that only one failure is possible.




Figure 5: A simple control example. The amount of fluid in the reservoir  $P_m$  is supposed to remain between 1 and 10.

We want to check whether even after a failure the fluid level in the reservoir is always between one and ten. For this purpose we model check the following STL formula at the starting time:

$$\mathbb{P}_{\leq 0.001}(tt \mathcal{U}^{[\alpha, \alpha+T]} (x \leq 1 \vee x \geq 10))$$

In this formula variable  $x$  represents the amount of fluid in the reservoir, and time bound  $T$  is the maximum time, after failure at time  $\alpha$ , for which we want to check that a state with less than 1 amount of fluid or at least 10 amounts of fluid is only reached with a very small probability.

Figure 6 shows, the STD of the control example, for the case of  $\alpha = 3$ . Green regions are representing those regions, in which the condition  $(x \leq 1 \vee x \geq 10)$  does not hold, and blue region are reached after the time boundary  $\alpha+T$ . It can be seen that for all the possible values of  $s$  ( $x$ -axis), which represents the general transition firing time, it is impossible to reach a region in which the condition  $(x \leq 1 \vee x \geq 10)$  holds within the maximum time  $T$ . In other words, there is no value of  $s$  for which the formula is satisfied. Hence, if we integrate over all values of  $s$ , as described in Section 3.2, the property holds with probability zero, and therefore the overall formula with the probability operator is satisfied.

## 5. OVERFLOW PLACES

In the HPnG formalism as presented before, rate adaptation prevents both overflow and underflow of a reservoir. However, for modelling real systems, especially water treatment and sewage facilities, sometimes we need to allow places to overflow. This was impossible before adding dynamic transitions and guard arcs to the definition of HPnGs. In case of a full continuous place, the overflow is defined as the




Figure 6: Stochastic time diagram of control example, for  $\alpha = 3$ . Regions in which  $(x \leq 1 \vee x \geq 10)$  does not hold are outlined in green, and regions in which the time boundary  $T + \alpha$  is reached, are outlined in blue color.

difference of the actual rate of the inflow and the outflow of that place. This is the reason for adding dynamic transitions to the definition of HPnG. More formally an overflow place is a structure in which when the continuous place reaches its upper boundary the dynamic transition becomes enabled with the rate that equals the difference of the actual rate of all the incoming and outgoing transitions.

Note that, the rate adaptation algorithm has no influence on an overflow place. This is because at the moment of reaching the upper boundary the state of the system is changed by the enabling of the dynamic transition. As a result, the drift of the place becomes zero and rate adaption is not necessary anymore.





Figure 7: Modelling of an overflow place using HPnG primitives.

As can be seen in Figure 7, where an overflow place is shown, whenever the continuous place reaches a certain boundary  $B$ , the immediate transition connected via the guard arc will fire, and as a result the dynamic transition becomes enabled. Note that the rate of this transition is adapted according to the inflow and outflow of the main continuous place. Also, whenever the fluid level in the continuous place is below the boundary  $B$ , due to a change either in inflow or outflow, the connected immediate transition will fire and the dynamic transition is disabled. For the case of modelling, we represent an overflow place graphically as follows,

in Figure 8.



**Figure 8: Overflow place component.**

In Figure 8, the sum of inflow and outflow rates are denoted by  $\sum f_i$  and  $\sum f_o$ , respectively. Moreover, the overflow rate is shown by  $R_o = \sum f_i - \sum f_o$ , and distinguished from ordinary outgoing transitions by a shaded rectangle connected to the overflow place.

## 6. THE ENSCHEDE SEWAGE TREATMENT FACILITY

In this section we analyse the survivability of a sewage treatment facility, inspired by the operational facility located in Enschede, the Netherlands (Figure 1).

### 6.1 System and Model

Waste water that is produced by citizens and nearby industries is directed towards the sewage treatment facility through the sewerage system of the community. The capacity of this system is, by design, limited. Since the waste water treatment facility by contract is only obliged to a certain maximum intake, in case of very heavy rainfall the sewerage system has shown to be too small, which results in flooding the street in front of the treatment facility, and thereby posing serious hinder to citizens and traffic [1, 2, 3]. We investigate under which circumstances flooding of the street occurs and which parameters of the system need to be changed in order to prevent this.

This case study models the various stages of the sewage treatment process in an abstract fashion. We are mainly interested in the capacity of each phase and the average amount of time the waste water stays in the different phases. We, however, do not aim at modelling the physical, chemical and biological processes in detail. Then, for a given failure of the system at a certain time, we analyse the survivability of the system for changing weather conditions. Fixing the failure to a specific time of the day results in a so-called Given the Occurrence Of Disaster (GOOD) model. Since our evaluation method is so quick, it is easily possible to parametrize the failure time, hence, analyse the system thoroughly.

The main goal of waste water treatment is to separate the input into water that can be safely released into the environment and into thickened sludge which is either used as fertilizer [17] or can be safely disposed [18]. This is done in several stages, where the primary stage mostly involves physical purification, the secondary stage involves chemical and biological treatment, and finally the sludge treatment phase aims at reducing the amount of sludge.

The HPnG model of the case study is depicted in Figure 9; volumes of tanks (continuous places) are indicated in  $1000 m^3$ , pump rates (continuous transitions) in  $1000 m^3/h$ , and delays (timed transitions) in hours. The capacity of the community sewerage system is modeled by an overflow

place denoted  $P_c$ , which has input rates that depend on the weather conditions. From this tank the water is pumped into the treatment facility with a maximum rate 12 and in case the input exceeds the capacity of the place and the intake of the treatment facility, the waste water flows into (overflow) place  $P_o$  which models the amount of water in the streets. The primary stage of the sewage treatment consists of two phases, namely the sand interceptor and the primary sedimentation tank. The first, as the naming suggests, is responsible for filtering solids like sand from the water. Then the sewage flows in a large tank, which is used to settle the sludge, while the lighter material, like oils, rise on the surface and are removed, and the remaining overflows. In the model the sand interceptor is abstracted through pump  $P_z$ , and the primary sedimentation tank is modelled through overflow place  $P_{ps}$ .

A sedimentation tank physically separates suspended solids from water using gravity [18]. While the dirt settles at the ground, cleaned water is forwarded to the second cleaning stage. This stage consists of several phases for removing chemical and biological contaminations, modelled by a sequence of continuous transitions and places, before a secondary sedimentation tank separates the biological material from the now environment friendly sewage water, that can safely be disposed to surface water. The second sedimentation tank is modelled by overflow place  $P_{ss}$ . The sludge that settles at the primary and secondary sedimentation tank is accumulated and forwarded to the sludge treatment stage. There it is thickened to reduce its volume for easier off-site transport. The sludge from the primary tank is pumped out and forwarded to the fresh sludge thickener. This is also modelled by an overflow place, denoted  $P_{ft}$ . Sludge is pumped out of the place with a small rate and discharged to the digestion tank which is considered a very large tank. The overflow is directed to the filtrate basement. The same procedure is repeated for the accumulated sludge in the second sedimentation tank.

### 6.2 Evaluations

In the following, we analyse the model in two different ways, namely, by changing the rate of the produced waste water after a random amount of time, and by introducing a stochastic failure at the sand interceptor  $P_z$ , which according to the plant operators is one of the most vulnerable components of the whole process. In Figure 9, the first scenario is depicted by the dashed box A and the second is shown in box B. Note that, we either analyse scenario A or B, but never both at the same time, due to restriction to single general one-shot transition. We analyse the influence of several system parameters on the measures of interest. Also, note that we start the analysis assuming that all tanks in the treatment facility (fluid places and overflow places) are full but the overflow place modelling the community sewage system is empty.

#### Scenario A

For the first scenario, as depicted by box A, we assume that the analysis starts at normal weather condition, i.e., the production rate of waste water is 3.3. However, after a while it starts to rain, due to firing the deterministic transition  $T_r$  at time  $\alpha$ . In the following we assume  $\alpha = 3$  (measures of interest can easily be derived for other values). After firing the deterministic transition the waste water production changes to 12.2, which is slightly more than the capacity of

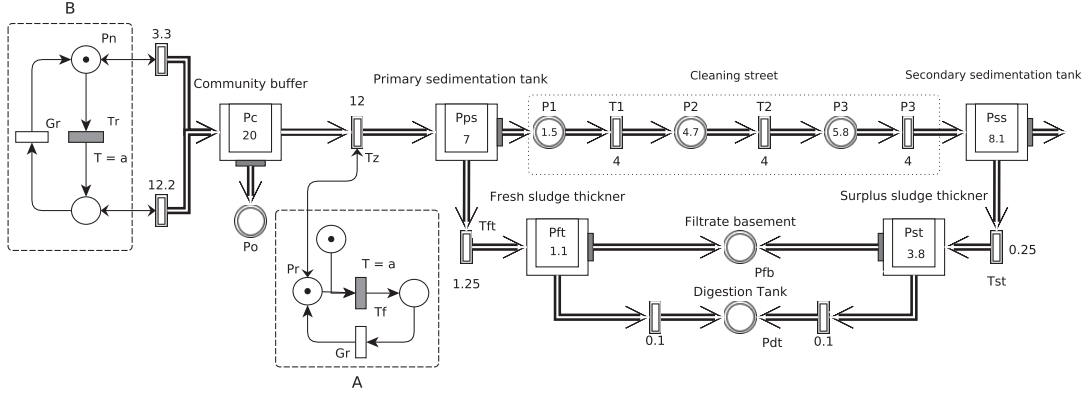


Figure 9: The abstract HPnG model of the sewage system in Enschede, Netherlands

the system, which is 12. Based on the region that the facility is placed in, the length of the rain could be distributed according to any probability density. We assume that the duration of the rain is normally distributed, however, based on any further knowledge this could be changed easily, due to modelling capabilities of HPnGs. This is modelled by the general transition  $G_r$ , which will fire according to the given normal distribution. Hence, if it continues to rain long enough, the capacity of the community sewage is exceeded and waste water will flood the streets.

We would like to analyse how long it may continue to rain without having water in the streets. Using the logic STL, as explained in Section 3.2, we want to ensure that the amount of water in the streets is very low until the rain stops, i.e.,

$$\Phi_A = (x_{P_o} < 0.01) \mathcal{U}^{[\alpha, \alpha+30]} (m_{P_n} = 1), \quad (1)$$

where  $m_{P_n} = 1$  means that the rain has stopped and we are back to the normal conditions. Formula  $\Phi_A$  is a typical expression of a survivability measure: the first term, before the Until operator is called the *safety condition*, whereas the one after the Until operator is called the *recovery condition*. In other words, as defined in Definition 3, for a specific system evolution,  $\Phi_A$  is satisfied if and only if the safety condition holds until we reach the recovery condition, before the given time bound. We have chosen time bound 30, which is considered to be big enough for this analysis, since it is reasonable that the rain stops within 30 hours.

In the following we investigate the influence of four different parameters, for varying average duration of rain  $\mu$ ; we consider the capacity of the community buffer ( $P_c$ ), the rate of the fresh sludge pump ( $T_{ft}$ ), the rate of the cleaning street ( $T_1$ ,  $T_2$  and  $T_3$ ) and the rate of the surplus sludge pump ( $T_{st}$ ). Figure 10 shows the probability that  $\Phi_A$  holds for varying mean durations of rain between 0.1 hour (6 minutes) and 4 hours. This parameter is the same for all four 3D-plots, and is depicted in the  $x$ -axis, while the parameters on the  $y$ -axis are different, as mentioned above. All the other characteristics of the model, which are not explicitly parametrized, keep their values according to Figure 9.

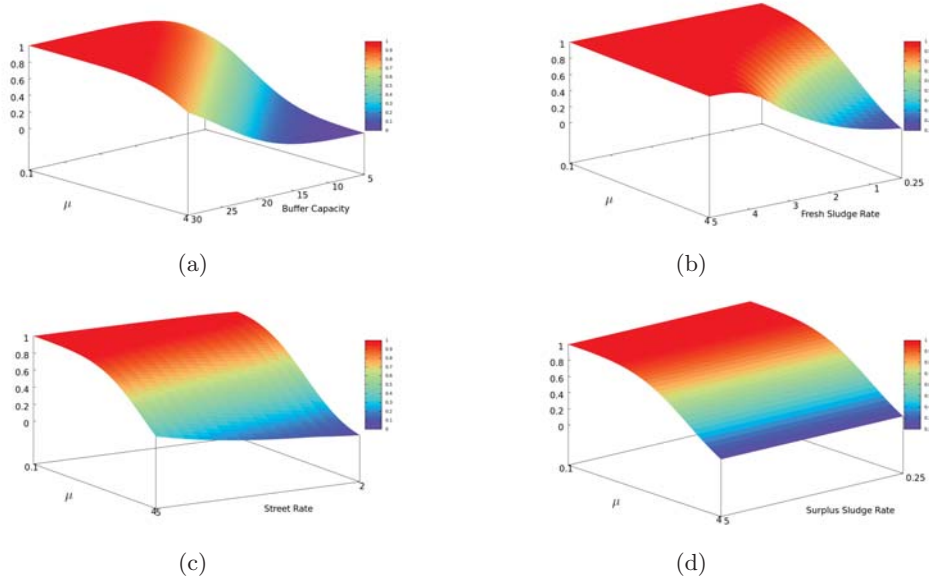
Figure 10(a) shows the influence of the capacity of the community buffer ( $P_c$ ), by varying its value from 5 to 30 (from right to left on the  $y$ -axis). By increasing this capacity, the probability that formula  $\Phi_A$  holds increases. We observe

that this increase is non-linear, especially for larger values of the capacity, we see a faster improvement. Furthermore, we observe that for long rain duration, even if we increase the buffer capacity to 30, still we have more than 20 percent probability of not satisfying the survivability property  $\Phi_A$ . This means that enlarging the buffer capacity alone is not enough for avoiding the flood in the area.

Figure 10(b) shows how the system survivability depends on the rate of the fresh sludge thickener pump,  $T_{ft}$  which is parametrized from 0.25 to 5 (from right to left on the  $y$ -axis). It can be seen that, by increasing the rate of this pump, the probability for  $\Phi_A$  to hold increases. Specially for long rain duration ( $\mu = 4$ ) this increase can be observed well. The increase is steeper than in Figure 10(a), hence, this pump plays a significant role; for larger values of its rate, e.g., larger than 3, even if it rains for more than four hours, formula  $\Phi_A$  holds with probability one. The reason for this is that the overflowed sewage from the primary tank,  $P_{ps}$  could be handled with a rate of at most 4, the rate of intake into cleaning street. So, the more we pump out of the primary tank, the more sewage intake the system can handle. However, since increasing the rate of this pump means pumping out sludge with more mixed water, this could be a disadvantage or even an obstacle for the next stage, i.e., sludge treatment.

Figure 10(c) shows the importance of the cleaning street pump rates, i.e., pumps  $T_1$ ,  $T_2$ ,  $T_3$ , of which we vary the rate from 2 to 5 (right to left). These pumps play a similar role as pump  $T_{ft}$ , but with lower impact. As can be seen, for long rain duration ( $\mu = 4$ ) the survivability probability remains low. Like in the previous case, also here increasing the pump rates could be problematic, because raising the rates involves pumping out water mixed with more dirt, since there may not have been enough time for the dirt to settle down in the primary sedimentation tank.

Finally, Figure 10(d), shows the influence of surplus sludge thickener pump,  $T_{st}$ . As can be seen, the rate of this pump has no effect on the survivability probability of the system. This can be explained by the fact that this pump plays a secondary role comparing to the cleaning street pumps. Since the rates of pumps  $T_1$ ,  $T_2$ ,  $T_3$  are constant, increasing any pump rate which is placed *after* them does not change the overall capacity of the system.



**Figure 10:** Probability for the survivability property  $\Phi_A$  to hold, while varying (a) the capacity of community buffer ( $P_c$ ), (b) the rate of fresh sludge pump ( $T_{ft}$ ), (c) the rate of the cleaning street pumps ( $T_1$ ,  $T_2$  and  $T_3$ ), and (d) the rate of surplus sludge pump ( $T_{st}$ );  $\mu$  is the mean rain duration, distributed according to a truncated Normal distribution with variance one.

In order to generate each diagram in Figure 10, for each combination of parameters, one STD has to be generated, followed by a model checking procedure. Table 1, shows number of points in each diagram in Figure 10 and the overall computation time for producing that diagram, i.e., generating the STD and the model checking. As can be seen, even for this big case study, generating and model checking 1000 STDs takes less than a second! This clearly shows the value and efficiency of our method.

**Table 1: Overall computation time for generating results depicted in Figure 10.**

	Number of points	Computation time (ms)
Figure 10(a)	1000	522
Figure 10(b)	800	573
Figure 10(c)	600	269
Figure 10(d)	800	459


### Scenario B

For the second scenario, shown by the dashed box B, we consider a failure in the sand interceptor pump,  $T_z$ , modelled by the deterministic transition  $T_f$ , firing at time  $\alpha$ , which again could be parametrized for any arbitrary value. After the occurrence of a failure, a repair crew will repair the pump with a duration distributed according to an exponential distribution, with mean 2 hours. For this case we investigate almost the same formula as before, only now the recovery condition is that the pump should be repaired:

$$\Phi_B = (x_{P_o} < 0.01) \mathcal{U}^{[\alpha, \alpha+30]} (m_{P_r} = 1), \quad (2)$$

where,  $m_{P_r} = 1$ , means that the sand interceptor pump is repaired. Here, we have chosen the time bound  $[\alpha, \alpha + 30]$  for the Until operator, since the pump is supposed to be repaired within 30 hours after its failure.


For this scenario, we consider two parameters, the time of failure and the intake rate. The result is shown in Figure 11. On the  $x$ -axis the intake rate is parametrized from 6 to 13, and the  $y$ -axis represents different times of failure, from 30 minutes to 5 hours (right to left). As expected, for larger rates of the intake, the probability for survivability property  $\Phi_B$  to hold decreases. However, it is interesting that for a late occurrence of failure, the probability is lower, especially for high intake rates. The reason for this is that the capacity of the system is equal to the sum of the cleaning street rate pumps (4) and the fresh sludge thickener pump rate (1.25), which is 5.25. Therefore, the buffer is filling up for intake rates greater than 5.25, and a late failure will cause a quicker violation of the safety condition. On the other hand, for early failures, we have a non-zero survivability probability, even for high intake rates.



**Figure 11:** Probability of survivability property  $\Phi_B$  to hold for varying intake rate ( $x$ -axis) and failure time ( $y$ -axis).

Figure 12 provides a better understanding of this case. Each curve in this figure represents the survivability probability ( $y$ -axis) for a given fixed intake rate (color) to the

system; the horizontal axis depicts the failure times. The time that the probability hits zero is the very exact moment that the community buffer has become full, hence, if the failure occurs at any time after that, the surrounding area will be flooded immediately. This is the reason that this probability equals zero for any time of failure after this point.



**Figure 12: Probability for survivability property  $\Phi_B$  to hold. Each curve represents a specific intake rate, and the horizontal axis depicts the failure occurrence time.**

The two last figures show the importance of fast maintenance in bad weather conditions, otherwise, soon after the occurrence of a failure, flooding of the surrounding area is inevitable.

## 7. CONCLUSIONS

This paper evaluates the survivability of a waste water sewage facility using Hybrid Petri nets with a single general one-shot transition.

In order to capture the behaviour of the system, two new concepts have been added to the formalism, namely guard arcs and continuous dynamic transitions. These concepts allow to express the dependency of the system evolution on the amount of fluid in a continuous place and on the rates of continuous transitions, respectively.

Using the underlying stochastic time diagram and recent algorithms for model checking the logic STL, it is possible to analyse the survivability of the system for two different scenarios and a wide range of parameter settings. We were able to estimate the entire capacity and performance of the system, for different intake rates. Moreover, we evaluated the importance of several components of the system, and provided suggestions for tuning their characteristics. This casestudy clearly showed the strength of HPnGs in both modelling capabilities and efficiency of computations, for this application area.

## Acknowledgements

This work has been supported by the ROCKS project through the NWO grant DN 63-257. Anne Remke is funded by a NWO Veni grant.

## 8. REFERENCES

- [1] “RTV Oost. Overijssel Vandaag,” July 2013, <http://www.rtvooost.nl/tv/uitzendingemist.aspx?uid=290892>.
- [2] “TV Enschede FM. TV Enschede Nieuws,” June 2013, <http://www.youtube.com/watch?v=DRIB6JTNvhA>.
- [3] “UT Nieuws. Wanneer kun je kanon op de Auke Vleerstraat?” July 2013, <http://www.utnieuws.nl/studenten/wanneer-kun-je-kanoen-op-de-auke-vleerstraat>.
- [4] M. Gribaudo and A. Remke, “Hybrid Petri Nets with General One-Shot Transitions for Dependability Evaluation of Fluid Critical Infrastructures,” in *2010 IEEE 12th International Symposium on High Assurance Systems Engineering*. IEEE CS Press, Nov. 2010, pp. 84–93.
- [5] H. Ghasemieh, A. Remke, B. Haverkort, and M. Gribaudo, “Region-Based Analysis of Hybrid Petri Nets with a Single General One-Shot Transition,” in *Formal Modeling and Analysis of Timed Systems*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2012, vol. 7595, pp. 139–154.
- [6] H. Ghasemieh, A. Remke, and B. Haverkort, “Survivability evaluation of fluid critical infrastructures using hybrid Petri nets,” 2013, To appear in 19th IEEE Pacific Rim International Symposium on Dependable Computing.
- [7] L. Cloth and B. Haverkort, “Model checking for survivability!” in *Proceedings of the Second International Conference on the Quantitative Evaluation of Systems, 2005*. IEEE, 2005, pp. 145–154.
- [8] P. E. Heegaard and K. S. Trivedi, “Network survivability modeling,” *Computer Networks*, vol. 53, no. 8, pp. 1215–1234, 2009.
- [9] J. C. Knight and K. Sullivan, “On the definition of survivability,” University of Virginia, Tech. Rep., 2000.
- [10] Dynamic simulation software for biological wastewater treatment modelling, ‘<http://holinger.com/index.php?id=748&l=10&type=98>’.
- [11] M. Faber and M. Stewart, “Risk assessment for civil engineering facilities: critical overview and discussion,” *Reliability Engineering and System Safety*, vol. 80, no. 2, pp. 173 – 184, 2003.
- [12] J. Derco, L. Cernochova, L. Krcho, and A. Lalai, “Dynamic simulations of waste water treatment plant operation,” *Chemical Papers*, vol. 65, no. 6, pp. 813–821, 2011.
- [13] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis, *Modelling with Generalized Stochastic Petri Nets*, 1st ed. John Wiley & Sons, Inc, 1994.
- [14] R. David and H. Alla, *Discrete, Continuous, and Hybrid Petri Nets*, 2nd ed. Springer Berlin Heidelberg, 2010.
- [15] M. Gribaudo and A. Remke, “Hybrid Petri nets with general one-shot transitions: model evolution,” University of Twente, Tech. Rep., 2010, <http://wwwhome.cs.utwente.nl/~anne/techreport/hpng.pdf>.
- [16] M. A. Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis, *Modelling with Generalized Stochastic Petri Nets*. John Wiley & Sons, Inc, 1995.
- [17] S. M. Khopkar, *Environmental Pollution Monitoring and Control*. New Age International, 2004.
- [18] “Primer for municipal wastewater treatment systems, ‘<http://www.epa.gov/npdes/pubs/primer.pdf>,’” 2004.