State-space Construction of Hybrid Petri Nets with Multiple Stochastic Firings

JANNIK HÜLS, CARINA PILCH, PATRICIA SCHINKE, HENNER NIEHAUS,
JOANNA DELICARIS, and ANNE REMKE, Westfälische Wilhelms-Universität Münster

Hybrid Petri nets have been extended to include general transitions that fire after a randomly distributed amount of time. With a single general one-shot transition the state space and evolution over time can be represented either as a *Parametric Location Tree* or as a *Stochastic Time Diagram*. Recent work has shown that both representations can be combined and then allow multiple stochastic firings. This work presents an algorithm for building the *Parametric Location Tree* with multiple general transition firings and shows how its transient probability distribution can be computed using multi-dimensional integration. We discuss the (dis-)advantages of an interval arithmetic and a geometric approach to compute the areas of integration. Furthermore, we provide details on how to perform a Monte Carlo integration either directly on these intervals or convex polytopes, or after transformation to standard simplices. A case study on a battery-backup system shows the feasibility of the approach and discusses the performance of the different integration approaches.

CCS Concepts: • **Theory of computation** → *Verification by model checking*;

Additional Key Words and Phrases: Petri nets, stochastic hybrid model, multi-dimensional integration, transient probability

ACM Reference format:

Jannik Hüls, Carina Pilch, Patricia Schinke, Henner Niehaus, Joanna Delicaris, and Anne Remke. 2021. State-space Construction of Hybrid Petri Nets with Multiple Stochastic Firings. *ACM Trans. Model. Comput. Simul.* 31, 3, Article 13 (August 2021), 37 pages.
<https://doi.org/10.1145/3449353>

1 INTRODUCTION

Hybrid Petri nets with general transitions (HPnG) [24] extend Hybrid Petri nets [1] by adding stochastic behavior through general transitions with a randomly distributed delay. HPnGs provide a high-level formalism for a class of stochastic hybrid systems with piece-wise linear continuous behavior without resets and a probabilistic resolution of non-determinism. Hybrid Petri nets have been shown useful for the evaluation of critical infrastructures [21, 35], even though they have previously been restricted to a single random variable. This article shows how the state-space of a Hybrid Petri net with multiple general transition firings can be constructed as a **Parametric Location Tree (PLT)** [24] for a predefined maximum time. As such, this article provides the missing

Authors' addresses: J. Hüls, C. Pilch, P. Schinke, H. Niehaus, J. Delicaris, and A. Remke, Westfälische Wilhelms-Universität Münster; emails: {jannik.huels, carina.pilch, patricia.schinke, henner.niehaus, joanna.delicaris, anne.remke}@wwu.de.

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1049-3301/2021/08-ART13 \$15.00

<https://doi.org/10.1145/3449353>

link to recent work [33, 34] that provides model checking capabilities for HPnGs with multiple random variables, assuming the existence of the PLT.

The approach of this article is twofold: First, a purely numerical iterative algorithm for the construction of a PLT in the presence of a finite but arbitrary number of stochastic firings before a certain maximum time is presented. This algorithm is based on introducing an order on the random variables that occur due to the firing of general transitions. Each of these firings affects the evolution of the Petri net. The idea of parametric analysis, as presented in Reference [24], is to collect values of random variables into intervals, which lead to a similar system evolution. These intervals are called potential domains. Second, to facilitate a transient analysis, we identify those parametric locations the HPnG can be in at a specific point in time (the so-called candidates) and compute those subsets of the domains of the random variables present, for which the HPnG is guaranteed to be in a specific candidate. A multi-dimensional integration over these so-called time-restricted domains yields a transient distribution over locations.

The evolution of the state space over time of an HPnG can be partitioned into sets of states with similar behavior. Two approaches exist for representing the time-bounded system executions symbolically, by conditioning their evolution on the firing times of the general transitions: One method creates a **Parametric Location Tree (PLT)**, which symbolically represents sets of states as *locations* in the nodes of the tree and events between states as edges between such nodes. The other method uses a geometric representation for sets of states in terms of convex polytopes (so-called regions) with similar characteristics [34]. The idea of a polyhedra-based representation of the state space has been explored before, e.g., for (flowpipe) approximations [14, 15] and to abstract uncountable-state stochastic processes [44, 45].

Both representations are symbolic, having different (dis-)advantages. The PLT-representation can be computed efficiently and lends itself well to compute simple measures like transient distributions. The purely geometric representation as presented in Reference [23] is computationally more expensive, as it heavily relies on hyperplane arrangement, for which to the best of our knowledge no efficient implementations for more than three dimensions exist. However, it is possible to compute a geometric representation for each location in a PLT [34] that fully relies on halfspace intersection, which can be solved e.g., using quickhull [4] also for higher dimensions. This combines the computational speed of the parametric reachability analysis with the ease that a geometric representation brings for model checking complex properties [33].

The current article extends the previously published version [32] by a thorough discussion on how the integration can be performed either on the restricted domains or on the equivalent geometric representation as regions that have to be intersected with a time constraint. The resulting arbitrary convex polytopes and their representation need to be adapted for the computation of the corresponding probabilities using multi-dimensional integration. This can be done by performing triangulation on these convex polytopes, resulting in a set of simplices. These are then affinely transformed into standard simplices, which allow the use of different multi-dimensional integration techniques. Hence, the change of representation ensures the extensibility of the presented approach to different integration methods.

The multi-dimensional integration in this article is performed using Monte Carlo methods, either over these standard simplices or directly on convex polytopes to compute transient probabilities. As the integration approach detailed in this article is applicable to arbitrary convex polytopes, it can also be applied to convex polytopes that form the satisfaction set of an STL formula on a Hybrid Petri net, as described in Reference [33].

A scalable case study on a battery backed-up system in the presence of short-term peak loads and grid failures shows the feasibility of the approach and compares the performance of the different integration methods, as implemented in the tool HPNMG [31].

Related Work. Hybrid Petri nets without general transitions form a subclass of non-initialized singular automata [1]. Hence, unbounded reachability is not decidable and References [22, 24] resort to computing time-bounded reachability—however, only in the presence of at most two random variables. Previous work had to restrict the number of general transition firings, since (i) identifying the potential domains of child locations requires linear optimization, (ii) transient analysis relies on multi-dimensional integration over polytopes. For one random variable, this was solved using the simplex method and discretization [24]. For two random variables, Reference [22] used hyperplane arrangement and triangulation to compute numerically exact results. Vertex enumeration was proposed in Reference [34] to construct a PLT for an arbitrary but finite number of stochastic firings, but did not provide a general algorithm. The strict order on the firings of the general transitions, as in this article, simplifies the computation of the child locations to solving linear inequalities. This allows the implementation of (i) an efficient iterative algorithm for constructing the PLT up to a certain maximum time and (ii) a transient analysis based on *Fourier-Motzkin* variable elimination.

The MODEST toolset [28] provides an extensible and comprehensive toolset for analytical and statistical model checking. It allows numerous model formalisms and in contrast to the approach presented in this article, can also deal with non-determinism. Statistical model checking for stochastic hybrid models has been presented in Reference [7] and for HPnGs in Reference [39] with various confidence intervals and hypothesis tests. A limited overapproximating model checker for stochastic hybrid automata has also been released as part of MODEST [26]. SPACEEx [15] and FLOW* [8] provide reachability analysis for non-linear hybrid automata without stochasticity. Also constraint solvers have been proposed for the automated analysis of probabilistic hybrid automata [16]. Approaches for hybrid automata extended with discrete probability distributions [36, 46, 47, 51] compute time-bounded reachability using abstraction. The tools FAUST² [45] and StochHy [44] provide model checking capabilities for discrete-time stochastic hybrid models via abstractions to Markov decision processes. While our approach is limited to a restricted class of stochastic hybrid automata, the resulting areas of integration can be computed without approximation or abstraction, both for transient and for reachability analysis. Approximation is only introduced via the multi-dimensional integration, for which an error estimate can be computed.

Modeling and analysis of reactive timed systems based on various classes of Petri nets has been proposed in Reference [6]. CPN TOOLS [13] is a toolset for simulating and analyzing **Coloured Petri nets (CPNs)**, for which extensions to **Stochastically and Dynamically Coloured Petri nets (SDCPNs)** exist [12]. Both approaches allow for quantitative evaluation and qualitative verification for systems with a discrete notion of state. Another variant of Petri nets with hybrid and stochastic behavior are **Fluid Stochastic Petri nets (FSPNs)** [25, 29], which only allow exponentially distributed and immediate discrete transitions. Related Petri net approaches are all restricted w.r.t. the number of continuous variables [29] or to Markovian jumps [12]. MÖBIUS [11] supports modeling formalisms such as **Stochastic Petri Nets (SPNs)**, **Stochastic Automata Networks (SANs)**, and Markov Chains and mainly implements discrete-event simulation. HPnGs instead allow for general probability distributions, as the analysis abstracts from the specific distribution as long as possible. Furthermore, the number of continuous variables in the model does not restrict the analysis capabilities of our approach. While the approach presented for HPnGs is also general w.r.t. the number of random variables present in the system, its performance clearly suffers with a growing number of dimensions.

Semi-Markov processes have been used to evaluate the dependability of **uninterrupted power supply (UPS)** systems [50] and Stochastic Activity networks have been used to study the resilience of smart grid distribution networks [2]. They are, however, restricted to negative exponential stochastic behavior.

Organization. This article is further organized as follows: Section 2 recalls the syntax and semantics of the modeling formalism, Section 3 repeats the Parametric Location Tree. Section 4 introduces its construction for an arbitrary but finite number of random variables. Transient analysis via interval arithmetic is explained in Section 5, and Section 6 introduces a geometric approach for the computation of the areas of integration. Section 7 presents a feasibility study and compares the performance of the different integration approaches, taking into account CPU and memory usage. Section 8 concludes the article.

2 THE MODEL FORMALISM OF HPNGS

HPnGs extend hybrid Petri nets [1] with so-called general transitions [24]. We present their syntax in Section 2.1 and the model evolution together with a definition of states and events in Section 2.2. A small example is provided to illustrate the notation.

2.1 Syntax

Hybrid Petri nets with multiple general transition firings (HPnG) consist of three components: *places*, *transitions*, and *arcs*, where places are either *discrete* or *continuous*. A discrete place holds a natural number of tokens and a continuous place contains a non-negative real amount of fluid. Note that in the following definition and the remainder of this article, $\mathbb{N}_0$ indicates the set of natural numbers including zero and $\mathbb{R}_0^+$ indicates the set of non-negative real numbers. Transitions that change the amount of tokens in discrete places are either *deterministic*, *immediate*, or *general*, i.e., they fire after a random delay. Together they are referred to as discrete transitions, as they only change the discrete marking. This is in contrast to continuous transitions, which only change the continuous marking.

Definition 2.1. A **Hybrid Petri net with general transitions (HPnG)** is a tuple $(\mathcal{P}, \mathcal{T}, \mathcal{A}, \mathbf{M}_0, \phi)$.

- $\mathcal{P} = \mathcal{P}^d \cup \mathcal{P}^c$ is a finite set of *places*, partitioned into disjoint sets of *discrete* and *continuous* places. Any discrete place $P_i^d \in \mathcal{P}^d$ holds a marking of $m_i \in \mathbb{N}_0$ tokens and any continuous place $P_j^c \in \mathcal{P}^c$ holds a continuous marking described by a fluid level $x_j \in \mathbb{R}_0^+$.
- A continuous place has a predefined upper bound, possibly infinity. The lower bound is always zero.
- The initial marking $\mathbf{M}_0$ is given by the initial number of tokens $\mathbf{m}_0$ and fluid levels $\mathbf{x}_0$ of all places.
- Transitions in the finite set $\mathcal{T}$ change the marking upon firing, i.e., they change the tokens or fluid level of discrete and continuous places.
- Places and transitions are connected via arcs in the finite set $\mathcal{A}$. Discrete arcs in $\mathcal{A}^d$ connect discrete places to discrete transitions and vice versa. Correspondingly, continuous arcs in $\mathcal{A}^c$ connect continuous places and continuous transitions. How the marking is changed is defined by assigning weights and priorities to those arcs. Depending on the direction of connection, places are called input or output places of transitions.
- A discrete transition in $\mathcal{T}^D \cup \mathcal{T}^I \cup \mathcal{T}^G$ is either deterministic, immediate, or general. It is enabled when its input places match or exceed the weight of their connecting arcs. Discrete transitions are each associated with a clock c_i , which if enabled evolves with $\frac{dc_i}{dt} = 1$, otherwise $\frac{dc_i}{dt} = 0$. Note that, upon disabling, the clock value is preserved. For general transitions this corresponds to the preemptive resume strategy. A deterministic transition $T_k^D \in \mathcal{T}^D$ fires when c_i reaches the predefined transitions firing time. For an immediate transition $T_k^I \in \mathcal{T}^I$ the predefined firing time is always zero. The firing time of a general

Fig. 1. Hybrid Petri net with general transition of a reservoir from Reference [24]. The continuous place P_0^c models the reservoir, which has a deterministically timed demand of 5 time units, modeled by the deterministic transition T_0^D , and a generally distributed pump breaking time, modeled by the general transition T_0^G . The initial marking of P_0^c is zero and for the discrete places as shown in the figure. All weights and priorities equal 1.

transition $T_m^G \in \mathcal{T}^G$ is modelled by a **cumulative distribution function (CDF)**, which is assumed to be absolutely continuous. Each *stochastic firing* results in a random variable that follows the CDF of the general transition. By construction, these random variables are independent of each other, since the firing behavior of the general transitions is solely described by its CDF.

- A continuous transition is enabled if all connected input places hold fluid and fires at its nominal rate if the connected input and output places are not at either boundary. Every static continuous transition $T_n^F \in \mathcal{T}^F$ has a constant nominal flow rate. Dynamic continuous transitions $T_o^{Dyn} \in \mathcal{T}^{Dyn}$ represent a set $D \subset \mathcal{T}^F$ of static continuous transitions. Hence, their nominal flow rate is a function of the actual flow rates of all static continuous transitions in D [21].
- The rates of continuous transitions that are connected to a continuous place that is at either of its boundaries require *rate adaptation* [24], changing the actual flow rate.
- Guard arcs $A^t \in \mathcal{A}^t$ may further influence the enabling of transitions. They carry a comparison operator and a weight. A discrete guard arc connects a discrete place to any transition and conditions its enabling on the comparison of the discrete marking and its weight. Correspondingly, continuous places may be connected to (only) discrete transitions via (continuous) guard arcs.
- All parameters are given by the parameter function ϕ . For details on ϕ , refer to Reference [24].

With the above definition of arcs, this results in bipartite Hybrid Petri nets, where places form one set and transitions form another set. Furthermore, without guard arcs, the discrete parts of the net and the continuous parts of the net are separate. Figure 1 presents an example for an HPnG, which models a water reservoir. Places are illustrated by circles, transitions by bars, and arcs by arrows.

2.2 Model evolution

Next to the syntax of a hybrid Petri net as defined above, a couple of important concepts exist, which define the evolution of the model over time. These concepts comprise the enabling and firing rules, the sharing of fluid, conflict resolution, and rate adaptation. We give a brief overview on these concepts but refer to Reference [24] for further details.

For continuous transitions that are connected to places that are at either boundary, a so-called *actual rate* is computed by *rate adaptation*: For places that are at the upper boundary, the inflow of continuous transitions is decreased to match the outflow. For places at the lower boundary, the outflow of continuous transitions is reduced accordingly. A continuous place then evolves with a drift, which equals the sum of the actual inflow rates minus the sum of the actual outflow rates.

It is possible that multiple immediate or deterministic transitions are scheduled to fire at the same time, this so-called *conflict* is resolved using priorities and weights. For details on *conflict resolution* and *the concept of enabling*, we refer to Reference [24]. The probability that a general transition fires at the same point in time at which a deterministic transition fires is zero. We exclude Zeno behavior by banning cycles of immediate and general transitions. Section 3 further describes the model evolution and the interplay between stochastic and deterministic transitions.

A definition of state is required to formally define the evolution of the model.

Definition 2.2. The state of an HPnG is a tuple $\Gamma = (\mathbf{m}, \mathbf{x}, \mathbf{c}, \mathbf{d}, \mathbf{g}, \mathbf{e})$, where $\mathbf{m}$ is the discrete marking, $\mathbf{x}$ the continuous marking, $\mathbf{c}$ contains for each deterministic transition the enabling time. The drift $\mathbf{d}$ describes the current change of fluid level per time unit for each continuous place, $\mathbf{g}$ contains the enabling time for the general transitions, and $\mathbf{e}$ describes the enabling status of all transitions.

Events trigger changes in states as introduced in References [24, 40]:

Definition 2.3. An event $\Upsilon(\Gamma_i, \Gamma_{i+1}) = (\Delta\tau_{i+1}, \varepsilon_{i+1})$ describes the change from one state Γ_i to another state Γ_{i+1} , with $\varepsilon_{i+1} \in \mathcal{P}^c \cup \mathcal{T}^I \cup \mathcal{T}^D \cup \mathcal{T}^G \cup \mathcal{A}^t$ specifying the model element that caused the event. Note that $\Delta\tau_{i+1} \in \mathbb{R}_0^+$ is a relative time between two events, such that one of the following conditions is fulfilled:

- (1) An immediate, deterministic or general transition fires, such that $\mathbf{m}_i \neq \mathbf{m}_{i+1} \wedge \varepsilon_{i+1} \in \mathcal{T}^I \cup \mathcal{T}^D \cup \mathcal{T}^G$.
- (2) A continuous place reaches its lower or upper boundary, such that $\mathbf{d}_i \neq \mathbf{d}_{i+1} \wedge \varepsilon_{i+1} \in \mathcal{P}^c$.
- (3) A guard arc condition is fulfilled or stops being fulfilled, such that $\mathbf{e}_i \neq \mathbf{e}_{i+1} \wedge \varepsilon_{i+1} \in \mathcal{A}^t$.

The set of all possible events that can occur in state Γ_i is finite and its size depends on the number of continuous places, the number of guard arcs, and the number of enabled discrete transitions. It is denoted $\mathcal{E}(\Gamma_i)$ and the set of events with minimum remaining time for that state is defined as follows:

$$\mathcal{E}^{\min}(\Gamma_i) = \left\{ \Upsilon_j(\Gamma_i, \Gamma_j) \in \mathcal{E}(\Gamma_i) \mid \nexists \Upsilon_k(\Gamma_i, \Gamma_k) \in \mathcal{E}(\Gamma_i) : \Delta\tau_k < \Delta\tau_j \right\}. \quad (1)$$

Note that multiple events can happen at the same point in time, e.g., due to conflicts between deterministic transitions or due to the scheduling of stochastic transitions. We split the set $\mathcal{E}^{\min}(\Gamma_i)$ into two subsets, representing the next random events $\mathcal{E}_{\text{ran}}^{\min}(\Gamma_i)$ and the set of all next deterministic events $\mathcal{E}_{\text{det}}^{\min}(\Gamma_i) = \mathcal{E}^{\min} \setminus \mathcal{E}_{\text{ran}}^{\min}(\Gamma_i)$. The size of $\mathcal{E}_{\text{ran}}^{\min}(\Gamma_i)$ is bounded by the number of enabled general transitions and the size of $\mathcal{E}_{\text{det}}^{\min}(\Gamma_i)$ is bounded by the number of continuous places plus the number of guard arcs plus the number of enabled deterministic and immediate transitions. The time at which the next events may happen can always be determined only depending on the values of the given random variables. The next minimum event time is defined as the minimum time to the next event(s), which can be the same for two or more events. Also, before the first general transition firing, it is unique and simply is the minimum of the remaining times of all upcoming deterministic events, i.e., deterministic or immediate transition firings, guard arc en-/disablings or continuous boundary events. After at least one stochastic firing, the occurrence of events, the clocks, and the continuous marking may linearly depend on the random variable(s). Since an arbitrary but finite number of random variables is present in the system and the computation of the particular next

minimum event (for all event types) may depend on all past transition firing times, it is based on polynomials and not scalars. Hence, after at least one stochastic firing the set of deterministic events with minimum remaining time may consist of more than one element. Their remaining time to fire is derived by minimizing the elements of $\Delta T_{\text{det}}^{\min}(\Gamma_i)$ over the support of the random variables that already have fired, where

$$\Delta T_{\text{det}}^{\min}(\Gamma_i) = \left\{ \Delta \tau_m \in \mathbb{R}_0^+ \mid \exists Y(\Gamma_i, \Gamma_m) \in \mathcal{E}_{\text{det}}^{\min}(\Gamma_i) \right\}. \quad (2)$$

One of the main contributions of this article is to create a **Parametric Location Tree (PLT)** for multiple general transition firings. The PLT symbolically represents all states, in which an HPnG can be before some predefined maximum time.

3 THE PARAMETRIC LOCATION TREE

The evolution over time up to a fixed maximum time $\tau_{\max}$ of an HPnG can be described as a so-called Parametric Location Tree (PLT) [24, 34]. Each node in the tree is called a parametric location and represents a set of potentially uncountably many states.

The edges of a PLT represent events. When computing the PLT for a predefined time bound $\tau_{\max}$, the number of parametric locations is bounded by the number of events, since we exclude Zeno behavior through the model construction. The PLT then contains by construction all possible execution traces in a symbolic representation, which relies on the random variables modeling the firing times of general transitions. For predefined intervals of the support of random variables (i.e., their domain of definition), every location can be reached at an entry time $\Lambda.t$. Choosing an assignment for all random variables from the support of the joint probability distribution function defines a deterministic execution path throughout the PLT. For each time point of this evolution, the system state is unique and is symbolically represented by a specific location.

Hence, every location contains a symbolic representation of an uncountable set of states, where each of these states can be reached for a specific assignment for all random variables and the evolution time. This concept also holds for the entry time and the entry state of a location, which can both only be uniquely specified for a fixed assignment of the random variables. Note that both the support of the random variables and the evolution time are subject to constraints that are computed throughout the construction of the PLT. This is further explained in Section 4.

The event corresponding to the edge is called *source event* w.r.t. the child node location. Extending previous notation, a parametric location is defined as follows:

Definition 3.1. A parametric location Λ is defined as a tuple $\Lambda = (ID, t, p, \Gamma, S, \mathbf{o})$, where the entry time is $\Lambda.t$ and the probability of choosing that specific location with identifier $\Lambda.ID$ is $\Lambda.p$ w.r.t. the parent node. This probability $\Lambda.p$ is also denoted as conflict probability, as it equals one if no conflict occurs and in the case of conflicting events, i.e., multiple immediate or multiple deterministic transitions of the same priority are supposed to fire at the same time, the probability is derived from the weights of the conflicting transitions [24]. $\Lambda.\Gamma$ denotes the state of the HPnG when entering the location and $\Lambda.S$ denotes the potential domain of the random variables, as described in the following: There can be multiple general transitions present in the system and each can possibly fire multiple times. For each parametric location the firing order that has led to this location is unique and stored as vector $\Lambda.\mathbf{o}$.

Each firing of a general transition corresponds to a random variable that equals the enabling time between two consecutive firings. The random variable corresponding to the j th firing of T_i^G is denoted s_i^j and is added to $\mathbf{o}$ upon the firing of the transition. The r th stochastic firing is then stored in $\mathbf{o}[r]$. A new random variable is instantiated each time a general transition becomes enabled for the first time after a previous firing. Hence, the number of random variables n equals the number

of stochastic firings m plus the number of general transitions that are currently enabled, thus $m \leq n$. At a given point in time, a random variable that corresponds to a past firing is said to be *expired*. Accordingly, we define a vector $\mathbf{s}_\Lambda$ of all (expired and not yet expired) random variables for each location Λ , where the expired random variables are ordered in $\mathbf{s}_\Lambda$ in the same order as specified by $\Lambda.\mathbf{o}$, i.e., $\forall k \in [1, m] : \mathbf{s}_\Lambda[k] = \Lambda.\mathbf{o}[k]$. Concurrently enabled general transitions yield competing random variables of potentially different distributions, whereas consecutive firings of a single general transition result in a series of identically distributed random variables.

The firing of general transitions at different times determine the evolution of the Petri net. The idea of parametric analysis, as presented in Reference [24], is to collect values of random variables into intervals $\mathbf{S}$ for each random variable s_i^j , which lead to the same next event. These intervals are called potential domains and need to be specified for each parametric location. Hence, the potential domain $\Lambda.\mathbf{S}$ contains a dedicated vector $\mathbf{S}_i$ for each general transition T_i^G . Each of these vectors then collects one interval for each previous firing j , stored in $\mathbf{S}_i^j$, and one additional interval for a potential next firing. An interval $\mathbf{S}_i^j$ has a lower and an upper boundary $\mathbf{S}_i^j.l \leq \mathbf{S}_i^j.u$, such that location Λ is reached if for every j , $s_i^j \in [\mathbf{S}_i^j.l, \mathbf{S}_i^j.u]$. This extends the definition presented in Reference [34] by separately collecting the possible intervals for each general transition. Furthermore, this defines a strict order on the dependencies of each random variable. As the number of stochastic firings differs per location, the size of the potential domain, as indicated by $|\mathbf{S}|$, also differs.

All parametric locations, i.e., all symbolic representations of states, of an HPnG are collected in a so-called parametric location tree as defined in the following:

Definition 3.2. The PLT of an HPnG is defined as a tree $(\mathbf{V}, \mathbf{E}, v_{\Lambda_0})$, where $\mathbf{V}$ is the set of nodes representing the parametric locations of the HPnG. $\mathbf{E}$ is the set of edges with $e_i = (v_{\Lambda_j}, v_{\Lambda_k}) \in \mathbf{E}$ for $v_{\Lambda_j}, v_{\Lambda_k} \in \mathbf{V}$ if an event $\Upsilon(\Lambda_j.\Gamma, \Lambda_k.\Gamma)$ exists that leads from the parametric location Λ_j to its child location Λ_k . The root node v_{Λ_0} represents the initial location, initialized as follows: $\Lambda_0 = (0, 0, 1, \Gamma_0, \mathbf{S}^\infty, \mathbf{o}_0)$, where $\mathbf{S}^\infty$ denotes the not yet limited domain for all currently enabled random variables. Since initially no general transitions has fired, the vector $\mathbf{o}_0$ is empty.

The PLT is iteratively constructed by adding a child location Λ_c for each possible event that can take place from a given parametric location. The finite set of possible next events (see Equation (1)) together with the exclusion of cycles of immediate and general transitions in the model definition yields a finite PLT if computed up to time $\tau_{\max}$.

The absolute point in time at which the event takes place is then stored as the locations entry time $\Lambda_c.t$. The other parameters of the location are derived by executing the event in the HPnG. The random variable domain for a specific child location is derived from the parent location by taking into account the event that leads to the specific location. Since the location entry times may depend on previously expired random variables, the set that consists of the remaining minimum times for all deterministic events $\Delta T_{\det}^{\min}(\Lambda_p.\Gamma)$ may contain more than one element and the order of the timed events depends on the values of the random variables that have expired. Since different child nodes symbolically describe sets of values leading to different system evolutions, the resulting potential domains are always disjoint for non-conflicting successors.

Figure 2 shows the resulting state space for the reservoir example from Figure 1. Figure 2(a) shows the PLT, and Figure 2(b) provides a geometrical representation of the same state-space, i.e., the Stochastic Time Diagram [22], which allows to illustrate a unique path that corresponds to the evolution of time for a given assignment of the random variable, e.g., the dashed vertical line that starts for $s_0^0 = 8$ and crosses through locations ℓ_0, ℓ_2, ℓ_5 and ℓ_8 .

Fig. 2. Parametric Location Tree and geometric representation as Stochastic Time Diagram [22] corresponding to the HPnG in Figure 1. For each location Λ the illustration of the PLT contains the ID , the entry time $\Lambda.t$, and the domain of the random variable s_0^0 . The Stochastic Time Diagram illustrates the relation of the possible values for s_0^0 and the global time t' for each location. Choosing a specific assignment for the random variable, i.e., $s_0^0 = 8$, leads to a unique path that corresponds to the dashed line indicated in (b). The parametric locations that are passed with this specific evolution are highlighted with a dashed border in the PLT representation in (a).

4 PLT CONSTRUCTION

When allowing more than one stochastic firing the construction of the PLT becomes more complex [34]. For each enabled general transition, at least one location is added as child node, where the edge represents its next firing before every other event. Additionally, each general transition needs to be scheduled after each minimum deterministic event. Hence, if there are n deterministic events, then for each enabled general transition, first $n+1$ child nodes are generated, with one edge for the general transition firing and n for all possible deterministic next events $\mathcal{E}_{\text{det}}^{\min}(\Gamma_i)$. The deterministic next events and the enabled general transitions can then be arranged in two different ways: (i) either deterministic events are scheduled first (ii) or a general transitions fires before the next deterministic events. We need to consider all possible combinations of deterministic events and general transitions firings, and we need to adjust the boundaries of the potential domains in all cases accordingly, which is discussed in the following:

First, Section 4.1 explains how the intervals of expired random variables are limited, such that the locations for the deterministic events, whose occurrence time might depend on those random variables, can be correctly scheduled. Section 4.2 then describes the scheduling of the child locations in detail.

4.1 Adjust Boundaries for Expired Random Variables

The support of each random variable is partitioned into the potential domains, as explained in Section 3. These intervals S_i^j contain possible values for every random variables s_i^j per location. The bounds of those intervals may depend on firings with a lower order. When scheduling the deterministic successors of a location, we first need to compute the set of minimum events. In case this set contains more than one element, their order depends on the values of other random variables. This is handled by adjusting the potential domain for the corresponding random variables

per event, such that the event takes place before all other minimum events for all values in the (adjusted) potential domain of the corresponding successor location.

The computation of the correct potential domain corresponds to minimizing the time to the minimum event over the domain of all random variables. In earlier work this has been solved, e.g., using the simplex method [24] in case of one random variable, hyperplane arrangement in the case of two [22] or as suggested by Reference [34] using vertex enumeration. Instead, this article defines a strict total order $<$ on the random variables that is based on their firing order: Let s_i^j and s_k^l be two random variables, then $s_i^j < s_k^l$ holds iff the i th general transition fires for the j th time before the k th general transition fires for the l th time:

$$\exists u, v \in \mathbb{N} : \mathbf{o}[u] = s_i^j \wedge \mathbf{o}[v] = s_k^l \wedge u < v \Rightarrow s_i^j < s_k^l. \quad (3)$$

This ensures that the interval bounds of a random variable s_k^l may only depend on earlier expired random variables s_i^j . The interval S_k^l of a random variable s_k^l can in general be described by two linear expressions:

$$s_k^l \in S_k^l = [S_k^l.l, S_k^l.u] = \left[a_0 + \sum_{s_i^j < s_k^l} a_i^j \cdot s_i^j, b_0 + \sum_{s_i^j < s_k^l} b_i^j \cdot s_i^j \right], \text{ with } a_0, a_i^j, b_0, b_i^j \in \mathbb{R}. \quad (4)$$

Each interval bound is given by a real constant plus the sum of previously expired random variables, each potentially multiplied by a real constant. In the following, we describe how these bounds are computed.

To schedule the deterministic successors of a parametric location Λ_p , the set of minimum events $\mathcal{E}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$ is required. When generating a child location Λ_c for a successor that results from a deterministic event $\Upsilon(\Lambda_p.\Gamma, \Lambda_c.\Gamma) \in \mathcal{E}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$, its potential domain is determined as follows: In a first step, the potential domain of Λ_c is set to the one of Λ_p . If the set $\Delta\mathbf{T}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$ contains more than one value, then the intervals for the expired random variables in the potential domain $\Lambda_c.S$ have to be limited for the corresponding deterministic successor, such that its source event $\Upsilon(\Lambda_p.\Gamma, \Lambda_c.\Gamma) = (\Delta\tau_c, \varepsilon_c)$, takes place before any other event. Thus, $\Delta\tau_c$ is required to be smaller than or equal to the remaining time for any other event:

$$\forall \Delta\tau^* \in \Delta\mathbf{T}_{\text{det}}^{\min}(\Lambda_p.\Gamma) : \Delta\tau_c \leq \Delta\tau^*. \quad (5)$$

Hence, the potential domain of the random variables in Λ_c is adjusted by comparing $\Delta\tau_c$ in a pairwise fashion to each other value $\Delta\tau^* \in \Delta\mathbf{T}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$ and limiting the intervals of the corresponding random variables such that Equation (5) holds for the affected potential domains. Having $\Delta\tau_c = \Delta\tau^*$ means that both linear expressions intersect. Note that this results in closed intervals for the potential domains with overlapping interval bounds, which, however, does not make a difference, since the probability for such case equals zero.¹

The computation can be done by considering only random variables that correspond to firings in the past, i.e., with a lower order, since only past firings can affect clocks and markings and the remaining time to fire, as stored in $\Delta\mathbf{T}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$. Since $\Delta\tau_c$ is compared to the other remaining times one-by-one, we first present the pairwise comparison in the following and then explain how the results are brought together to obtain the potential domain for the child location Λ_c . The remaining times to fire, $\Delta\tau_c$, and each $\Delta\tau^* \in \Delta\mathbf{T}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$ for $\alpha_0, \alpha_z, \beta_0, \beta_z \in \mathbb{R}$ and $m = |\mathbf{o}|$, can

¹Working with only closed intervals facilitates the technical representation of the intervals in the implementation.

be written as:

$$\Delta\tau_c = \alpha_0 + \sum_{z=1}^m \alpha_z \cdot \mathbf{o}[z], \text{ and } \Delta\tau^* = \beta_0 + \sum_{z=1}^m \beta_z \cdot \mathbf{o}[z]. \quad (6)$$

Note that the coefficients α_i follow from the computation of the source event leading to this parametric location Λ_c . Specifically, the time to this event $\Delta\tau_c$ is computed according to the event type, defined in Section 2. Accordingly, the coefficients β_i follow from the events in the set $\mathcal{E}_{\text{det}}^{\min}(\Lambda_p, \Gamma)$.

According to Equation (4), the expressions in Equation (6) are also dependent on the random variables in firing order. Each variable may again be multiplied by a real constant, whose computation is presented in the following: The interval for any random variable $\mathbf{o}[z]$ with the same multiplier in both linear expressions, i.e., $\alpha_z = \beta_z$, does not need to be adapted, since the validity of $\Delta\tau_c \leq \Delta\tau^*$ is independent of this random variable.

Using these linear expressions, the inequality $\Delta\tau_c \leq \Delta\tau^*$ can be solved for a specific random variable, i.e., the random variable with the highest order for which the multipliers differ in both inequalities. Hence, to ensure that the condition $\Delta\tau_c \leq \Delta\tau^*$ is fulfilled, it is sufficient to limit the intervals of the potential domain of this specific random variable, such that the rearranged inequality (that is solved for this random variable) is fulfilled. Precisely, we determine the maximum index k with respect to the order $<$, for which $\alpha_k \neq \beta_k$ holds. If no such index k exists, then it follows that $\Delta\tau_c = \Delta\tau^*$, which indicates two events with the same remaining occurrence time. In this case, only the event with the higher order (see Table 1 in Reference [24]) is considered. If both events have the same order, then the conflict resolution (see Section 3.4 in Reference [24]) is used to compute probabilities and $\Lambda_c.p$ is updated accordingly and the potential domain does not need to be adjusted. Due to the definition of $\mathcal{E}_{\text{det}}^{\min}(\Gamma_i)$, if such a maximum index k exists, then it needs to be larger than zero. Otherwise, the linear expressions of the remaining times $\Delta\tau_c$ and $\Delta\tau^*$ would be parallel functions and one of the remaining event times is always larger and hence the corresponding event is not included in $\mathcal{E}_{\text{det}}^{\min}(\Gamma_i)$. For any $k > 0$, the potential domain of the random variable $\mathbf{o}[k]$ is adjusted by rearranging the inequality $\Delta\tau_c \leq \Delta\tau^*$ as:

$$\alpha_0 + \sum_{z=1}^m \alpha_z \cdot \mathbf{o}[z] \leq \beta_0 + \sum_{z=1}^m \beta_z \cdot \mathbf{o}[z] \Leftrightarrow (\alpha_k - \beta_k) \cdot \mathbf{o}[k] \leq (\beta_0 - \alpha_0) + \sum_{z=1}^{k-1} (\beta_z - \alpha_z) \cdot \mathbf{o}[z]. \quad (7)$$

First $\Delta\tau_c$ as well as $\Delta\tau^*$ are replaced by the expressions presented in Equation (6). Then both summations are combined on the right side of the inequality and the largest differentiating term, i.e., the one with index k , is pushed to the left side of the inequality. The sum $\sum_{z=k+1}^m (\beta_z - \alpha_z) \cdot \mathbf{o}[z]$ can be omitted, since $\forall z > k : \beta_z - \alpha_z = 0$. The adjusted potential domain of $\mathbf{o}[k]$ is obtained by rearranging the inequality and dividing it by $\alpha_k - \beta_k$, which is always possible, as $\alpha_k - \beta_k \neq 0$ holds by definition of the index k . We need to distinguish two cases:

Case $\alpha_k > \beta_k$:

$$\begin{aligned} (\alpha_k - \beta_k) \cdot \mathbf{o}[k] &\leq (\beta_0 - \alpha_0) + \sum_{z=1}^{k-1} (\beta_z - \alpha_z) \cdot \mathbf{o}[z] \\ \Rightarrow \mathbf{o}[k] &\leq \frac{\beta_0 - \alpha_0}{\alpha_k - \beta_k} + \sum_{z=1}^{k-1} \frac{\beta_z - \alpha_z}{\alpha_k - \beta_k} \cdot \mathbf{o}[z]. \end{aligned} \quad (8)$$

The upper bound for the random variable s_i^j stored in $\mathbf{o}[k]$ is adjusted as follows:

$$\Lambda_c.S_i^j = \left[\Lambda_c.S_i^j.l, \frac{\beta_0 - \alpha_0}{\alpha_k - \beta_k} + \sum_{z=1}^{k-1} \frac{\beta_z - \alpha_z}{\alpha_k - \beta_k} \cdot \mathbf{o}[z] \right]. \quad (9)$$

ALGORITHM 1: scheduling child locations

```

1: for all (deterministicEvent  $\Upsilon : \mathcal{E}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$ ) do
2:    $\Lambda_c := \text{generateChild}(\Lambda_p, \Upsilon)$ ;
3:    $\Lambda_c.\text{setExpiredRVBounds}(\Upsilon)$ ;
4:   for all (stochasticEvent  $\Upsilon_s : \mathcal{E}_{\text{ran}}^{\min}(\Lambda_p.\Gamma)$ ) do
5:      $\Lambda_c.s.l = \Lambda_p.s.l + \Delta\tau_c$ ; {Equation 12}
6:      $\Lambda_{c'} := \text{generateChild}(\Lambda_p, \Upsilon_s)$ ;
7:      $\Lambda_{c'}.\text{copyExpiredRVBounds}(\Lambda_c)$ ;
8:      $\Lambda_{c'}.s.u = \Lambda_p.s.l + \Delta\tau_c$ ; {Equation 13}
9:     for all (stochasticEvent  $\Upsilon_{s'} : \mathcal{E}_{\text{ran}}^{\min}(\Lambda_p.\Gamma)$ ) do
10:      if ( $s \neq s'$ ) then
11:         $\Lambda_{c'}.s'.l = \Lambda_p.s'.l + s - \Lambda_p.s.l$ ; {Equation 14}

```

Case $\alpha_k < \beta_k$:

$$\begin{aligned}
 (\alpha_k - \beta_k) \cdot \mathbf{o}[k] &\leq (\beta_0 - \alpha_0) + \sum_{z=1}^{k-1} (\beta_z - \alpha_z) \cdot \mathbf{o}[z] \\
 \Rightarrow \mathbf{o}[k] &\geq \frac{\beta_0 - \alpha_0}{\alpha_k - \beta_k} + \sum_{z=1}^{k-1} \frac{\beta_z - \alpha_z}{\alpha_k - \beta_k} \cdot \mathbf{o}[z].
 \end{aligned} \tag{10}$$

The lower bound for the random variable s_i^j stored in $\mathbf{o}[k]$ is adjusted as follows:

$$\Lambda_c.s_i^j = \left[\frac{\beta_0 - \alpha_0}{\alpha_k - \beta_k} + \sum_{z=1}^{k-1} \frac{\beta_z - \alpha_z}{\alpha_k - \beta_k} \cdot \mathbf{o}[z], \Lambda_c.s_i^j.u \right]. \tag{11}$$

Repeating the pairwise comparison for every $\Delta\tau^* \in \Delta T_{\text{det}}^{\min}(\Lambda_p.\Gamma)$, we compute the resulting upper and lower bounds and update the potential domain for child Λ_c by taking the maximum of the pair-wisely computed lower bounds and the minimum of the corresponding upper bounds. This ensures that the resulting potential domain for each random variable equals the intersection of all intervals obtained by the pairwise comparison of $\Delta\tau_c$ and any $\Delta\tau^* \in \Delta T_{\text{det}}^{\min}(\Lambda_p.\Gamma)$.

For each scheduled child location Λ_c , the potential domain further needs to be adjusted in case one or more general transitions are currently enabled in the parent location. This is discussed in the following section, where we assume that for all expired random variables, the intervals in $\Lambda_{c'}.S$ equal the corresponding intervals in $\Lambda_c.S$, that have been determined in the procedure described above.

4.2 Scheduling the child locations

After computing the potential domains for each minimum next event, the corresponding child location needs to be scheduled for each combination of possible next deterministic event and each enabled stochastic firing. This requires iterating over all elements in the set of minimum deterministic events, generating the corresponding child locations and computing the potential domain, as explained in the previous section. Then the firings of all enabled general transitions need to be scheduled before each minimum next event, and the corresponding child locations are generated. To ensure that all random variables are scheduled once before and once after each minimum next event, the intervals of the currently enabled random variables in the potential domain need to be further restricted.

Algorithm 1 presents pseudo code for the scheduling of all child locations for a given parametric location Λ_p . In line 2, the function $\text{generateChild}(\Lambda_p, \Upsilon)$ creates a location for each child Λ_c of parent Λ_p for every event Υ in the set of minimum deterministic events $\mathcal{E}_{\text{det}}^{\min}$. The entry time of each

child location $\Lambda_c.t$ is set to $\Lambda_p.t + \Delta\tau_c$, and the interval bounds for expired random variables are adjusted, as explained in the previous section by calling the function $\Lambda_c.\text{setExpiredRVBounds}(\Upsilon)$, where Υ is the source event of Λ_c (line 3). To schedule all competing general transition firings after the event $\Upsilon(\Lambda_p.\Gamma, \Lambda_c.\Gamma) = (\Delta\tau_c, \varepsilon_c)$, all stochastic events $\Upsilon_s \in \mathcal{E}_{\text{ran}}^{\min}(\Lambda_p.\Gamma)$ are iterated (line 4) and the lower interval bound for each of the still enabled random variables is increased by $\Delta\tau_c$ in the location Λ_c (lines 5–6). Formally, for the continuously enabled random variable s_i^j , the lower interval bound $\Lambda_c.S_i^j.l$ of the potential domain $\Lambda_c.S$ is adapted by increasing the lower bound $\Lambda_p.S_i^j.l$ of the potential domain of the parent location:

$$\Lambda_c.S_i^j.l = \Lambda_p.S_i^j.l + \Delta\tau_c. \quad (12)$$

After adapting the already scheduled successors that correspond to deterministic events, we need to schedule each enabled stochastic firing before each of the deterministic events. This requires the construction of new parametric locations. Hence, for each enabled stochastic firing, i.e., for each event $\Upsilon(\Lambda_p.\Gamma, \Lambda_c.\Gamma) = (\Delta\tau_{c'}, \varepsilon_{c'}) \in \mathcal{E}_{\text{ran}}^{\min}(\Lambda_p.\Gamma)$ one new parametric location $\Lambda_{c'}$ is constructed (line 6) for each instance of the for loop. Furthermore, the intervals of the expired random variables are inherited from the corresponding deterministic location Λ_c by calling the function $\Lambda_{c'}.\text{copyExpiredRVBounds}(\Lambda_c)$ (line 7).

Furthermore, the upper interval bound of the random variable s_i^j that related to the stochastic firing is decreased in the new child location $\Lambda_{c'}$, such that the firing takes place before any other competing event occurs (line 8). Consequently, s_i^j cannot take any value that is larger than the sum of its previous lower bound in the parent location Λ_p and the time to the next competing event. The upper interval bound $\Lambda_{c'}.S_i^j.u$ of the child location is decreased to ensure that the enabled random variable s_i^j fires first:

$$\Lambda_{c'}.S_i^j.u = \Lambda_p.S_i^j.l + \Delta\tau_{c'}. \quad (13)$$

Next, the potential domains for all other enabled random variables have to be adjusted, such that the corresponding firings do not occur before the firing related to s_i^j takes place. Hence, for all enabled random variables (except s_i^j), the potential domain is adjusted (lines 9–11). For every random variable $s_u^v \neq s_i^j$ that corresponds to the v th firing of transition $T_u^G \in \mathcal{T}^G$, which is enabled in $\Lambda_p.\Gamma$, the lower interval bound $\Lambda_{c'}.S_u^v.l$ is increased:

$$\Lambda_{c'}.S_u^v.l = \Lambda_p.S_u^v.l + s_i^j - \Lambda_p.S_i^j.l. \quad (14)$$

Furthermore, the support of the random variable s_i^{j+1} that corresponds to the next firing of T_i^G has to be initialized if T_i^G is still enabled in $\Lambda_{c'}.\Gamma$. The interval is then always set to $\Lambda_{c'}.S_i^{j+1} = [0, \infty)$. All of the new child locations Λ_c and $\Lambda_{c'}$ together create the set of children of location Λ_p , i.e., $\text{children}(\Lambda_p)$.

The computational complexity of Algorithm 1 follows the nested for loops and is in $O(n^2 \times |\mathcal{E}_{\text{det}}^{\min}(\Lambda_p.\Gamma)|)$, where $n = |S|$ forms an upper bound on the number of enabled random variables. The size of $\mathcal{E}_{\text{det}}^{\min}(\Lambda_p.\Gamma)$ depends on the number of continuous places, guard arcs, and immediate and deterministic transitions, which are present in the model, as discussed in Section 2.

5 TRANSIENT ANALYSIS USING INTERVAL ARITHMETIC

Transient analysis computes the probability to be in a certain state of the model at time t' . This can easily be extended to computing the probability that the model fulfills a certain atomic property, w.r.t. the discrete or continuous marking. As the PLT is a symbolic representation of the state space evolution over time, the system can be in different parametric locations at time t' , depending on

the values of the random variables. For each of these so-called *candidate locations*, the subset of the potential domains needs to be computed for which the system actually is in that location at time t' . The transient probabilities can then be computed by integrating over the computed subsets.

This section is further organized as follows: Section 5.1 explains how the candidate locations are obtained and their corresponding support is restricted. In Section 5.2 the computation of the transient probabilities is detailed; Section 5.4 discusses the numerical integration. The computational complexity of this approach is briefly discussed in the respective sections.

5.1 Candidate Locations and Restricting their Potential Domains

Since the entry time $\Lambda.t$ of a location Λ depends on the potential domain of the expired random variables $\Lambda.S$, the entry time point is a linear function of the random variables. However, we can specify the minimum entry time, i.e., the earliest possible entry time $\Lambda.t_{\min}^{\text{entry}}$ of a location Λ . Accordingly, $\Lambda.t_{\max}^{\text{entry}}$ specifies the latest possible entry time into a specific location. To identify the candidate locations for t' , also the latest possible exit time of a location needs to be considered, which is obtained as the latest maximum entry time of all child locations:

$$\Lambda.t_{\max}^{\text{exit}} = \max \left\{ \Lambda_c.t_{\max}^{\text{entry}} \mid \Lambda_c \in \text{children}(\Lambda) \right\}. \quad (15)$$

Note that the number of locations in the tree for a maximum considered time $\tau_{\max}$ is finite, according to Reference [34]. To obtain candidate locations for a specific time t' , we check for each location of the PLT, whether the potential domains are such that values for random variables exist that bring the system in that location at time t' . This is done by checking whether $\Lambda.t_{\min}^{\text{entry}} \leq t'$ and $\Lambda.t_{\max}^{\text{exit}} \geq t'$. If the first condition is violated, then the entry time of the location Λ is guaranteed to be later than the considered time point, and hence its children do not need to be considered as candidates either. The second condition ensures that it is possible to still be in location Λ at time t' . In case this condition is violated, the child locations of Λ still need to be considered as candidates. Hence, the set of candidate locations $C_{t'}$ at time t' equals

$$C_{t'} = \left\{ \Lambda \in \mathbf{V} \mid \Lambda.t_{\min}^{\text{entry}} \leq t' \leq \Lambda.t_{\max}^{\text{exit}} \right\}. \quad (16)$$

The earliest and latest possible entry time of a parametric location Λ are computed iteratively by replacing random variables by their lower or upper boundaries. The boundaries of a random variable may depend on random variables with a lower order, i.e., that fired earlier. These dependencies need to be resolved when minimizing the entry time of a location or maximizing the entry time of its children. Using the previously defined order, the linear equation defining $\Lambda.t$ can be rearranged such that the expired random variables together with their factors occur according to their firing order:

$$\Lambda.t = f_0 + \sum_{k=1}^m f_k \cdot \mathbf{o}[k]. \quad (17)$$

By iterating over those random variables, it is possible to resolve their dependencies and to minimize the entry time at the same time. In case random variable s_i^j is stored in $\mathbf{o}[l]$, the random variable s_i^j needs to be replaced within $\Lambda.t$ by the corresponding interval bound $S_i^j.l$ or $S_i^j.u$, depending on the sign of f_l . This step is repeated for all $l \leq m$ in descending order and the resulting $\Lambda.t$ after each step then only depends on all $\mathbf{o}[k] < \mathbf{o}[l]$:

$$\Lambda.t = f_0 + \sum_{k=1}^{l-1} f_k \cdot \mathbf{o}[k] + \begin{cases} f_l \cdot S_i^j.l, & \text{if } f_l > 0, \\ f_l \cdot S_i^j.u, & \text{if } f_l < 0. \end{cases} \quad (18)$$

Once all random variables have been resolved by the corresponding lower or upper boundary, $\Lambda.t$ has been minimized within the potential domain of location Λ . When computing the maximum of the entry time for all children, i.e., the latest exit time point, the conditions to replace the random variables are both reversed.

In all candidate locations the potential domain of all random variables present up to t' need to be restricted, such that only those values remain for which the system certainly is in that location at time t' . The computation of the restricted potential domains per location corresponds to a *cylindrical algebraic decomposition* [9], which returns the restricted potential domains S' as a set of multi-dimensional intervals. For linear inequalities, this can be computed in $O(n^2)$ using a variant of the *Fourier-Motzkin elimination* [10].

Under the assumption that no conflicts have occurred, the restricted potential domain of a location and the restricted potential domains of its children form disjoint subsets. In case a conflict occurs, it needs to be resolved by adapting the value $\Lambda_c.p$, for all children that participate in the conflict. In that case, the restricted domains of all conflicting child locations may overlap. Note that conflicting child locations have the same minimum entry time.

5.2 Computing Transient Probabilities

When the candidate nodes and the restricted intervals at time t' have been determined, the transient probabilities can be computed. Recall that each general transition was assigned a continuous probability distribution that describes how the probabilities are distributed over the values of the random variables. The probability to be in a specific candidate location Λ at time t' is computed by first integrating over the joint probability density function of all random variables and then multiplying the result with the *accumulated conflict probability* $p_{\text{acc}}(\Lambda)$ towards that location Λ . The latter is computed recursively by multiplying the conflict probability of each location visited when traversing the tree from the root location to Λ :

$$p_{\text{acc}}(\Lambda) = \begin{cases} p_{\text{acc}}(\Lambda_p) \cdot \Lambda.p, & \text{if } \exists \Lambda_p : \Lambda \in \text{children}(\Lambda_p), \\ \Lambda.p, & \Lambda \text{ is the root location.} \end{cases} \quad (19)$$

Recall that the potential domain $\Lambda.S$ consists of one multi-dimensional interval for each location Λ . The restricted potential domain $\Lambda.S'$ w.r.t. the time t' , however, may consist of several multi-dimensional intervals, as the Fourier-Motzkin elimination might split the potential domain in several dimensions. Hence, the restricted domain $\Lambda.S'$ is a set of multi-dimensional intervals. Each element $\mathbf{r} \in \Lambda.S'$ is represented by a vector of intervals defining an upper and a lower bound for each random variable present. Note that this does not influence the number of random variables present in location Λ : In case $|s_\Lambda| = n$ (expired and not yet expired) random variables are present in Λ , we obtain n -dimensional integrals. The random variables that correspond to the particular next firing of each enabled general transition need to be included to ensure that this firing has not taken place yet.

Recall that having $m \leq n$ expired random variables, for each $k \in [1, m]$ and $\mathbf{r} \in \Lambda.S'$, $\mathbf{r}[k]$ can only depend on the random variables $\mathbf{r}[1], \mathbf{r}[2], \dots, \mathbf{r}[k-1]$. Similarly, for each $l \in [m+1, n]$, $\mathbf{r}[l]$ can only depend on $\mathbf{r}[1], \mathbf{r}[2], \dots, \mathbf{r}[m]$. Hence, the lower bound of $\mathbf{r}[1]$ is a constant as no dependencies can exist. For all $i \in [1, n]$, the upper bound of $\mathbf{r}[i]$ might be infinite as well. Further recall that the bounds of the (restricted) potential domains of each random variable are linear expressions. We denote the lower bound of $\mathbf{r}[i]$ with $\mathbf{r}[i].l$ and the upper bound with $\mathbf{r}[i].u$, respectively.

Let function $g_i(s)$ denote the probability density function of the random variables $s_\Lambda[i]$, which is uniquely determined by the probability distribution assigned to the corresponding general

transition. All random variables are accessed in their firing order, as specified in $\mathbf{s}_\Lambda$. The resulting transient probability at t' for parametric location Λ then equals the probability that the values of all random variables in $\mathbf{s}_\Lambda$ lie within $\Lambda.S'$, multiplied by its accumulated conflict probability, as given by Equation (20):

$$Prob(\Lambda, \mathbf{s}_\Lambda, \Lambda.S') = p_{acc}(\Lambda) \cdot \sum_{\mathbf{r} \in \Lambda.S'} \int_{\mathbf{r}[1].l}^{\mathbf{r}[1].u} \cdots \int_{\mathbf{r}[n].l}^{\mathbf{r}[n].u} G(\mathbf{s}_\Lambda) \, d\mathbf{s}_\Lambda[n] \cdots d\mathbf{s}_\Lambda[1], \quad (20)$$

where $G(\mathbf{s}_\Lambda) = \prod_{i=1}^n g_i(\mathbf{s}_\Lambda[i])$ is the joint probability density function, which equals the product over the n probability density functions, due to the independence of the random variables, as explained in Section 2. Note that multi-dimensional integration over dependent intervals is required, as the upper and lower bounds per random variable may depend on other random variables, as explained in Section 4.1.

5.3 Transformation of Variables

There exist a variety of implementations for (numerical) multi-dimensional integration. However, to the best of our knowledge there are no libraries available for multi-dimensional integration over dependent intervals, but mostly over (bounded) rectangular regions of integration, i.e., with constant limits only. To use, e.g., the *GNU Scientific Library (GSL)*² [18], which provides integration over rectangular regions, a transformation of each $\mathbf{r} \in \Lambda.S'$ onto a rectangular region is required. Based on the approach of Reference [38], we present such a transformation for the integral from Equation (20) in the following:

The basic idea is to use the Transformation Theorem by Hammer and Wymore [5, 27], which is the generalization of integration by substitution on functions of higher dimensions. It is used for calculating integrals if the integral can be calculated more easily after transformation to another coordinate system. The goal in the following is to use a transformation $\mathcal{T}$ to transform the arbitrary region of integration V into a simpler, aligned polytope V^* and to calculate the integral with integrand $f(x)$ by using the Transformation Theorem:

$$\int_V f(x) \, dx = \int_{V^*} f(\mathcal{T}(v)) \cdot |\det(D\mathcal{T}(v))| \, dv, \quad (21)$$

where $\mathcal{T}$ is a bijective transformation from V^* to V and $D\mathcal{T}$ is the Jacobian matrix of $\mathcal{T}$. For Equation (21) to be applicable f must be measurable on V and $\mathcal{T}$ must be a diffeomorphism. Function f is measurable, since the (product of) probability density functions over V is measurable. The transformation $\mathcal{T}$ is a differentiable, bijective function whose inverse is differentiable as well and hence fulfills the second condition.

Let U^n be the n -cube with coordinates in $[-1, 1]$ for every dimension. A transformation $\mathcal{T}_r$ for any $\mathbf{r} \in \Lambda.S'$ with $|\mathbf{r}| = n$ from U^n onto $\mathbf{r}$ is given by:

$$\begin{aligned} \mathcal{T}_r: U^n &\rightarrow \mathbb{R}^n, \\ v_i &\mapsto \frac{\mathbf{r}[i].u + \mathbf{r}[i].l}{2} + v_i \frac{\mathbf{r}[i].u - \mathbf{r}[i].l}{2}, \quad i = 1, \dots, n, \end{aligned} \quad (22)$$

where v_i denotes the variable transformed onto the random variable $\mathbf{s}_\Lambda[i]$.

²<https://www.gnu.org/software/gsl/>.

Note that this does not work for the case $\mathbf{r}[i].u = \infty$, as discussed later. The Jacobian determinant of this transformation is given by:

$$|\det(D\mathcal{T}_{\mathbf{r}}(v))| = \prod_{i=1}^n \left(\frac{\mathbf{r}[i].u - \mathbf{r}[i].l}{2} \right). \quad (23)$$

Hence, it follows:

$$\begin{aligned} Prob(\Lambda, \mathbf{s}_{\Lambda}, \Lambda.S') &= p_{\text{acc}}(\Lambda) \cdot \sum_{\mathbf{r} \in \Lambda.S'} \int_{\mathbf{r}[1].l}^{\mathbf{r}[1].u} \cdots \int_{\mathbf{r}[n].l}^{\mathbf{r}[n].u} G(\mathbf{s}_{\Lambda}) \, d\mathbf{s}_{\Lambda}[n] \cdots d\mathbf{s}_{\Lambda}[1], \\ &= p_{\text{acc}}(\Lambda) \cdot \sum_{\mathbf{r} \in \Lambda.S'} \int_{-1}^1 \int_{-1}^1 \cdots \int_{-1}^1 (F(\mathbf{v}) \cdot |\det(D\mathcal{T}_{\mathbf{r}}(v))|) \, dv_n \cdots dv_1, \end{aligned} \quad (24)$$

where $F(\mathbf{v})$ is $G(\mathbf{s}_{\Lambda})$ under the transformation stated by Equation (22).

Since the transformation of variables in Reference [38] is not defined for a random variable s_i with an upper interval bound $\mathbf{r}[i].u = \infty$, we handle this special case as follows: We make use of the maximum time $\tau_{\max}$, up to which the PLT is created, and in a first step, we set $\mathbf{r}[i].u = \tau_{\max}$. However, in case that the probability density of s_i is greater than zero in the interval $[\tau_{\max}, \infty]$, we further need to adjust the result of the integration to include the missing probability mass in $[\tau_{\max}, \infty]$. This can be realized by both: (i) adding the integral over $s_i \in [0, \infty]$ and (ii) subtracting the integral over $s_i \in [0, \tau_{\max}]$, where the integrand and the interval bounds for all other bounded random variables remain the same.

5.4 Numerical Integration Based on Monte Carlo

Solving Equation (24) for unknown antiderivatives efficiently in multiple dimension requires numerical integration. We use approximative methods together with an estimation of the (statistical) error for bounded rectangular regions.

Monte Carlo methods [37] provide algorithms for the approximation of such integrals including a reliable error estimate, which are well-suited to compute integrals of higher dimensions, since their convergence rates, i.e., the speed at which the estimated probability approaches the real probability, are independent of the dimension, cf. Reference [37]. The main idea for estimating an integral as in References [18, 41] is to select N random points from a distribution of points in the (rectangular) integration region U^n .

Let x_i with $i \in \mathbb{N} : i \leq N$ denote those random points, and let V_{U^n} denote the volume of region U^n . An estimate $E(f; N)$ for the integral of the function $f(x)$ over the U^n is then given by:

$$E(f; N) = \frac{V_{U^n}}{N} \sum_{i=1}^N f(x_i). \quad (25)$$

The error on the estimate $\sigma(E; N)$ results from the square root of the estimated variance of the mean $\sigma^2(E; N)$, as stated by Equation (26):

$$\sigma^2(E; N) = \frac{(V_{U^n})^2}{N^2} \sum_{i=1}^N (f(x_i) - \langle f \rangle)^2, \quad (26)$$

where $\langle f \rangle$ denotes the arithmetic mean of the evaluations of f over N sample points.

Our implementation uses the adaptive Monte Carlo scheme VEGAS [37] as implemented in GSL to compute the transient probabilities. VEGAS is based on *importance sampling*, where the probability distribution of the random points is iteratively improved to reduce the variance. We assume that a data structure d is provided by implementation of the approach from Section 5.1. For every

candidate node and for every firing in the past, i.e., for every random variable s_i , the corresponding functions $r[i].l$ and $r[i].u$ and its probability distribution are stored and subsequently transformed onto the cube $[-1, 1]^n$ (see Section 5.3).

Since the transformation of variables is done for every sample point, the total complexity of the transformation is in $O(N \cdot n)$. When performing transient analysis, integration is performed once for each candidate location with a non-empty restricted potential domain S' . When computing the probability for more complex properties, the restricted potential domain may consist of more intervals, whose number is bounded by the number of constraints considered in the cylindrical decomposition, which is given by $O(2 \cdot 4^{n-2})$ for a location with n random variables present [30].

6 GEOMETRICAL APPROACH TO TRANSIENT ANALYSIS

In addition to the state-space representation as PLT, a geometric representation of locations has been proposed together with an automated transformation from locations into geometric regions [34]. The geometric state set representation relies on closed convex polytopes and the region corresponding to the location is constructed in $\mathcal{H}$ -representation. This representation allows for model checking more complex properties, as described in Reference [33].

In the following, we extend the concept presented in Section 5 to the geometric representation of locations to facilitate the computation of probabilities of more complex properties as required for model checking. Section 6.1 presents the concept of regions that can be represented using convex polytopes. Sections 6.2 and 6.3 discuss two approaches on how to integrate over these polytopes to compute probabilities. The computational complexity of both is examined in Section 6.4.

6.1 The Concept of Regions

The geometric representation of a location Λ has $n + 1$ dimensions, given a vector s_Λ of size n . One additional dimension is included to represent time. Following Reference [22], a region is defined as the maximal connected set of states whose discrete parts remain unchanged as long as no event occurs.

Definition 6.1. A region R is a maximal convex set of points (s, t) , for which the following three conditions are fulfilled for all $(s_1, t_1), (s_2, t_2) \in R, (s_1, t_1) \neq (s_2, t_2)$:

- (1) $\Gamma(s_1, t_1).m = \Gamma(s_2, t_2).m$,
- (2) $\Gamma(s_1, t_1).d = \Gamma(s_2, t_2).d$,
- (3) $\Gamma(s_1, t_1).e = \Gamma(s_2, t_2).e$.

Hence, within a region R , the discrete marking $\Gamma(s, t).m$, the drift vector $\Gamma(s, t).d$, and the enabling status vector $\Gamma(s, t).e$ are equal for any point $(s, t) \in R$. The amount of fluid and the clock valuations are linear equations of s and t . As a consequence, each boundary between regions, which represents the occurrence of an event, is characterized by a linear polynomial of s and t and thus restricts a *half-space* in $n + 1$ dimensions. Due to the linearity of events and the restriction with respect to $\tau_{\max}$, these half-spaces form convex polytopes. For a proof, we refer to Reference [20]. A polytope P can be represented as the intersection of a finite set of half-spaces stored as a pair of a matrix A and a vector b such that $P = \{x \mid A \cdot x \leq b\}$. This representation is referred to as $\mathcal{H}$ -representation and the resulting polytope is called $\mathcal{H}$ -polytope.

Transformation into regions and implementation. A transformation from locations into geometric regions has been presented in Reference [34] that fully relies on halfspace intersection, which is dual to the convex-hull problem. We used the C++ library HyPRO [42] to create a convex polytope representing a region using quickhull [4] algorithms. This allows the analysis of Hybrid Petri nets with a finite number of multiple stochastic firings. HyPRO does not limit the number of dimensions and provides most of the set operations required for the analysis of HPnGs. We are

aware that there exist various implementations of convex polytopes, e.g., PPL [3], POLYMAKE [19], CGAL [48], CDD [17], which can be used as well but do not provide intuitive, convenient interfaces and conversion functions. Furthermore, we can exploit the features of HyPRO for model checking.

Computing the transient probability for a given time point t' requires intersecting every region with the hyperplane representing time t' . As convex polytopes are closed w.r.t. intersection, the result (if not empty) is again a convex polytope that represents the possible states of the system at time t' . To compute the joint probability distribution the dimension corresponding to time is removed from the polytope representation by projection.

6.2 Integration over Convex Polytopes

The integration bounds for polytopes that are aligned to the coordinate axes can be trivially deduced from the vertices. In general, however, polytopes created for a parametric location are not aligned. To integrate over general polytopes, two different approaches are discussed in this Section. The first approach is based on integration over general simplices as presented in Section 6.2.1. To use this integration method, Section 6.2.2 presents how an arbitrary convex polytope can be divided into standard simplices. The presented approach maintains a strict order of the integration intervals and allows the use of general integration methods. Alternatively, Section 6.3 shows how Monte Carlo integration can be applied directly to convex polytope using a bounding box.

6.2.1 Transformation to Simplices. Computing the Jacobian Matrix $D\mathcal{T}$ as presented in Section 5.3 for an arbitrary transformation $\mathcal{T} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is computationally expensive, due to the quadratic nature of the matrix. Alternatively, an affine transformation may be used [5]. An affine transformation $A \cdot R + b$ for a d -dimensional set R with $A \in \mathbb{R}^{d \times d}$ and $b \in \mathbb{R}^d$ is defined as the set:

$$\{A \cdot x + b \mid x \in R\}$$

and can be described as a combination of scalings, rotations, and skewings by a matrix $A \in \mathbb{R}^{d \times d}$ and a translation given by the vector $b \in \mathbb{R}^d$. For such an affine transformation $\mathcal{A}$ the Jacobian matrix $D\mathcal{A}$ simply is A , hence, only the determinant of a matrix needs to be calculated. Additionally, $\mathcal{A}$ is a diffeomorphism if A is invertible, because all linear functions $A \cdot x$ are differentiable. However, a limiting factor of affine transformations is their preservation of parallelism, which greatly limits the types of polytopes that they can align to the coordinate axes. For example, a general polytope with four vertices cannot be affinely transformed into a square, because the opposing sites of a quadrilateral are, in general, not parallel. A polytope for which such an affine transformation is always possible is a so-called n -simplex. A simplex is the generalization of the notion of a triangle to arbitrary dimensions. A unit n -simplex $\Delta_n^* \subseteq \mathbb{R}^n$ is defined by the convex hull $\text{cHull}(\{0, e_1, \dots, e_n\})$ of its $n + 1$ vertices denoted $\{0, e_1, \dots, e_n\}$:

$$\begin{aligned} \Delta_n^* &:= \text{conv}(\{0, e_1, \dots, e_n\}) \\ &= \left\{ x \in \mathbb{R}^n \mid \sum_{i=1}^n x_i \leq 1 \wedge \forall i \in \{1, \dots, n\} : x_i \geq 0 \right\}. \end{aligned} \quad (27)$$

For a unit n -simplex Δ_n^* the region of integration is given by the set of intervals as presented in the following: The integration limits are either constant or depend on previous variables, according to Section 5.2. This order is given for unit simplices, as shown by the following equation:

$$\forall x \in \Delta_n^* : \forall i \in \{1, \dots, n\} : x_i \in \left[0, 1 - \sum_{j=1}^{i-1} x_j \right]. \quad (28)$$

ALGORITHM 2: getProbability(s: Δ)

```

1: matrix  $A : \mathbb{R}^{n \times n}$ 
2:  $\{v_0, \dots, v_n\} = s.\text{getVertices}()$ 
3: for all column  $A_i$  of  $A$  do
4:    $A_i = (v_i - v_0)$ 
5:  $\mathcal{A} = s \rightarrow A \cdot s + v_0$ 
6:  $G = s \rightarrow \prod_{i=1}^n g(s)$ 
7:  $f = s \rightarrow G(\mathcal{A}(s)) \cdot |\det(A)|$ 
8:  $I : \text{list}()$ 
9: for  $i = 0$  to  $o.\text{size}()$  do
10:   $I.\text{append}([0, 1 - \sum_{j=0}^i o[j]])$ 
11: return  $\text{integrate}(f, I)$ 

```

Each coordinate is non-negative and must be at most one minus all of the preceding coordinates accumulated. This ensures that the total sum of the coordinates is at most one and that the sequence of intervals describes exactly the set given in Equation (27).

Finally, for any n -simplex $\Delta_n = \text{cHull}(\{v_0, \dots, v_n\}) \subseteq \mathbb{R}^n$, an affine transformation $\mathcal{A}_{\Delta_n}$ from the unit n -simplex Δ_n^* onto Δ_n , is defined by

$$\mathcal{A}_{\Delta_n} : \mathbb{R}^n \rightarrow \mathbb{R}^n$$

$$x \mapsto ((v_1 - v_0) \dots (v_n - v_0)) \cdot x + v_0. \quad (29)$$

The exact assignment of a random variable s_i^j in a vertex x is denoted $s_i^j \in x$. Recall that the function $g(s_i^j)$ denotes the probability density function of a random variables s_i^j , as presented above. The joint probability density function of all random variables present in a parametric location is denoted as $G(x)$ in the following: Since the firing times of general transitions are independently distributed, the joint probability of all random variables represented by a vertex x is the product of their individual densities, i.e.,

$$G(x) = \prod_{\forall s_i^j \in x} g(s_i^j). \quad (30)$$

For any simplex $\Delta_n \in \text{Sat}^t(\Phi)$ the probability over the joint probability density function G can be computed as:

$$\int_{\Delta_n} G(x) \, dx = \int_0^1 \dots \int_0^{1 - \sum_{i=1}^{n-1} x_i} G(\mathcal{A}_{\Delta_n}((x_1, \dots, x_n))) \cdot |\det(A)| \, dx_n \dots dx_1, \quad (31)$$

with normalized integration bounds. The affine transformation $\mathcal{A}_{\Delta_n}$ is defined by the $n \times n$ matrix A , as defined in Equation (29).

Algorithm 2 provides pseudo-code for the function getProbability to obtain the probability to be in a simplex, with a single function parameter s . First, the matrix A needed for the affine transformation is initialized (line 1) and all vertices are obtained from the simplex s (line 2). These vertices are used to create the transformation matrix (lines 3–4) as presented in Equation (29). Then three functions are initialized, each of which returns values for a vector of random variables. The function $\mathcal{A}$ denotes the affine transformation (line 5), the function G returns the joint probability density of all random variables (line 6), and the function f is used to apply the transformation theorem during integration (line 7). The region of integration then is initialized (lines 8–10) as

presented in Equation (28). To obtain the probability, a function integrate is called with the variables f and I . Hence, the integration method can easily be exchanged.

6.2.2 Extension to Arbitrary Convex Polytopes. Algorithm 2 computes $\Pr(\Delta_n)$ for arbitrary simplices Δ_n and needs to be generalized to arbitrary convex polytopes R . To compute the probability $\Pr(R)$, the polytope R needs to be subdivided into simplices for which the probability can be calculated individually. In the following, a polytope is partitioned using triangulation denoted as $\bar{\Delta}(R)$. More formally, the n -dimensional Delaunay triangulation [49] is used as presented in the following:

Definition 6.2 (n-dimensional Delaunay-triangulation). Given a convex polytope $R \subseteq \mathbb{R}^n$ in n -dimensional space, the triangulation $\bar{\Delta}(R)$ is a simplicial complex

$$\bar{\Delta}(R) := \{\Delta \subseteq \mathbb{R}^n \mid \Delta \text{ is a simplex}\},$$

where R is the union of all simplices, i.e., $R = \bigcup_{\Delta \in \bar{\Delta}(R)} \Delta$, and the intersection of two simplices is either empty or a facet, i.e., $\forall \Delta_i, \Delta_j \in \bar{\Delta}(R)$ with $i \neq j$, Δ_i and Δ_j are disjoint.

The probability to be in R can then be computed in three consecutive steps:

- (1) Construct a triangulation $\bar{\Delta}(R)$ for R .
- (2) For each resulting simplex $\Delta \in \bar{\Delta}(R)$, calculate the probability to be in that simplex, i.e., $\Pr(\Delta)$, using Algorithm 2.
- (3) The total probability can be accumulated, as the resulting simplices are pairwise disjoint:
 $\Pr(R) = \sum_{\Delta_i \in \bar{\Delta}(R)} \Pr(\Delta_i)$.

6.3 Direct Numerical Integration over Convex Polytopes

The previously presented effort to generate integration regions with very specific boundaries from arbitrary convex polytopes allows the use of different integration methods. When focusing on numerical integration methods, the integration can be carried out directly on convex polytopes, as presented in the following. This concept is implemented via Monte Carlo integration as included in the GSL library and estimates the multidimensional integral of the form

$$\int_{x_1.l}^{x_1.u} \dots \int_{x_n.l}^{x_n.u} f(x_1, \dots, x_n) dx_n \dots dx_1$$

over a hypercubic region $((x_1.l, x_1.u), \dots, (x_n.l, x_n.u))$. The minimum axis-aligned bounding box of a convex polytope can be used as a hypercubic region for Monte Carlo integration, as defined below.

Definition 6.3 (Axis-aligned bounding box). Let $P \subseteq \mathbb{R}^n$ be a convex polytope. The bounding box of a convex polytope P is an axis-aligned box within which all points of P lie. The minimum value in a dimension i is obtained by $\min_{x \in P} x_i$ and the maximum value in a dimension i is obtained by $\max_{x \in P} x_i$. The bounding box $\text{bb}(P)$ of a polytope P then is defined as the cartesian product of n intervals:

$$\text{bb}(P) := [\min_{x \in P} x_1, \max_{x \in P} x_1] \times \dots \times [\min_{x \in P} x_n, \max_{x \in P} x_n].$$

The bounding box of a candidate region R defines the region of integration. To compute the probability over a candidate region R , an indicator function evaluates whether a point x lies within R :

$$1_R(x) := \begin{cases} 1, & \text{if } x \in P, \\ 0 & \text{otherwise.} \end{cases}$$

Using the joint probability density function $G(x)$ as presented above, the probability for a point x to be in a candidate region R can be computed as:

$$\int_{\text{bb}(R)} \mathbf{1}_R(x) \cdot G(x) \, dx.$$

Hence, the Monte Carlo integration samples a vertex $x \in \text{bb}(R)$ and the joint density $G(x)$ is only accumulated in case the vertex x lies within R .

6.4 Complexity

In the following, we first discuss the complexity of transforming a single parametric location into a region, as briefly explained in Section 6.1. Then, we indicate the computational complexity of integration over convex polytopes, relying on triangulation, as indicated in Section 6.2. Finally, we discuss the complexity when integrating directly over convex polytopes using a bounding box, as explained in Section 6.3.

Constructing a region. Since the construction of a region solely relies on the geometric intersection of half-spaces, we limit our discussion to the estimation of the number of intersections that is required. As presented in Reference [52], intersection of polytopes in $\mathcal{H}$ -representation can be performed in polynomial time, but heavily depends on the number of dimensions. A single intersection is performed for the source event of the location and multiple intersections for every child of the location. According to Section 4.2, the number of children of a parametric location Λ is in

$$\mathcal{O}\left(\left|\mathcal{E}_{\text{det}}^{\min}(\Lambda, \Gamma)\right| + n\right). \quad (32)$$

Further, half-spaces are created for the constraints of the potential domain of a parametric location Λ . According to Section 4.2, the number of random events, i.e., general transition firings, that may occur is in $\mathcal{O}(n)$ with $n = |\Lambda.S|$. This results in $\mathcal{O}(2 \cdot n)$ half-spaces, i.e., upper and lower bounds for each potential domain. However, this does not change the worst-case complexity, denoted in Equation (32).

Pre-computations required for geometric integration. Geometrically, the restricted domains required for integration are computed by halfspace intersection followed by a projection to remove the dimension of time from the polytope representation. Internally, this requires to change the $\mathcal{H}$ -representation to a $\mathcal{V}$ -representation and back. Relying on HyPro, this is realized by the quickhull algorithm [4]. The worst-case complexity of this algorithm is quadratic with a linear space complexity in the worst case.

Integration over simplices. For determining the complexity of the integration over simplices, we consider the complexity of the different steps. For a d -dimensional polytope with n_f facets, the worst-case computational complexity of the Delauney triangulation is in $\mathcal{O}((n_f)^2)$ and results in $\mathcal{O}((n_f)^{\frac{d}{2}})$ simplices [43, 49]. Note that the number of facets n_f is bounded by the number of events, which may occur in a parametric location.

Since Monte Carlo integration has to be performed for each simplex, the resulting number of integrations is in

$$\mathcal{O}\left(\left(\left|\mathcal{E}_{\text{det}}^{\min}(\Lambda, \Gamma)\right| + n\right)^{\frac{n}{2}}\right). \quad (33)$$

Direct integration. The direct integration of a region requires the computation of one bounding box for the set of convex polytopes per candidate location. This has an exponential worst-case complexity in the number of dimensions. Integration is then called once for the area defined by

each bounding box, and a containment test is required for each integration run to indicate whether the respective sample lies in the currently considered set of convex polytopes. Containment testing has a worst-case complexity that is linear in the number of half-spaces.

7 CASE STUDY: BATTERY BACK-UP

This section presents a case-study model of a factory with battery back-up in Section 7.1 and validates the methods discussed in this article with statistical model checking in Section 7.2, before discussing their performance in more detail in Section 7.3. Section 7.4 presents additional measurements on CPU and memory usage and allows to identify the current bottleneck of the approach.

7.1 Model Description

The system model presented in the following is an HPnG model of a factory with battery back-up. We carry out transient analysis to evaluate the availability of power supply from the grid and to show the feasibility of the presented approach. The analysis is carried out end-to-end, i.e., the model description is parsed, the state space is created, transient analysis is performed, and the resulting probability is returned.

As a feasibility study, we model a factory whose power intake from the grid is limited by a **service level agreement (SLA)** with the energy provider. The factory relies on battery back-up that is discharged in the case of a peak-demand or a power outage. Figure 3 shows the core of the HPnG model, consisting of the battery, the power grid, and costs, and Figure 4 shows the power demands. The continuous place P_0^c in Figure 3 models the battery with capacity B , an inflow (T_0^{Dy}) from the power grid and an outflow (T_1^{Dy}) that corresponds to the power that exceeds the constant peak demand P_A , specified by the SLA. The rates of the dynamic inflow and outflow are specified as $f_{T_0^{Dy}}(\mathbf{d}) = \max((P_A - \sum_{d_i \in \mathbf{d}} d_i), 0)$ and $f_{T_1^{Dy}}(\mathbf{d}) = \max((\sum_{d_i \in \mathbf{d}} d_i - P_A), 0)$, where $\mathbf{d} \in \mathbb{R}^{|\mathcal{T}^F|}$ is the vector of currently active demands (which are modeled in Figure 4).

The upper left part of Figure 3 models the status of the grid (indicated by a box) and enables the inflow of the battery. Initially, the grid is on, which is indicated by the token in P_0^d . The failure of the grid (T_0^G) is modeled by a general transition, moving the token to P_1^d , indicating that the grid is off. T_0^D models the deterministic recovery of the grid.

Initially, the place P_2^d holds a token, which indicates that the battery still has power. When the battery is completely drained, the immediate transition T_0^I fires and moves the token to place P_3^d , such that T_1^I gets enabled. When the battery is filled again, T_1^I moves the token back to P_2^d . Hence, a token in P_3^d indicates an empty battery, which in turn enables the dynamic transition T_2^{Dy} . This transition models the cost flow that occur due to service-level violations when the accumulated demand exceeds the amount specified in the SLA. Its flow rate is $f_{T_2^{Dy}}(\mathbf{d}) = \max((\sum_{d_i \in \mathbf{d}} d_i - P_A), 0)$ and the extra cost is then collected in the continuous place P_1^c (with infinite capacity).

Figure 4 models three levels of demand via static continuous transitions, i.e., standard ($T_{d_1}^F$), reduced ($T_{d_0}^F$), and extended ($T_{d_2}^F$), that get enabled via guard arcs that are connected to the discrete places $P_{d_0}^d$, $P_{d_1}^d$, and $P_{d_2}^d$. General transitions between those places switch from standard to reduced demand ($T_{d_{10}}^G$) and from standard to extended demand ($T_{d_{12}}^G$). Deterministic transitions model the possibility to switch back to standard demand (from reduced with $T_{d_{01}}^D$ and from extended demand with $T_{d_{21}}^D$). The token indicates the current demand by enabling one of the corresponding continuous transitions.

The number of random variables in the system depends on the considered time t' of the transient analysis and the repair and switching times. At least three random variables are

Fig. 3. HPnG model of a factory that is connected to the power grid and has an additional battery storage.

Fig. 4. HPnG model for the different power demands of the factory.

present, since all three general transitions are initially enabled concurrently. Hence, the deterministic firing times of transition T_0^D , modeling the repair of the grid, and transitions $T_{d_0}^D$ and $T_{d_2}^D$, which reset the power demand to standard, have a high impact on the complexity of the model.

Table 1. Parameters Used for the Single Components of the Case Study Model in Its Initial Setting

Parameter	Value
Peak demand P_A	700 kW
Initial level $m_0(P_0^c)$	1000 kWh
Capacity B of P_0^c	1500 kWh
Initial level $m_0(P_1^c)$	0 kWh
Continuous rate d_0 of $T_{d_0}^F$	500 kW
Continuous rate d_1 of $T_{d_1}^F$	700 kW
Continuous rate d_2 of $T_{d_2}^F$	800 kW
Distribution of T_0^G	$\mathcal{U}(0 \text{ h}, 10 \text{ h})$
Distribution of $T_{d_{10}}^G$	$\mathcal{U}(0 \text{ h}, 10 \text{ h})$
Distribution of $T_{d_{12}}^G$	$\mathcal{U}(0 \text{ h}, 10 \text{ h})$
Firing time at t_1 of T_{01}^D	11 h
Firing time at t_2 of T_0^D	11 h

The discrete initial markings are indicated in the figure.

The parameters of the specific model components are presented in Table 1, as used in the initial setting of the case study. All of the general transitions follow a uniform distribution $\mathcal{U}(0 \text{ h}, 10 \text{ h})$, and the firing time of T_0^D takes natural values between 8 h and 1 h. All arcs (except the guard and inhibitor arc connected to P_0^c) have weight one, which is not indicated in the figures.

7.2 Computation and Validation of Results

The probability that certain properties hold is computed separately using all three methods previously presented, which were implemented within the tool HPNMG [31]. The method **NUM** *intervals* performs integration on the intervals computed from the PLT, **NUM** *simplices* integrates over general polytopes transformed to standard simplices. **NUM** *polytopes* is a direct integration on convex polytopes. All three methods rely on Monte Carlo integration, for which error estimates are computed. For comparison, we include results obtained from the statistical model checker HYPEG [39], whose technique is based on simulation and thus denoted as **SIM** in the following.

In this subsection, we compute the transient probability for two different measures, namely, (i) the probability that the grid is available at time $t' = 8 \text{ h}$ for decreasing grid repair times and (ii) the probability that the standard demand is enabled at time $t' = 8 \text{ h}$ for randomly distributed demand changing times with different distributions.

Results for the first property obtained by the three numerical methods are summarized in Table 2 and validated with the statistical model checker HYPEG. Each column in this table represents a different value for the firing time of the deterministic transition, which models the duration of grid repair. Every entry in Table 2 then shows the probability that P_0^d contains one token at time $t' = 8 \text{ h}$, indicating that the (possibly repaired) power grid is available.

The number of random variables present in the system increases, since T_0^G can fire more often with decreasing firings times of transition T_0^D . The first row of the table represents the resulting state-space dimension. The other rows show the results for the three different numerical computation approaches and the simulation. A cell represents the results of the transient analysis verifying that $\Phi_1 := m(P_0^d) = 1$ holds at time $t' = 8 \text{ h}$, where $m(P_0^d)$ denotes the marking of the discrete place P_0^d . The resulting probability is denoted $p(t', \Phi_1)$, and the resulting error is denoted $e(t', \Phi_1)$. The

Table 2. Probability of an Available Grid, i.e., $\Phi_1 := m(P_0^d) = 1$, for $t' = 8$ h and Decreasing Grid Repair Times T_0^D , Computed Numerically with Different Integration Approaches and via Simulation

T_0^D		8 h	7 h	5 h	3 h	2 h	1 h
	dim	4	4	5	6	8	13
NUM <i>intervals</i>	$p(t', \Phi_1)$	0.200	0.295	0.454	0.593	0.689	0.817
	$e(t', \Phi_1)$	0	$5.5 \cdot e^{-6}$	$5.8 \cdot e^{-6}$	$1.2 \cdot e^{-5}$	$1.0 \cdot e^{-5}$	$1.9 \cdot e^{-5}$
	comp. time	1.67 s	17.74 s	38.92 s	90.22 s	250.15 s	2145.66 s
NUM <i>simplices</i>	$p(t', \Phi_1)$	0.200	0.295	0.455	0.594		
	$e(t', \Phi_1)$	$5.2 \cdot e^{-6}$	$7.7 \cdot e^{-6}$	$6.1 \cdot e^{-5}$	$3.9 \cdot e^{-5}$	-	-
	comp. time	2.85 s	5.66 s	33.45 s	344.54 s		
NUM <i>polytopes</i>	$p(t', \Phi_1)$	0.200	0.295	0.455	0.593	0.691	
	$e(t', \Phi_1)$	0	$4.2 \cdot e^{-6}$	$2.2 \cdot e^{-4}$	$5.6 \cdot e^{-4}$	$9.6 \cdot e^{-4}$	-
	comp. time	0.77 s	1.53 s	3.94 s	21.78 s	1721.95 s	
SIM	$p(t', \Phi_1)$	0.200	0.296	0.455	0.594	0.690	0.818
	comp. time (CI: ± 0.001)	27.61 s	31.30 s	39.87 s	39.29 s	33.57 s	26.04 s
	comp. time (CI: ± 0.0001)	1477.39 s	1983.97 s	2599.05 s	2761.17 s	2630.43 s	1728.41 s

We state the computed probability $p(t', \Phi_1)$, the error from Monte Carlo integration $e(t', \Phi_1)$, and the computation time.

last value in every cell is the computation time. Note that simulation is executed for two 99% confidence intervals with two different widths, i.e., ± 0.001 and ± 0.0001 .

The numerical results are well supported by the simulation **SIM** and illustrate that the computation times of simulation are independent of dimensions, whereas the numerical computation times grow exponentially with increasing dimension. While the **NUM intervals** method is able to compute the probability for a model setting with up to 13 dimensions, the simplices method can not provide results for the model with more than 6 dimensions and the polytope method can only deal with up to 8 dimensions. In the cases indicated with -, the computation was terminated after two hours, as the integration was taking too much time.

NUM polytopes performs better except for the two highest dimensions. For at most 6 dimensions it outperforms the other methods by an order of magnitude. **NUM simplices** takes longer than direct integration over polytopes due to the required transformation to standard simplices, which results in a larger number of objects that are each integrated over individually. The advantage of the simplices method is, however, that it would allow the use of other multi-dimensional integration methods. While the **NUM intervals** method performs best in this setting, it can only be used to compute transient distributions. Model checking, as proposed in Reference [33], requires geometric operations and hence, integration on polytopes.

The number of dimensions affects the computation time of the numerical analysis, whereas the width of the confidence interval significantly influences the computation time of the simulation. A confidence interval of ± 0.00002 would be comparable to the largest error made by the numerical analysis. This is, however, not feasible, since the simulation takes too long for this confidence interval width. For ± 0.001 and ± 0.0001 , the statistical simulation takes (significantly) more time

Table 3. Probability of a Standard Demand, i.e., $\Phi_2 := m(P_{d_1}^d) = 1$, at Time $t' = 8$ h for Randomly Distributed Demand Changing Times with different Distributions, Computed Numerically with different Integration Approaches and via Simulation

$T_{d_{10}}^G, T_{d_{12}}^G$		$\mathcal{U}(0 \text{ h}, 10 \text{ h})$	$\mathcal{U}(6 \text{ h}, 10 \text{ h})$	$\mathcal{N}(\mu=8 \text{ h}, \sigma=1 \text{ h})$	$\mathcal{N}(7 \text{ h}, 1 \text{ h})$	$\mathcal{N}(7 \text{ h}, 2 \text{ h})$
	dim	4	4	4	4	4
NUM <i>intervals</i>	$p(t', \Phi_2)$	0.040	0.241	0.249	0.023	0.070
	$e(t', \Phi_2)$	$1.3 \cdot e^{-5}$	$8.5 \cdot e^{-5}$	$5.5 \cdot e^{-6}$	$9.8 \cdot e^{-6}$	$2.4 \cdot e^{-5}$
	comp. time	4.16 s	4.04 s	8.17 s	8.21 s	7.62 s
NUM <i>simplices</i>	$p(t', \Phi_2)$	0.040	0.242	0.247	0.024	0.068
	$e(t', \Phi_2)$	$1.1 \cdot e^{-6}$	$6.4 \cdot e^{-6}$	$3.5 \cdot e^{-6}$	$1.5 \cdot e^{-7}$	$1.1 \cdot e^{-6}$
	comp. time	3.13 s	3.06 s	3.91 s	3.14 s	1.84 s
NUM <i>polytopes</i>	$p(t', \Phi_2)$	0.040	0.241	0.250	0.023	0.069
	$e(t', \Phi_2)$	$1.0 \cdot e^{-6}$	$6.5 \cdot e^{-6}$	$3.5 \cdot e^{-6}$	$1.5 \cdot e^{-7}$	$1.2 \cdot e^{-6}$
	comp. time	2.95 s	2.89 s	2.94 s	2.95 s	1.62 s
SIM	$p(t', \Phi_2)$ comp. time (CI: ± 0.001)	0.041 8.89 s	0.240 19.54 s	0.250 32.91 s	0.025 7.09 s	0.072 16.87 s

We state the computed probability $p(t', \Phi_2)$, the error from Monte Carlo integration $e(t', \Phi_2)$, and the computation time.

than the numerical analysis for four or five dimensions. For 13 dimensions, the computation time of numerical analysis and simulation with a confidence interval of ± 0.0001 are of the same magnitude. However, the integration error of the numerical analysis is still approximately five times smaller than the half width of the confidence interval.

Results for the second property analyzed in this section are summarized in Table 3. The probability that the standard demand is enabled, i.e., $\Phi_2 := m(P_{d_1}^d) = 1$, at time $t' = 8$ h is computed for varying uniform and folded normal CDFs assigned to the competing general transitions $T_{d_{10}}^G$ and $T_{d_{12}}^G$. Each column holds the results for one CDF. The firing time of T_0^D is set to $\text{dat}_b = 11$ h. Again, the results are computed for all three integration methods and compared to simulation. In this execution run, simulation was performed with the single confidence interval ± 0.001 . Note that the exact CDF used barely affects the runtime of the analysis. The results are well supported by simulation and in most cases the computation time for all numerical approaches is considerably slower. The firing times for the deterministic transitions $T_{d_{10}}^D$ and $T_{d_{12}}^D$ are chosen such that there is no possibility to return to standard demand once it has changed. Hence, the probability of being in standard demand mode results in $T_{d_{10}}^D$ and $T_{d_{12}}^D$ not having fired in the interval $[0 \text{ h}, 8 \text{ h}]$.

7.3 Performance Evaluation

After focusing on the validation of the computed probabilities in the first part of the case study, the second part further analyzes the performance of the different integration methods. We investigate the number of dimensions, the number of locations, and the computation times, when checking whether $\Phi_3 := x(P_1^c) > 0$ kWh holds for increasing time t' . In this setup, the transition T_0^G follows a folded normal distribution with mean $\mu = 14$ h and variance $\sigma = 4$ h. Once the grid failed, it

Fig. 5. Number of dimensions of the state space for different times t' (measured in hours), independent of the property and for $\tau_{\max} = t'$ (left y-axis). Probability that $\Phi_3 := x(P_1^c) > 0 \text{ kWh}$ holds at time t' , computed with **NUM** intervals (right y-axis).

Fig. 6. Number of locations: (i) locations in the parametric location tree created up to $\tau_{\max} = t'$ (black), (ii) the locations the system can be in at time t' (measured in hours) (blue), and (iii) the subset of locations satisfying $\Phi_3 := x(P_1^c) > 0 \text{ kWh}$ at time t' (green).

needs 22 hours to be repaired. $T_{d_{10}}^G$ and $T_{d_{12}}^G$ follow a uniform distribution $\mathcal{U}(0 \text{ h}, 2 \text{ h})$ and the firing times of both $T_{d_{01}}^D$ and $T_{d_{21}}^D$ equals 5 h.

We evaluate the impact of a larger time on the size of the state space and on the computation times. The state space is computed for $\tau_{\max} = t'$, where a larger t' results in more random variables present in the system. Figure 5 indicates the number of dimensions for different values of t' . Additionally, it shows the probability that extra cost have been accumulated, as the available battery charge was not sufficient. For $t' \leq 5 \text{ h}$, we obtain four dimensions, as all three general transitions can fire only once and the time also adds one dimension. In general, two dimensions are added every 5 hours, since the two general transitions that change the demand can both fire once more. At $t' = 22 \text{ h}$ the number of dimensions increases to 13, because the grid can then fail twice.

The computed probability $p(t', \Phi_3)$ is zero for $t' \leq 4$ h and then increases with the possibility of a grid failure. Note that the steeper increase after $t' = 14$ h coincides with the mean of the grid failure distribution. The growing probability is, however, due to the battery level, which has been drained before and is no longer able to provide the facility with enough power in case of a grid failure. For $t' = 21$ h the value 0.586 is reached, which is the last probability that could be computed. For later time points, neither of the integration methods is able to compute the corresponding probability in less than 2 hours.

The y-axis in Figure 6 indicates the number of locations (i) in the PLT created up to time $\tau_{\max} = t'$, (ii) the system can be in at time t' , and (iii) that satisfy Φ_3 at t' . Similar to the number of dimensions, the number of locations also increases periodically. At $t' \geq 22$ h the number of locations rises, considerably. Both, the total number of locations, i.e., the size of the state space, as the dimensionality of the model has a great impact on the computation time of the transient probabilities. While the number of dimensions directly depends on the number of random variables present in the system and is hard to reduce, the number of locations could be reduced by, e.g., a property-guided exploration of the state space.

Figure 7 separately indicates the time needed to (i) create the PLT up to time t' , (ii) compute the restricted domains, and (iii) perform the integration for all presented integration methods w.r.t. constraint Φ_3 . For all integration methods, the time for integration takes an order of magnitude longer than creating the complete PLT and preparing the intervals for final integration. Note that the time for creating the PLT is equal for all methods, since the creation is independent of the chosen method. Further, the time for restricting the domains is equal in Figures 7(b) and 7(c), as both integration approaches rely on the same geometric operations to prepare the intervals for integration. Transformation to standard simplices as well as all operations performed for direct integration are included into the total integration time. Comparing the computation times indicated in Figure 7(a) with Figures 7(b) and 7(c), shows that geometrically restricting the domains takes longer than restricting the intervals arithmetically. However, the geometric approach is more general and enables model checking more complex properties.

Similar to the results shown in Table 2, the **NUM simplices** method was unable to terminate for time t' larger than 9 h and the **NUM polytopes** method for t' larger than 14 h. Only for a time $t' \geq 22$ h, the **NUM intervals** method was unable to complete within 2 hours. Hence, the limiting factor of the presented approach clearly are the number of dimensions and the number of required triangulations and integrations.

Summarizing, we can conclude for lower dimensions that **NUM simplices** and **NUM polytopes** perform better than **NUM intervals**, which becomes more efficient for larger dimensions. This can be explained by the computational and storage complexity of the geometrical operations. Both computing the Delauney triangulation for **NUM simplices** and the bounding box for **NUM polytopes** require the transformation of the hyperplane-representation into the vertex representation for polytopes.

Reproducibility. All computations have been performed on a MacBook Pro with 2.5 GHz Inter Core i7 and 16 GB of memory. The source code for the tool HPNMG can be accessed at <https://zivgitlab.uni-muenster.de/ag-sks/tools/hpnmng>. The tags num-simplices and num-polytopes lead to the different versions using the corresponding integration method. Furthermore, scripts for reproducing the above results can be obtained from <https://zivgitlab.uni-muenster.de/ag-sks/tools/misc/execution-scripts/2020-09-qest-special-issue>.

7.4 CPU and Memory Usage

As shown in the previous section, both versions of the geometric computation are more restricted w.r.t. the number of dimensions than the interval integration method. This is expected

(a) Computation time of the **NUM intervals** integration method for different time points.

(b) Computation time of the **NUM simplices** integration method for different time points.

(c) Computation time of the **NUM polytopes** integration method for different time points.

Fig. 7. Computation times for the probability that $\Phi_3 := x(P_1^c) > 0$ kWh holds at time t' (measured in hours), with three different integration methods. In each plot, the overall computation time equals the sum of the three components: (i) PLT creation (green), (ii) Restricting domains (black), and (iii) Integration (blue).

to be a consequence of the geometric operations necessary for computing the restricted domains and the inherent transformations from the $\mathcal{H}$ -representation to $\mathcal{V}$ -representation of polytopes and vice versa, which is known to be a memory-intensive task. Note that these changes of representation are not required for the interval method. However, the interval method can only

be used to compute transient probabilities and is not suited for model checking more complex properties.

To further analyze memory and CPU usage for the geometric methods, we use the tool `PSRECORD`.³ We record the CPU and memory activity for a set of experiments with both geometric integration approaches. On every experiment, `PSRECORD` is directly attached to the specific process and only measures the accumulated process-specific memory and CPU activity.

The measurements shown in Figure 8 correspond to the case-study setting also discussed in Figure 7. The plots from Figures 8(a)–8(e) show CPU and memory usage for the computations required for $t' = 3$ h, 9 h, and 10 h, where the computation time (in seconds) is provided on the x -axis and CPU and memory usage are depicted on the respective y -axes. In contrast, Figure 8(f) shows the number of constraints resulting from the computation of the restricted domains, as required for the computation of the transient probability at different times t' , as indicated on the x -axis, as well as the resulting number of dimensions of the state space.

The plots in Figures 8(a) and 8(b) show that for $t' = 3$ h memory usage is constant around 15.4 MB for both integration methods, while CPU usage varies considerably for **NUM** *simplices* between 0% and 300% and only rarely drops from the usual 100% for the **NUM** *polytopes* method. Observing the plots for $t' = 9$ h in Figures 8(c) and 8(d), one can see that memory usage slightly increases over the runtime of the computation for both methods. Also CPU usage starts to vary for **NUM** *polytopes* and even more so for $t' = 10$ h (Figure 8(e)). The transient probability for $t' = 10$ h could only be obtained from the **NUM** *polytopes* method. The corresponding memory usage for that computation clearly shows an increased memory usage over the duration of the computation from 15 MB to 30.1 MB in three steps with an additional peak at the end of the computation around 34 MB. More detailed measurements have shown that a new step is introduced when integration is called for an additional candidate location. The geometric computations required to remove the time dimension from the representation of the convex polytopes accounts for the additional memory usage.

Overall, the CPU usage for **NUM** *simplices* is considerably higher than for **NUM** *polytopes*, while the memory usage is almost similar for both methods up to $t' = 9$ h. Hence, we expect that the combination of high CPU usage with the highly increased memory usage, as observed in Figure 8(e), is the reason why **NUM** *simplices* is unable to compute results for $t' \geq 10$ h.

A direct dependency can be observed between the number of dimensions and constraints, as illustrated in Figure 8(f), and the corresponding memory and CPU usage. We have illustrated CPU and memory usage over the full computation time for different times t' resulting in 8, 14, and 20 constraints.

For a more complete comparison of memory and CPU usage, Figure 9 shows boxplots for memory and CPU usage for both geometric integration methods for all times t' , which terminate within 2 h. This results in different scales on the x -axes of the plots. Note that, due to different measurements, also the scales on y -axes differ.

Figures 9(a) and 9(b) show the distribution of CPU usage for varying values of t' . While the median remains stable around 100% for **NUM** *polytope*, it varies slightly for **NUM** *simplices*. It can be seen that in Figure 9(a) CPU usage varies with an increasing number of constraints. Especially when reaching 20 constraints (i.e., for $t' \geq 10$ h), we see that half of the CPU measurements now lie approximately between 50% and 135%, resulting in an interquartile range of 85%. The whiskers indicate those measurements that are just included in the quartiles extended by 1.5 times the interquartile range. For 20 constraints the lower whiskers are at 0% CPU usage and the upper whiskers vary between 237% and 270%. In contrast, Figure 9(b) shows that the interquartile ranges

³<https://github.com/astrofrog/psrecord>.

Fig. 8. CPU and memory usage for the integration methods **NUM polytopes** and **NUM simplices** for three different values for t' (a)–(e) and the average number of constraints required for computing the transient probability at time t' and the number of dimensions (f). Note that for each of the plots the same amount of points was used to plot the CPU and memory usage, even though the x-axis have different maximum values for the time (depending on the total computation time).

for the CPU measurements with **NUM simplices** are considerably larger, starting at time $t' = 1$ h. The effect of the constraints is less visible in this boxplot, than for **NUM polytopes**.

Considering the boxplots for memory usage for both integration methods, as illustrated in Figures 9(c) and 9(d), one can again observe a direct relationship between the number of constraints

(a) CPU usage for **NUM polytopes** for $t' = 1 \text{ h}, \dots, 14 \text{ h}$. (b) CPU usage for **NUM simplices** for $t' = 1 \text{ h}, \dots, 9 \text{ h}$. (c) Memory usage for **NUM polytopes** for $t' = 1 \text{ h}, \dots, 14 \text{ h}$. (d) Memory usage for **NUM simplices** for $t' = 1 \text{ h}, \dots, 9 \text{ h}$.

Fig. 9. Boxplots illustrating the CPU (a) and (b) and memory usage (c) and (d) for the two integration methods **NUM polytopes** and **NUM simplices** for increasing values of t' . The boxplots show the median, the lower and upper quartile of the values, as well as the lower and upper whiskers, indicating the distribution of the measured values.

and the variance in the measurements. For less than 20 constraints (i.e., $t' \leq 9 \text{ h}$) the interquartile range is around 0.5 MB with an outlier of 2 MB for time $t' = 1 \text{ h}$. For more constraints, the interquartile range increases to 15 MB ($t' = 10 \text{ h}$ and $t' = 11 \text{ h}$) and 20 MB ($t' = 12 \text{ h}, 13 \text{ h}, 14 \text{ h}$), while the median also increases to approximately 21 MB and 29.5 MB, respectively. Taking into account the different axes, Figure 9(d) illustrates similar quartiles for **NUM simplices**.

These results can be explained as follows: When applying **NUM simplices** the integration is carried out by GSL, which samples from the unit cube and internally compares the sample with the bounds obtained from triangulation. Note that integration is called once for every simplex, which each results in a peak in CPU usage. Since the integration is solely carried out by GSL, we did not further look into the performance of these calls.

In contrast, **NUM polytopes** calls the integration via GSL once for each pre-computed bounding box and then relies on HyPro to geometrically test containment for each sample. Further detailed measurements have shown that, for a parameter setting resulting in 8 constraints, the CPU is mainly utilized for containment testing. With an increasing number of constraints, the CPU and memory usage required for geometric operations and the corresponding change of representation

becomes inhibitive. It is possible that the combination of third-party libraries results in additional overhead, which could be reduced in the future.

8 CONCLUSION

We proposed and implemented a general algorithm for building a Parametric Location Tree for HPnGs with an arbitrary but finite number of general transition firings and presented the computation of transient probabilities in three stages. First the candidate locations, i.e., parametric locations the system can be in at time t' , were obtained. Second, the potential domain of all candidate locations was restricted, such that only those values of the random variables remained for which the system certainly can be in that location at time t' . Third, the probability to be in a specific candidate location at time t' was computed by integrating over the joint probability density function and multiplying the result with the accumulated conflict probability. We presented two approaches to compute the areas of integration, either using interval arithmetic or using the geometric representation of regions. Both methods then allow direct integration using Monte Carlo simulation on intervals or convex polytopes, respectively. Alternatively, the resulting objects can be transformed into standard simplices, which then allows the use of other multi-dimensional integration methods, which is, however, out of the scope of this article.

The feasibility study shows that the interval approach performs better for higher dimensions than the geometrical approach. This is due to the dedicated implementation of Fourier-Motzkin, which is tailored towards the specific requirements of our algorithms. In contrast, the library HyPro is more general and can be used to model check more complex properties. Both geometrical integration methods, however, require the transformation of the halfspace representation into the vertex representation of polytopes, which can be exponential in the state space dimension.

A case study on a battery-backup system showed the feasibility and current limitations of the approach.

After a validation of the computed probabilities, we performed dedicated measurements for a range of parameter settings. It could be seen that the presented geometric integration approaches are not able to perform in as many dimensions as the presented interval integration, since the geometric approaches can be used to compute not only the transient probability, but also more complex properties, like, e.g., model checking. Hence, we performed further measurements on CPU- and memory-usage for the geometric integrations.

These measurements have shown that this bottleneck currently lies at the pre-computations required for integration, i.e., removing time from the state representation, computing bounding boxes, and the inherent changes between $\mathcal{H}$ - and $\mathcal{V}$ -representations and vice versa.

We plan to further improve the efficiency of the computation by better integrating the use of third-party libraries. Furthermore, we will investigate a property-guided exploration of the state space and by investigating how multi-dimensional integration can be performed without the need to compute the convex hull.

ACKNOWLEDGMENTS

A special thanks goes to Stefan Schupp for his help with the integration of HyPro and performing dedicated measurements on the relevant function calls. Furthermore, we would like to thank the reviewers for their very constructive and positive remarks.

REFERENCES

- [1] Hassane Alla and René David. 1998. Continuous and hybrid Petri nets. *J. Circ. Syst. Comput.* 8, 1 (1998), 159–188.
- [2] Tesfaye Amare, Bjarne E. Helvik, and Poul E. Heegaard. 2018. A modeling approach for dependability analysis of smart distribution grids. In *21st Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN'18)*. IEEE, 1–8.

- [3] Roberto Bagnara, Patricia M. Hill, and Enea Zaffanella. 2008. The Parma Polyhedra Library: Toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems. *Sci. Comput. Prog.* 72, 1–2 (2008), 3–21.
- [4] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. 1996. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.* 22, 4 (1996), 469–483.
- [5] Heinz Bauer. 2011. *Measure and Integration Theory*. De Gruyter Studies in Mathematics, Vol. 26. De Gruyter, Berlin, Boston.
- [6] Giacomo Bucci, Laura Carnevali, Lorenzo Ridi, and Enrico Vicario. 2010. ORIS: A tool for modeling, verification and evaluation of real-time systems. *Int. J. Softw. Tools Technol. Transf.* 12, 5 (2010), 391–403.
- [7] Carlos E. Budde, Pedro R. D’Argenio, Arnd Hartmanns, and Sean Sedwards. 2018. A statistical model checker for nondeterminism and rare events. In *24th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’18) (LNCS)*, Vol. 10806. Springer, Cham, 340–358.
- [8] Xin Chen, Erika Ábrahám, and Sriram Sankaranarayanan. 2013. Flow*: An analyzer for non-linear hybrid systems. In *25th International Conference on Computer Aided Verification (CAV’13) (LNCS)*, Natasha Sharygina and Helmut Veith (Eds.), Vol. 8044. Springer, Berlin, 258–263.
- [9] George E. Collins and Hoon Hong. 1991. Partial Cylindrical Algebraic Decomposition for quantifier elimination. *J. Symb. Comput.* 12, 3 (1991), 299–328.
- [10] George B. Dantzig and B. Curtis Eaves. 2003. Fourier-Motzkin elimination and its dual. In *The Basic George B. Dantzig*, Richard W. Cottle (Ed.). Stanford University Press, Stanford, 255–264.
- [11] Daniel D. Deavours, Graham Clark, Tod Courtney, David Daly, Salem Derisavi, Jay M. Doyle, William H. Sanders, and Patrick G. Webster. 2002. The Möbius framework and its implementation. *Trans. Softw. Eng.* 28, 10 (2002), 956–969.
- [12] Mariken H. C. Everdij and Henk A. P. Blom. 2006. Hybrid Petri nets with diffusion that have into-mappings with generalised stochastic hybrid processes. In *Stochastic Hybrid Systems: Theory and Safety Critical Applications*, Henk A. P. Blom and John Lygeros (Eds.). LNCIS, Vol. 337. Springer, Berlin, 31–63.
- [13] Mariken H. C. Everdij and Henk A. P. Blom. 2008. Enhancing hybrid state Petri nets with the analysis power of stochastic hybrid processes. In *9th International Workshop on Discrete Event Systems (WODES’08)*. IEEE, 400–405.
- [14] Goran Frehse, Rajat Kateja, and Colas Le Guernic. 2013. Flowpipe approximation and clustering in space-time. In *16th ACM International Conference on Hybrid Systems: Computation and Control (HSCC’13)*. ACM, New York, 203–212.
- [15] Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. 2011. SpaceEx: Scalable verification of hybrid systems. In *23rd International Conference on Computer Aided Verification (CAV’11) (LNCS)*, Vol. 6806. Springer, Berlin, 379–395.
- [16] Martin Fränzle, Tino Teige, and Andreas Eggers. 2010. Engineering constraint solvers for automatic analysis of probabilistic hybrid automata. *J. Log. Algebr. Prog.* 79, 7 (2010), 436–466.
- [17] Komei Fukuda. 2017. Cdd library website. Retrieved from https://www.inf.ethz.ch/personal/fukudak/cdd_home/.
- [18] Mark Galassi, Jim Davies, James Theiler, Brian Gough, Gerard Jungman, Patrick Alken, Michael Booth, and Fabrice Rossi. 2018. *GNU scientific library. Release 2.5*. Network Theory Ltd.
- [19] Evgenij Gawrilow and Michael Joswig. 2000. polymake: A framework for analyzing convex polytopes. In *Polytopes - Combinatorics and Computation (DMV Seminar)*, Vol. 29. Birkhäuser, Basel, Switzerland, 43–73.
- [20] Hamed Ghasemieh. 2016. *Analysis of Hybrid Petri Nets with Discrete Random Events*. PhD Thesis. University of Twente, Enschede.
- [21] Hamed Ghasemieh, Anne Remke, and Boudewijn R. Haverkort. 2016. Survivability analysis of a sewage treatment facility using hybrid Petri nets. *Perf. Eval.* 97 (2016), 36–56.
- [22] Hamed Ghasemieh, Anne Remke, Boudewijn R. Haverkort, and Marco Gribaudo. 2012. Region-based analysis of hybrid Petri nets with a single general one-shot transition. In *10th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS’12) (LNCS)*, Vol. 7595. Springer, Berlin, 139–154.
- [23] Adrian Godde and Anne Remke. 2017. Model checking the STL time-bounded until on hybrid Petri nets using Nef polyhedra. In *14th European Workshop on Performance Engineering (EPEW’17) (LNCS)*, Vol. 10497. Springer, Cham, 101–116.
- [24] Marco Gribaudo and Anne Remke. 2016. Hybrid Petri nets with general one-shot transitions. *Perf. Eval.* 105 (2016), 22–50.
- [25] Marco Gribaudo, Matteo Sereno, András Horváth, and Andrea Bobbio. 2001. Fluid stochastic Petri nets augmented with flush-out arcs: Modelling and analysis. *Discr. Event Dynam. Syst.* 11, 1 (2001), 97–117.
- [26] E. Moritz Hahn, Arnd Hartmanns, Holger Hermanns, and Joost-Pieter Katoen. 2013. A compositional modelling and analysis framework for stochastic hybrid systems. *Form. Meth. Syst. Des.* 43, 2 (2013), 191–232.

- [27] Preston C. Hammer and A. Wayne Wymore. 1957. Numerical evaluation of multiple integrals I. *Math. Tables Aids Comput.* 11, 58 (1957), 59–67. DOI: <https://doi.org/10.2307/2002154>.
- [28] Arnd Hartmanns and Holger Hermanns. 2014. The modest toolset: An integrated environment for quantitative modelling and verification. In *10th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'14) (LNCS)*, Vol. 8413. Springer, Berlin, 593–598.
- [29] Graham Horton, Vidyadhar G. Kulkarni, David M. Nicol, and Kishor S. Trivedi. 1998. Fluid stochastic Petri nets: Theory, applications, and solution techniques. *Eur. J. Oper. Res.* 105, 1 (1998), 184–201.
- [30] Jannik Hüls. 2020. *Model Checking Hybrid Petri Nets with Multiple General Transition Firings*. PhD Thesis. University of Münster, Münster.
- [31] Jannik Hüls, Henner Niehaus, and Anne Remke. 2020. hpnmg: A C++ tool for model checking hybrid Petri nets with general transitions. In *12th International NASA Formal Methods Symposium (NFM'20) (LNCS)*, Vol. 12229. Springer, Cham, 369–378.
- [32] Jannik Hüls, Carina Pilch, Patricia Schinke, Joanna Delicaris, and Anne Remke. 2019. State-space construction of hybrid Petri nets with multiple stochastic firings. In *16th International Conference on Quantitative Evaluation of Systems (QEST'19) (LNCS)*, Vol. 11785. Springer, Cham, 182–199.
- [33] Jannik Hüls and Anne Remke. 2019. Model checking HPnGs in multiple dimensions: Representing state sets as convex polytopes. In *19th IFIP WG 6.1 International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE'19) (LNCS)*, Vol. 11535. Springer, Cham, 148–166.
- [34] Jannik Hüls, Stefan Schupp, Anne Remke, and Erika Ábrahám. 2018. Analyzing hybrid Petri nets with multiple stochastic firings using HyPro. In *11th EAI International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS'17)*. ACM, New York, 178–185.
- [35] Marijn R. Jongerden, Jannik Hüls, Anne Remke, and Boudewijn R. Haverkort. 2016. Does your domestic photovoltaic energy system survive grid outages? *Energies* 9, 9 (2016), 736–744.
- [36] Marta Z. Kwiatkowska, Gethin Norman, Roberto Segala, and Jeremy Sproston. 2002. Automatic verification of real-time systems with discrete probability distributions. *Theor. Comput. Sci.* 282, 1 (2002), 101–150.
- [37] G. Peter Lepage. 1978. A new algorithm for adaptive multidimensional integration. *J. Comput. Phys.* 27, 2 (1978), 192–203.
- [38] John McNamee and Frank Stenger. 1967. Construction of fully symmetric numerical integration formulas. *Numer. Math.* 10 (1967), 327–344.
- [39] Carina Pilch, Fabian Edenfeld, and Anne Remke. 2017. HYPEG: Statistical model checking for hybrid Petri nets. In *11th EAI International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS'17)*. ACM, New York, 186–191.
- [40] Carina Pilch and Anne Remke. 2017. Statistical model checking for hybrid Petri nets with multiple general transitions. In *47th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'17)*. IEEE, 475–486.
- [41] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. 1992. *Numerical Recipes in C: The Art of Scientific Computing* (2nd ed.). Cambridge University Press, Cambridge.
- [42] Stefan Schupp, Erika Ábrahám, Ibtissam B. Makhoul, and Stefan Kowalewski. 2017. HyPro: A C++ library of state set representations for hybrid systems reachability analysis. In *9th International NASA Formal Methods Symposium (NFM'17) (LNCS)*, Vol. 10227. Springer, Cham, 288–294.
- [43] Raimund Seidel. 1995. The upper bound theorem for polytopes: An easy proof of its asymptotic version. *Comput. Geom.* 5, 2 (1995), 115 – 116.
- [44] Sadegh E. Z. Soudjani and Alessandro Abate. 2013. Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM J. Appl. Dynam. Syst.* 12, 2 (2013), 921–956.
- [45] Sadegh Esmaeil Zadeh Soudjani, Caspar Gevaerts, and Alessandro Abate. 2015. FAUST2: Formal abstractions of uncountable-STate STochastic processes. In *21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'15) (LNCS)*, Vol. 9035. Springer, Berlin, 272–286.
- [46] Jeremy Sproston. 2000. Decidable model checking of probabilistic hybrid automata. In *6th International Symposium on Formal Techniques in Real-time and Fault-tolerant Systems (FTRFT'00) (LNCS)*, Vol. 1926. Springer, Berlin, 31–45.
- [47] Tino Teige and Martin Fränzle. 2009. Constraint-based analysis of probabilistic hybrid systems. *IFAC Proc.* Vol. 42, 17 (2009), 162–167.
- [48] The CGAL Project. 2017. *CGAL User and Reference Manual* (4.10 ed.). CGAL Editorial Board. Retrieved from <http://doc.cgal.org/4.10/Manual/packages.html>.
- [49] Csaba D. Toth, Joseph O'Rourke, and Jacob E. Goodman. 2017. *Handbook of Discrete and Computational Geometry* (3rd ed.). CRC Press, Boca Raton.
- [50] Liang Yin, Ricardo M. Fricks, and Kishor S. Trivedi. 2002. Application of semi-Markov process and CTMC to evaluation of UPS system availability. In *Reliability and Maintainability Symposium (RAMS'02)*. IEEE, 584–591.

- [51] Lijun Zhang, Zhikun She, Stefan Ratschan, Holger Hermanns, and E. Moritz Hahn. 2012. Safety verification for probabilistic hybrid systems. *Eur. J. Contr.* 18, 6 (2012), 572–587.
- [52] Günter M. Ziegler. 1995. *Lectures on Polytopes*. Graduate Texts in Mathematics, Vol. 152. Springer, New York, New York.

Received March 2020; revised September 2020; accepted February 2021