Decoupling Passenger Flows for Improved Load Prediction

Stefan Haar, Simon Theissing

Decoupling Passenger Flows for Improved Load Prediction

Stefan Haar and Simon Theissing

MExICo team, INRIA and LSV, CNRS & ENS de Cachan,
Cachan, France

Abstract. This paper continues our work on perturbation analysis of multimodal transportation networks (TNs) by means of a stochastic hybrid automaton (SHA) model. We focus here on the approximate computation, in particular on the major bottleneck consisting in the high dimensionality of systems of stochastic differential balance equations (SDEs) that define the continuous passenger-flow dynamics in the different modes of the SHA model. In fact, for every pair of a mode and a station, one system of coupled SDEs relates the passenger loads of all discrete points such as platforms considered in this station, and all vehicles docked to it, to the passenger flows in between. In general, such an SDE system has many dimensions, which makes its numerical computation and thus the approximate computation of the SHA model intractable. We show how these systems can be canonically replaced by lower-dimensional ones, by decoupling the passenger flows inside every mode from one another. We prove that the resulting approximating passenger-flow dynamics converges to the original one, if the replacing set of balance equations set up for all decoupled passenger flows communicate their results among each other in vanishing time intervals.

Keywords: Stochastic Hybrid Automata, Transportation Networks

1 Introduction

Apart from some exceptions, the different modes and lines in modern multimodal transportation networks do not share infrastructure elements, but are loosely connected through passenger transfers. Understanding how these passenger transfers connect their modes and lines is thus crucial if one wants to analyse how perturbations spread across such TNs. In this context, the present work is a contribution to our SHA model from [6] that we have developed for the computation of passenger load forecasts in multimodal TNs; given (i) estimations for all uncertain initial passenger loads (platforms, vehicles, etc.) and uncertain *continuous* passenger arrival flows, and (ii) the possibility to track individual vehicles so as to study the impact of well-directed interventions to their operation such as early departures.

Our SHA Model. Our SHA model from [6] extends a previous *deterministic* hybrid automaton (DHA) model from [4]. In this, a finite set of vehicles is operated, and every vehicle is confined to a particular mode or line which does not share infrastructure elements with any other modes or lines. Passengers are grouped into a finite set of *trip profiles*, which define routes in the TN at hand, together with preferences for choosing different vehicle missions. Every mode of the DHA model corresponds to a particular configuration of the vehicles' discrete positions and discrete operational states. With these parameters, every mode defines which passenger flows between stations and vehicles are possible. In this way, a system of coupled ordinary differential equations (ODEs), one equation per station, is associated to every mode. This system relates the passenger loads of all stations and of all *stopped* vehicles docked to these stations, to the passenger flows such as boarding and alighting in between. Transitions *between* modes are triggered by (i) vehicles that *must depart*, i.e. whose elapsed driving and dwell times exceed some deterministic thresholds fixed by operation rules, and (ii) by passenger load trajectories hitting some pre-defined regions and thus triggering the departure of some vehicle (examples: boarding a train must stop if the train is full, or if the number of passengers on the platform is small and the train is scheduled to leave, etc.).

Now a TN is everything but deterministic: The influx of passengers into the system is a random process (from a macroscopic point of view, in fact a very continuous and measurable random process as compared to e.g. single passenger incidents), and the distribution of the passengers over the different possible trip profiles - is also unknown and can only be given statistically. This motivated the stochastic hybrid automaton (SHA) model that we introduced in [6]: Compared to our above DHA model, we replaced all systems of ODEs by systems of (Itô-) stochastic differential equations (SDEs), so as to be able to (i) start our analyses with uncertain initial passenger loads, and (ii) include uncertain passenger arrival flows into the model's continuous time dynamics. The mechanism of triggering mode transitions via thresholds remains the same; however, these hitting times are not deterministic, isolated points in time any more, but rather random variables with a continuous range of values.

Our SHA model thus does not fully cover the dynamical spectrum of the stochastic hybrid system (SHS) from [8], but only implements a particular realization thereof: In our SHA model, there are no mode transitions which are exponentially distributed w.r.t. time. In this context, also note that the SHS from [8] is an abstract mathematical model for a system with a mixed discrete and continuous dynamics; no more no less. The definition of e.g. all vector fields or possible mode transitions therein might be non-trivial and often cannot be done by pen and paper. That is why, we employ artefacts from the Petri nets formalism so as to e.g. derive all differential balance equations in a canonical way; which was proposed in many papers such as [10] before.

Problem Formulation. In [5], we introduced a strategy for the approximate computation of our above SHA model: We let the automaton change its mode only at equidistantly-spaced discrete points in time. Several challenges then arise. On the one hand, we are confronted with an explosion of the SHA model's timed mode graph, that - as its name suggests - captures the evolution of the SHA model's mode in discrete time steps; but we do not consider this combinatorial problem here, it will be treated in another work. Rather, our present paper focusses on another major bottleneck, namely the high dimensionality of the SDE systems defining the passenger flow dynamics in any given mode. The dimension of theses systems of coupled SDEs that we set up for every pair (mode, station) in the SHA model from [6] corresponds to the number of the passengers' different trip profiles, multiplied by the number of different discrete positions for the passengers within this station and the vehicles docked to it. Our major concern with this high dimensionality then is the fact that all algorithms that we have found so far are prone to what is known as the curse of dimensionality.

Simulation of SDEs. Monte Carlo simulations [9] require to sample realizations of the uncertain initial states of the considered random variables. For one-dimensional RVs subjected to one-dimensional SDEs this sampling might be trivial e.g. by employing the inverse transform sampling. However, it seems that sampling the uncertain initial state of multidimensional RVs is a non-trivial task that is active and still an open problem. Among the algorithms proposed thus far, we mention the Metropolis-Hastlings and the Gibbs sampler, which can be integrated into what is called a Markov Chain Monte Carlo simulation [1]. Other more exotic sampling techniques might involve e.g. neural networks [7].

Analytic Methods. Instead of sampling as above, another approach that we shall study elsewhere is to numerically integrate a multivariate Fokker-Planck equation. Such a system of partial *ordinary* differential equations is derived from the original multidimensional SDE, and describes the time evolution of an initial probability density function (PDF) under the system's dynamics; here, it concerns the passenger load vector's density function, giving the distribution of the number of passengers in the different trip profiles. However, many computational drawbacks also come along with this method, or more specifically with the numerical integrations required. First, not all numerical integration schemes can ensure the conservation of the probability flux in their basic set up; with the Finite Volume method [2] being one exception. Second, those schemes which can ensure the conservation of the probability flux are not easily extendible from common two or three dimensional applications to higher-dimensional problems.

Alternative Approaches. Alternatives to the computation or simulation of high-dimensional SDEs might involve their discrete approximation, which we do not pursue here. The technique studied here aims at decoupling the dynamics in the SDEs, so as to produce an alternative set of lower-dimensional SDEs that reproduces, or at least approximates, the original model dynamics. For instance, the authors of [3] mention the *local* specification of flows in a fluid stochastic Petri net model as a means for the decoupling. However, in contrast to our approach, they look at *scalar* rather than vectorial (passenger) flows.

In the rest of this paper, we shortly review our SHA model from [6] in Sec. 2 together with the discrete time computation of its state space from [5]. We also discuss the set up of all high-dimensional SDEs for the passenger flow dynamics in the SHA model's different modes. We then explain in Sec. 3 how the passenger flows in all modes can be systematically decoupled so as to replace the original systems of SDEs by approximating lower-dimensional ones. In this context, we also proof asymptotic convergence of the dynamics produced by the lower-dimensional SDEs w.r.t. the original dynamics. Last but not least, we summarize the contribution of our approach, and give a brief outlook on future work in Sec. 4.

2 Our SHA Model

2.1 Model Structure

Infrastructure. Basic modelling blocks of the SHA model are place/transition nets (= Petri nets with the token flow left out), which capture the structure of a finite set of stations $\mathcal{S}$ and a finite set of transportation grids $\mathcal{G}$ (TGs).

Every station $s \in \mathcal{S}$ is made up of a finite set P_s of gathering points $p \in P_s$ (= places; represented by double circles) that can accommodate a limited number of passengers, and a finite set T_s of corridors $t \in T_s$ (= transitions; represented by double boxes) connecting (i) GPs to other GPs, or (ii) GPs to the station's exterior (cf. Fig. 1 below). Here, connected means "possibility of a passenger flow" in the direction of the edges that connect the corridors with the GPs.

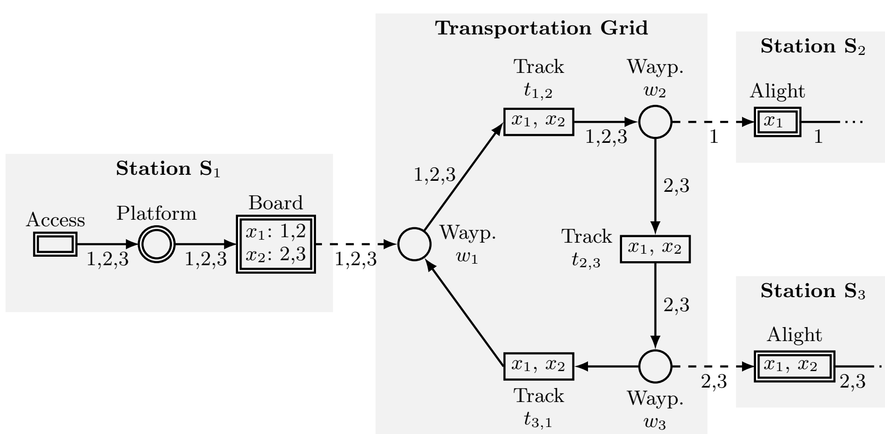

Fig. 1. Representation of the infrastructure of a sample TN in our SHA model, together with (i) the paths of two different vehicle missions x_1 and x_2, and (ii) an indication of the stops along these paths for the specification of three different trip profiles (TPs).

Every TG $g \in \mathcal{G}$ captures the structure of a particular mode or line; and in doing so, all possible vehicle movements between its finite set W_g of discrete waypoints $w \in W_g$ (= places; represented by simple circles) which accommodate the vehicle tokens (at maximum one vehicle per waypoint) via tracks (= transitions; represented by simple boxes).

A finite set of tuples $(a, b) \in \mathcal{I}$, with $\mathcal{I} \subseteq (T \times W) \cup (W \times T)$, $T := \bigcup_{s \in \mathcal{S}} T_s$ and $W := \bigcup_{g \in \mathcal{G}} W_g$, composed of a transition in a station and a waypoint in a TG, defines the interface between the stations and the TGs (represented by dashed arcs in Fig. 1 above): Every tuple $(a, b) \in \mathcal{I}$ either connects some GP in a station $s \in \mathcal{S}$ to a waypoint in a TG $g \in \mathcal{G}$, in which case $a \in P_s$ and $b \in W_g$; or vice versa. In this way, every tuple defines which passenger flow between a vehicle stopped at a waypoint in a TG and a GP (= platform) in a station is possible for the purpose of boarding & alighting; see below.

Vehicle Operation. At the heart of the operation of a finite set $\mathcal{V}$ of all vehicle tokens $v \in \mathcal{V}$ considered in the SHA model are missions: Every mission defines a path in a particular transportation grid, together with (i) a sequence of stops at the waypoints along that path; (ii) deterministic-timed (minimum & maximum dwell times) and passenger load-dependent departure conditions from the stops which might state that a vehicle cannot depart from a stop as long as some passengers still want to alight from or board it; and (iii) driving times between all waypoints which might be functions of the positions of all vehicle tokens.

Passenger Routing. We group all passengers into a finite set $\mathcal{Y} := \{1, 2, \ldots, n\}$ of $n \in \mathbb{N}$ different trip profiles (TPs): Every $y \in \mathcal{Y}$ defines a particular path in TN's infrastructure, together with the passengers' preferences for the different vehicle missions (cf. Fig. 1 above). However, this does not mean that the passengers cannot change their TPs as we will highlight in a short (see Sec. 2.3).

2.2 Hybrid State

As common in the literature of hybrid automata, we refer to the discrete state of our SHA model at any time $\tau \geq 0$ as its mode: A particular mode $q \in Q$ from the finite set of all different modes Q defines for every $v \in \mathcal{V}$ (i) the position of v in form of a waypoint in a TG; (ii) the driving condition of v which is either parked, stopped or driving; (iii) the operational state of v in form of a mission to be executed, a discrete state therein, and a sequence of missions to be accomplished. Thus, every $q \in Q$ tells us which vehicle is docked to which station; and in doing so, defines the (continuous) passenger flow dynamics in TN.

Remark 1. We say that a vehicle $v \in \mathcal{V}$ is docked to a station $s \in \mathcal{S}$ iff (i) v is stopped at a waypoint $w \in W_g$ in some TG $g \in \mathcal{G}$; (ii) acc. to $\mathcal{I}$, either passengers can board v stopped at w from some GP in s, or alight from it to some GP in s. Moreover, we denote by $\mathcal{V}(s, q) \subseteq \mathcal{V}$ the subset of all vehicles that are docked to s in q.

Remark 2. If k is a row (column) vector, then we denote by $k[i]$ the element in its i-th column (row).

The continuous state of the SHA model at any $\tau \geq 0$, defines (i) the elapsed dwell times of all stopped vehicles, (ii) the elapsed driving times of all moving vehicles, and (iii) the passenger load $\mathrm{M}(b, \tau)$, with $\mathrm{M} : (b, \tau) \in (P \cup \mathcal{V}) \times \mathbb{R}_{\geq 0} \to \mathcal{M}(b)$ and

$$\mathcal{M}(b) := \left\{ k \in (\mathbb{R}_{\geq 0})^{|\mathcal{Y}|} : \sum_{i=1}^{|\mathcal{Y}|} k[i] \leq c(b) \right\}, \tag{1}$$

for every vehicle $b \in \mathcal{V}$ and every GP in a station $b \in P$. Therein, $P := \bigcup_{s \in \mathcal{S}} P_s$, $\mathrm{M}(b, \tau)[i]$ gives the number of passenger at/on-board b, who travel acc. to the TP $i \in \mathcal{Y}$, and $c(b)$, with $c : P \cup \mathcal{V} \to \mathbb{R}_{>0}$, gives the maximum number of passengers b can accommodate at the same time.

2.3 Balance Equations

For any $q \in Q$, we adapt the notation ${}^{\bullet}b(q)$ for the preset and $b^{\bullet}(q)$ for the postset of any $b \in P \cup \mathcal{V}'(q)$, with $\mathcal{V}'(q) := \bigcup_{s \in \mathcal{S}} \mathcal{V}(s, q)$, from the Petri nets literature for our purposes: ${}^{\bullet}b(q)$ denotes the set of all corridors in the stations that are connected by an arc pointing towards b. Accordingly, $b^{\bullet}(q)$ denotes the set of all corridors in the stations that are connected by an arc pointing away from b. For $b \in \mathcal{V}'(q)$, those arcs (dashed arcs in Fig. 1 above) point towards/away from the waypoint which accommodates b.

Note that all corridors in the stations of our SHA model are connected in a special way to the rest of the modelled infrastructure (GPs in the stations and waypoints in the TGs).

Remark 3. For any $t \in T$, we denote by ${}^{\star}t(q) := b$ the single GP in a station or vehicle docked to a station $b \in P \cup \mathcal{V}'(q)$ which is connected to t in q by an arc pointing towards t iff $t \in b^{\bullet}(q)$. Accordingly, we denote by $t^{\star}(q) := a$, for any $t \in T$, the single GP or vehicle docked to a station $a \in P \cup \mathcal{V}'(q)$ which is connected to t in q by an arc pointing away from t iff $t \in {}^{\bullet}a(q)$.

This special structure allows us to decompose all corridors in $q \in Q$ into three disjoint sets; implementing inflows, transfer flows, and outflows: Inflows model the arrival processes of the passengers who join the SHA model from TN's exterior.

Definition 1 (Inflow). *An inflow is a passenger flow assigned to any $t \in T_1$, with*

$$\begin{aligned} T_1 := \big\{ t \in T : \exists\, p \in P \ \ s.t.\ t \in {}^{\bullet}p \,\wedge \\ \nexists\, p' \in P \ \ s.t.\ t \in p'^{\bullet} \wedge \nexists\, w \in W \ \ s.t.\ \ (w, t) \in \mathcal{I} \big\}. \end{aligned} \tag{2}$$

Transfer flows model passenger flows within the SHA model; including passenger transfers between the GPs in the stations, as well as passenger transfers between GPs in the stations and vehicles docked to the stations.

Definition 2 (Transfer Flow). *A transfer flow in $q \in Q$ is a passenger flow assigned to any $t \in T_2\left(q\right)$, with*

$$T_2\left(q\right) := \big\{t \in T : \exists\, b \in P \cup \mathcal{V}\left(q\right) \;\; s.t.\; t \in {}^{\bullet}b \wedge \\ \exists\, b' \in P \cup \mathcal{V}'\left(q\right) \;\; s.t.\; t \in \left(b'\right)^{\bullet}\big\}. \tag{3}$$

Finally, outflows model the SHA model's drain of passengers to TN's exterior.

Definition 3 (Outflow). *An outflow is a passenger flow assigned to any $t \in T_3$, with*

$$T_3 := \big\{t \in T : \exists\, p \in P \;\; s.t.\; t \in p^{\bullet} \wedge \\ \nexists\, p' \in P \;\; s.t.\; t \in p^{\bullet} \wedge \nexists\, w \in W \;\; s.t.\; \left(t, w\right) \in \mathcal{I}\big\}. \tag{4}$$

With that said, we denote by $T'\left(q\right)$, with $T'\left(q\right) := T_1 \cup T_2\left(q\right) \cup T_3$, the set of all corridors active in $q \in Q$; and by $\gamma\left(\tau\right)$, with $\gamma : \mathbb{R}_{\geq 0} \to Q$, the mode of our SHA model at time $\tau \geq 0$.

$$\mathrm{dM}\left(b, \tau\right) := \sum_{t \in {}^{\bullet}b \cap T'\left(\gamma\left(\tau\right)\right)} \mathrm{R}\left(t\right) \overbrace{\left[\phi\left(t, \tau\right)\, \mathrm{d}\tau + \delta\left(t\right)\, \mathrm{d}\mathcal{W}\left(\tau\right)\right]}^{\text{Passenger flow into b}} - \\ \sum_{t \in b^{\bullet} \cap T'\left(\gamma\left(\tau\right)\right)} \underbrace{\left[\phi\left(t, \tau\right)\, \mathrm{d}\tau + \delta\left(t\right)\, \mathrm{d}\mathcal{W}\left(\tau\right)\right]}_{\text{Passenger flow leaving b}} \tag{5}$$

then defines the time evolution of the passenger load of every GP in a station and of every vehicle docked to a station $b \in P \cup \mathcal{V}'\left(q\right)$ at any time $\tau \geq 0$ when the SHA model is in $q \in Q$. This balance equation relates $\mathrm{M}\left(b, \tau\right)$ to all passenger flows into b and leaving it: We capture the routing of all passengers along the different TPs as well as their local re-routing among these TPs in so-called routing matrices.

Remark 4. We denote by $\Psi^{d_1 \times d_2}$, for some $d_1, d_2 \in \mathbb{N}_{>0}$ and any set Ψ, the set of all matrices with d_1 rows and d_2 columns, whose elements are from Ψ. In the case that $d_2 = 1$, we drop d_2 in $\Psi^{d_1 \times d_2}$ and write Ψ^{d_1} instead.

The i-th row and the j-th column of a particular routing matrix $\mathrm{R}\left(t\right)$ assigned to $t \in T$, with

$$\mathrm{R} : T \to \left\{K \in \left(\mathbb{R}_{\geq 0}\right)^{|\mathcal{Y}| \times |\mathcal{Y}|} : \sum_{i=1}^{|\mathcal{Y}|} K[i, j] = 1, \forall j = \mathcal{Y}\right\},$$

defines the relative amount of the flow of passengers who join t acc. to the TP $j \in \mathcal{Y}$, and who leave t acc. to the TP $i \in \mathcal{Y}$; and the fact that every column of $\mathrm{R}\left(t\right)$ must either sum up to one or to zero, implies that all passenger flows are conserved.

Remark 5. Time could be included in the domain of the routing matrices above so that they might change values during mode transitions of the SHA model depending on the hybrid state; so as to account e.g. for loudspeaker announcements.

We next write down the passenger flow assigned to every corridor $t \in T(q)$ in q acc. to its impact on $\mathrm{M}(p,\tau)$ as the sum of a drift term $\phi(t,\tau)$, with

$$\phi : (t,\tau) \in \bigcup_{q \in Q} T'(q) \times \mathbb{R}_{\geq 0} \to \left\{ v \in (\mathbb{R}_{\geq 0})^{|\mathcal{Y}|} : \sum_{i=1}^{|\mathcal{Y}|} v[i] \leq \phi_{\max}(q,t) \right\},$$

and a constant diagonal diffusion term

$$\delta : \bigcup_{q \in Q} T'(q) \to \left\{ K \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|} : K[i,j] = 0, \forall i \neq j \right\}.$$

Therein, $\phi_{\max}(q,t)$, with $\phi_{\max} : q \in Q \times T'(q) \to \mathbb{R}_{\geq 0}$, is the maximum passenger throughput of the corridor $t \in T'(q)$, when the SHA model is in $q \in Q$.

Remark 6. Let X be a continuous RV. Then, $\mathrm{pdf}(X)$ denotes its PDF; $\sigma(X)$ denotes its state space; and $\mathrm{pdf}(X,x)$ denotes the evaluation of $\mathrm{pdf}(X)$ at x for some $x \in \sigma(X)$.

We discuss the specification of $\phi(\cdot)$ and $\delta(\cdot)$ in more detail in the rest of this paper. Here, only note that the drift term of a flow into some $b \in P \cup \mathcal{V}'(q)$ shifts the density of $\mathrm{M}(b,\tau)$ in its domain. The flow's diffusion term narrows or broadens the density of $\mathrm{M}(b,\tau)$.

2.4 Grouping of Balance Equations

In principle, the passenger flows in (5) can be defined as any functions of the SHA model's complete hybrid state as long as they are capacity- and demand-sensitive; crucial properties that we assume for all passenger flows in our SHA model: We say that some passenger flow is capacity-sensitive iff its drift does not cause the passenger load of some GP or vehicle to exceed the capacity limit of that GP or vehicle.

Definition 4 (Capacity-Sensitive Flow). *A passenger flow assigned to some $t \in T'(q)$ in $q \in Q$ is capacity-sensitive iff $t \in T_3$ or*

$$\sum_{i=1}^{|\mathcal{Y}|} \mathrm{M}(t^\star,\tau)[i] \to c(t^\star)$$

implies that $\phi(t,\tau) \to 0$ for any $\tau \geq 0$.

Additionally, we say that a passenger flow is demand-sensitive iff its drift does not cause any passenger load to become negative.

Definition 5 (Demand-Sensitive Flow). *A passenger flow assigned to some $t \in T'(q)$ in $q \in Q$ is demand-sensitive iff $t \in T_1$ or*

$$\mathrm{M}\left({}^{\star}t, \tau\right)[j] \sum_{i=1}^{|\mathcal{Y}|} \mathrm{R}(t)[i, j] \rightarrow 0$$

implies that $\phi(t, \tau)[j] \rightarrow 0$ for all $j \in \mathcal{Y}$ and for any $\tau \geq 0$.

Remark 7. Def. 4 and 5 taken alone cannot ensure the non-negativity and capacity limits of the passenger loads assuming non-zero diffusion terms in (5). Instead both properties must be explicitly ensured during the computation or simulation of (5) in form of reflecting boundary conditions. See e.g. [6], where we derive reflecting boundary conditions for the numerical integration of a multivariate Fokker-Planck equation obtained from (5).

For our purposes however, we do not need this kind of global inclusion of the SHA model's complete hybrid state into the specification of the passenger flows: We restrict the domains of their drift terms to the passenger loads in their presets and postsets.

Definition 6 (Local Flow). *A passenger flow assigned to some $t \in T'(q)$ in $q \in Q$ is local iff for any $\tau \geq 0$,*

- *$t \in T_1$, and the flow's drift term only depends on $\mathrm{M}(t^{\star}, \tau)$, or*
- *$t \in T_2(q)$, and the flow's drift term only depends on $\mathrm{M}({}^{\star}t, \tau)$ and $\mathrm{M}(t^{\star}, \tau)$, or*
- *$t \in T_3$, and the flow's drift term only depends on $\mathrm{M}({}^{\star}t, \tau)$.*

This *local specification* of all passenger flows produces a natural decomposition of all SDEs set up for any $q \in Q$: The balance equations in form of (5) set up for the passenger loads of all GPs $p \in P_s$ and vehicles $v \in \mathcal{V}(s, q)$, for some station $s \in \mathcal{S}$, are independent from the passenger loads of all GPs outside s and vehicles not docked to s. We can thus group them into one common system of coupled SDEs of dimension $k := (|P_s| + |\mathcal{V}(s, q)|)\ |\mathcal{Y}|$, which latter system is decoupled from those systems set up for all other stations.

Remark 8. In practice, we do only have to consider all those TPs in the domain specification for the passenger load of a particular GP or vehicle, whose paths cover this GP or vehicle. Thus, k as defined above only defines an upper bound for the dimension of the system of SDEs set up for s in q.

2.5 Mode Transitions

We assume that at the initial simulation time $\tau = 0$, with $\tau \geq 0$, our SHA model is in one particular mode with marginal probability one, and we know the elapsed driving & dwell times of all vehicles. We then let our SHA model transition between its discrete modes only at discrete time steps $\tau = i\,{}_{\Delta}\tau$, with $i \in \mathbb{N}_{>0}$, of fixed length ${}_{\Delta}\tau > 0$. In this context, we also let the elapsed driving

& dwell times of all vehicles only evolve at $\tau = i \,{}_\Delta\tau$ by ${}_\Delta\tau$. A directed acyclic graph (DAG) then captures the time evolution of our SHA model's vehicle load (= particular mode and particular realization of all elapsed discrete driving & dwell times). We do not go into details of its computation here, but only stress some important points. Refer to [5] for more information: Every node, say m, in this DAG, say G, represents a particular vehicle load for our SHA model in the half-closed time interval $\big[h_m \,{}_\Delta\tau, (h_m + 1)\, {}_\Delta\tau\big)$ iff $h_m \in \mathbb{N}_{\geq 0}$ is the height of m in G. Thus, two nodes with the same height $h' \in \mathbb{N}_{>0}$ in G represent two alternatives for our SHA model's vehicle load in $\big[h' \,{}_\Delta\tau, (h' + 1)\, {}_\Delta\tau\big)$. Two or more branches away from m indicate the possibility of mode transitions; with one branch for every alternative mode transition, and one additional branch for the continuation of m-th mode. Several nodes with the same height in G can have the same mode and thus the same passenger flow dynamics in common.

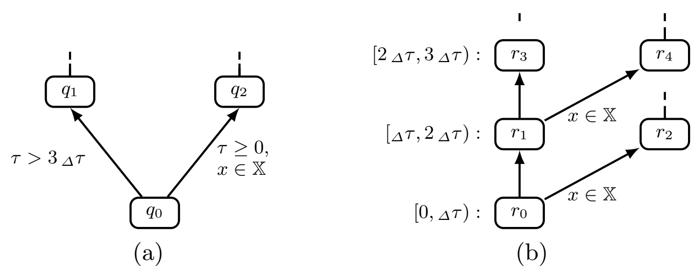

Fig. 2. Schematic comparison of a (classical) mode graph (a) and a timed mode graph (b) for our SHA model: $\mathbb{X}$ denotes a compact region in the SHA model's complete passenger load space as entrance condition for a not further specified passenger load-driven mode transition, and ${}_\Delta\tau > 0$ is the fixed time step that separates every pair of two consecutive time layers when the SHA model can change its mode

2.6 Propagation of Passenger Loads

At any simulation time $\tau = i \,{}_\Delta\tau$, with $i \in \mathbb{N}_{\geq 0}$ and ${}_\Delta\tau > 0$, one single marginal joint PDF, say $\mathrm{pdf}\,(i)$, defines the passenger loads of all GPs in the stations and of all vehicles. For $i = 0$, we assume that $\mathrm{pdf}\,(i)$ is known with marginal probability one. Then, starting from $i = 0$, all passenger loads have to be propagated forward in time from one time layer in the SHA model's DAG to the next: For the computation of $\mathrm{pdf}\,(i + 1)$, for some $i \in \mathbb{N}_{\geq 0}$, all high-dimensional systems of SDEs defined by our SHA model's different modes in the time layer $\big[i \,{}_\Delta\tau, (i + 1)\, {}_\Delta\tau\big)$ of the DAG, must be computed from $\tau = i \,{}_\Delta\tau$ to $\tau = (i + 1)\, {}_\Delta\tau$ with $\mathrm{pdf}\,(i)$ as common initial PDF. Depending on the particular use case at hand so as to e.g. forecast the risk of overcrowded platforms, this forward propagation is normally terminated once the simulation time exceeds some constant threshold. Refer to [5] for more details.

3 The Decoupling of All Passenger Flows

3.1 Overview

Our decoupling approach is perhaps best described by the following sequence of images: We assume that every GP in a station and every vehicle $b \in P \cup \mathcal{V}$ has the shape of a circular area, say A_b. We next assume that the passenger load of b is equally distributed on A_b at any simulation time step $\tau = i \,{}_{\Delta}\tau$, with $\tau \geq 0$, $i \in \mathbb{N}_{\geq 0}$, and ${}_{\Delta}\tau > 0$; in which ${}_{\Delta}\tau$ is the fixed time step that separates every pair of two consecutive time layers confining all mode transitions.

Remark 9. We denote by $\Gamma(\tau)$, with $\Gamma : \mathbb{R}_{\geq 0} \to 2^Q \setminus \emptyset$, the subset of all modes our SHA model can be in at time $\tau \in \mathbb{R}_{\geq 0}$.

For any time $\tau \in \mathcal{H}_i$, from the time interval $\mathcal{H}_i := \left[i \,{}_{\Delta}\tau, (i+1) \,{}_{\Delta}\tau\right)$, any mode $q \in \Gamma(\tau)$, and any $b \in P \cup \mathcal{V}'(q)$, we divide A_b into $|({}^{\bullet}b \cup b^{\bullet}) \cap T'(q)|$ non-overlapping slices (cf. Fig 3 below); in which one slice is attributed to every passenger flow into or leaving b, i.e., the passenger flow assigned to every corridor $t \in ({}^{\bullet}b \cup b^{\bullet}) \cap T'(q)$. Our assumptions above then imply that at $\tau = i \,{}_{\Delta}\tau$ (i) the surface area of a particular slice defines how many passengers it accommodates at τ, and (ii) the distribution of this latter number of passengers w.r.t. the passengers' different TPs is identical to the distribution of the total number of passengers at b and τ w.r.t. the different TPs. We moreover assume that a retractable wall is installed along every frontier separating two neighbouring slices (dashed lines in Fig. 3 below). These walls prevent the equidistant re-distribution of the slices' passenger loads at any $\tau \in \mathcal{H}_i$, which *diffusion* is restricted to the discrete time step $\tau = (i+1) \,{}_{\Delta}\tau$ when all walls are removed.

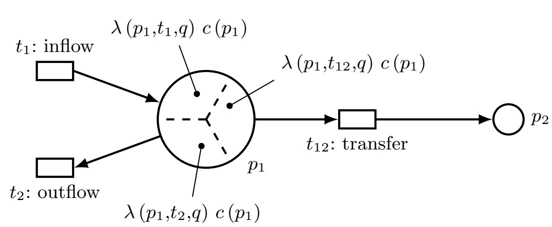

Fig. 3. Schematic representation of our decoupling approach: all GPs and vehicles docked to the stations in a particular mode, say q, of the SHA model are divided into slices, with impenetrable walls separating neighbouring slices until the next discrete point in time, say τ, arrives when the SHA model can change its mode. As long as the SHA model stays in q, all passengers flow into or out of the slices. They do not flow into or out of the original GPs and vehicles. A re-distribution of the slices' passenger loads occurs at τ.

So in our physically-touched model above, the slices' passenger loads are decoupled at any $\tau \in \mathcal{H}_i$, which implies that they might be filled and emptied at different rates if we assume that the passengers flow into and leave the slices of b; instead of flowing into and leaving b itself. For the specification of the slices' surface areas, we use the maximum passenger throughputs assigned to the corridors for the different modes; see below.

3.2 Decoupled Balance Equations

General Structure. The system of SDEs that we will set up for the decoupled passenger flow assigned to every $t \in T'(q)$ for any $q \in Q$ next, defines how this flow manipulates the passenger load $\mathrm{M}_{q,t}(^{\star}t, \tau)$ of the isolated slice from $^{\star}t$ attributed to t in q and/or the passenger load $\mathrm{M}_{q,t}(t^{\star}, \tau)$ of the isolated slice from $t^{\star}$ attributed to t in q; when our SHA model is in q. We write it down in the very general form of

$$\mathrm{d}X_{q,t}(\tau) := A_{q,t}(X_{q,t}(\tau))\,\mathrm{d}\tau + B_{q,t}(X_{q,t}(\tau))\,\mathrm{d}\mathcal{W}(\tau), \tag{6}$$

with the state vector $X_{q,t}$, the drift vector $A_{q,t}$, the diffusion matrix $B_{q,t}$, and the vector of $|\mathcal{Y}|$ uncorrelated Wiener processes $\mathcal{W}$.

Remark 10. We write the tuple of a mode $q \in Q$ and a transition $t \in T'(q)$ in form of subscript separating both in the given order by a comma next to a variable or constant iff we refer to the projection of that variable or constant in (6) set up for the decoupled passenger flow assigned to t in q.

Projection of Passenger Loads & Flows. As outlined in the figurative overview of our decoupling approach above, we project $\mathrm{M}(b, \tau)$, for any $b \in P \cup \mathcal{V}'(q)$ and $q \in Q$, to $\mathrm{M}_{q,t}(t, \tau)$, with $\mathrm{M}_{q,t} : T'(q) \times \mathbb{R}_{\geq 0} \rightarrow \mathcal{M}_{q,t}(b)$ and

$$\mathcal{M}_{q,t}(b) := \left\{ k \in (\mathbb{R}_{\geq 0})^{|\mathcal{Y}|} : \sum_{i=1}^{|\mathcal{Y}|} k[i] \leq \lambda(b, t, q)\; c(b) \right\},$$

at $\tau = i\,_{\Delta}\tau$, with $i \in \mathbb{N}_{\geq 0}$, acc. to

$$\mathrm{M}_{q,t}(b, i\,_{\Delta}\tau) := \lambda(b, t, q)\; \mathrm{M}(b, i\,_{\Delta}\tau) \tag{7}$$

iff our SHA model is in mode q at $\tau = i\,_{\Delta}\tau$. Therein, $\lambda(b, t, q)$, with

$$\lambda(b, t, q) := \frac{\phi_{\max}(q, t)}{\sum\limits_{t' \in (^{\bullet}b \cup b^{\bullet}) \cap T'(q)} \phi_{\max}(q, t')}, \tag{8}$$

defines the maximum number $\lambda(b, t, q)\; c(b)$ of passengers the isolated slice from $b \in P \cup \mathcal{V}'(q)$ assigned to $t \in (^{\bullet}b \cup b^{\bullet}) \cap T'(q)$ in q can accommodate (cf. Fig. 3 above). This simple projection implies

$$\mathrm{pdf}(\mathrm{M}_{q,t}(b, i\,_{\Delta}\tau) = \lambda(b, t)\; k) = \mathrm{pdf}(\mathrm{M}(b, i\,_{\Delta}\tau) = k), \forall k \in \mathcal{M}(b), \tag{9}$$

with $\mathcal{M}(b)$ from (1). We also use (8) to project $\phi(t, \tau)$ - which we assume to be local, demand- & capacity sensitive - to $\phi_{q,t}(t, \tau)$ acc. to Tab. 1 below, which implies that all qualitative properties of $\phi(t, \tau)$ such as demand-sensitiveness are adopted by $\phi_{q,t}(t, \tau)$.

Table 1. Specification of $\phi_{q,t}(t,\tau)$ assigned to $t \in T'(q)$ in $q \in Q$

Inflow:	$\phi\left(\lambda^{-1}(t^{\star},t,q)\ \mathrm{M}_{q,t}(t^{\star},\tau)\right)$
Transfer Flow:	$\phi\left(\lambda^{-1}({}^{\star}t,t,q)\ \mathrm{M}_{q,t}({}^{\star}t,\tau),\lambda^{-1}(t^{\star},t,q)\ \mathrm{M}_{q,t}(t^{\star},\tau)\right)$
Outflow:	$\phi\left(\lambda^{-1}({}^{\star}t,t,q)\ \mathrm{M}_{q,t}({}^{\star}t,\tau)\right)$

Inflows. In general, we neither know the passengers' exact arrival times, nor the TPs of the new arriving passengers. However, in most situations we know some reference values, and we can estimate quite reasonably fluctuations around them (e.g. from statistical considerations); which latter knowledge we can then map to the systems of SDEs set up for all decoupled inflows. More specifically, we set up for every $t \in T_1$ a balance equation in form of (5), which defines the impact of the inflow assigned to t, to the passenger load of $t^{\star}$; and integrate this balance equation into (6). Tab. 2 lists the corresponding ingredients.

Transfer Flows. Once having joined the SHA model, we assume that the passenger transfer dynamics regarded in isolation within the SHA model in a particular mode is deterministic; which implies zero diffusion terms for the specification of all decoupled passenger transfer flows: For every $t \in T_2(q)$ in $q \in Q$, we set up two balance equations in form of (5). The first balance equation defines the impact of the transfer flow assigned to t, to the passenger load of ${}^{\star}t$. Accordingly, the second balance equation relates the passenger load of $t^{\star}$ to the same decoupled transfer flow. We then integrate both balance equations into (6) acc. to Tab. 2.

Table 2. Specification of the system of SDEs set up for the decoupled inflow, transfer flow, or outflow assigned to $t \in T'(q)$ in mode $q \in Q$ of our SHA model

	Inflow	**Transfer Flow**	**Outflow**
Schematic structure	$t\ \square \rightarrow \bigcirc\ t^{\star}$	${}^{\star}t\ \bigcirc \rightarrow \overset{t}{\square} \rightarrow \bigcirc\ t^{\star}$	${}^{\star}t\ \bigcirc \rightarrow \square\ t$
$X_{q,t}(\tau)$	$\mathrm{M}_{q,t}(t^{\star},\tau)$	$\begin{bmatrix}\mathrm{M}_{q,t}({}^{\star}t,\tau)\\ \mathrm{M}_{q,t}(t^{\star},\tau)\end{bmatrix}$	$\mathrm{M}_{q,t}({}^{\star}t,\tau)$
$A_{q,t}(\tau)$	$\mathrm{R}(t)\ \phi_{q,t}(t,\tau)$	$\begin{bmatrix}-\phi_{q,t}(t,\tau)\\ \mathrm{R}(t)\ \phi_{q,t}(t,\tau)\end{bmatrix}$	$-\phi_{q,t}(t,\tau)$
$B_{q,t}$	$\delta(t)$	0	0

Outflows. Similar to the specification of all transfer flows above, we demand zero diffusion terms for all passenger outflows: For every $t \in T_3$, we set up a balance equation in form of (5) and integrate it into (6). This balance equation relates the passenger load of ${}^{\star}t$, to the outflow assigned to t (cf. Tab. 2).

3.3 Correctness of Our Decoupling Approach

Assume that our SHA model is in mode $q \in Q$ at time $\tau = i\,{}_{\Delta}\tau$, for some $i \in \mathbb{N}_{\geq 0}$; in which ${}_{\Delta}\tau > 0$ is the fixed time step that separates every pair of two consecutive time layers confining all mode transitions. Moreover, assume that we like to compute the probability of a particular mode transition of the SHA model at time $\tau = (i+1)\ {}_{\Delta}\tau$; which is triggered by the passenger load trajectory of some GP in a station or vehicle docked to a station $b \in P \cup \mathcal{V}'(q)$ taking a value from $k \in K$, with $K \subseteq \mathcal{M}(b)$ and $\mathcal{M}(b)$ from (1). More formally speaking, we thus like to compute the probability

$$\mathbb{P}\left(\mathrm{M}\left(b,(i+1)\ {}_{\Delta}\tau\right) \in K\right) := \int_K \mathrm{pdf}\left(\mathrm{M}\left(b,(i+1)\ {}_{\Delta}\tau\right) = k\right) \mathrm{d}k \tag{10}$$

with $\mathrm{M}(b,\tau)$ specified at $\tau = i\,{}_{\Delta}\tau$ by $\mathrm{pdf}\left(\mathrm{M}\left(b, i\,{}_{\Delta}\tau\right)\right)$ acc. to (9).

Remark 11. Let $X_1,\ X_2,\ \ldots,\ X_n$ be a vector of $n \in \mathbb{N}_{>0}$ continuous RVs. Then, $\mathrm{pdf}\left(X_j; j \in \{1,2,\ldots,n\}\right)$ denotes the joint PDF of $X_1,\ X_2,\ \ldots,\ X_n$; $\mathrm{pdf}\left(X_j = x_j; j \in \{1,2,\ldots,n\}\right)$ denotes the evaluation of $\mathrm{pdf}\left(X_j; j \in \{1,2,\ldots,n\}\right)$ at $(x_1, x_2, \ldots, x_n)$, with $x_j \in \sigma(X_j)$, $\forall j \in \{1,2,\ldots,n\}$.

Look at

$$\begin{aligned}\mathbb{P}\left(\sum_{t \in \left({}^{\bullet}b \cup b^{\bullet}\right) \cap T'(q)} \mathrm{M}_{q,t}\left(b,(i+1)\ {}_{\Delta}\tau\right) \in K\right) = \\ \int_K \mathrm{pdf}\left(\sum_{t \in \left({}^{\bullet}b \cup b^{\bullet}\right) \cap T'(q)} \mathrm{M}_{q,t}\left(b,(i+1)\ {}_{\Delta}\tau\right) = k\right) \mathrm{d}k\end{aligned} \tag{11}$$

instead, which is the probability that the sum of the decoupled passenger loads of the different isolated slices from b (isolated in q) takes a value from K at $\tau = (i+1)\ {}_{\Delta}\tau$. Let

$$l := \left|\left({}^{\bullet}b \cup b^{\bullet}\right) \cap T'(q)\right|, \tag{12}$$

and introduce the set $\overline{\mathcal{M}}(b,k)$, with

$$\overline{\mathcal{M}}(b,k) := \left\{(k_1, k_2, \ldots, k_l) \in \left(\mathcal{M}(b)\right)^l : \sum_{j=1}^{l} k_j = k\right\} \tag{13}$$

Moreover, let $\{t_1, t_2, \ldots, t_l\} := \left({}^{\bullet}b \cup b^{\bullet}\right) \cap T'(q)$. Then, write down (11) in form of

$$\begin{aligned}\mathbb{P}\left(\sum_{t \in \left({}^{\bullet}b \cup b^{\bullet}\right) \cap T'(q)} \mathrm{M}_{q,t}\left(b,(i+1)\ {}_{\Delta}\tau\right) \in K\right) = \\ \int_K \int_{\overline{\mathcal{M}}(b,k)} \mathrm{pdf}\left(\mathrm{M}_{q,t_j}\left(b,(i+1)\ \tau\right) = k_j; j \in \{1,2,\ldots,l\}\right) \mathrm{d}\left(k_1, k_2, \ldots, k_l\right) \mathrm{d}k\end{aligned} \tag{14}$$

Therein, note that $\mathrm{M}_{q,t_1}\left(b,(i+1)\ {}_{\Delta}\tau\right),\ldots,\mathrm{M}_{q,t_l}\left(b,(i+1)\ {}_{\Delta}\tau\right)$ are independent RVs. Thus, (14) simplifies to

$$\begin{aligned}&\mathbb{P}\left(\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\mathrm{M}_t\left(b,(i+1)\ {}_{\Delta}\tau\right)\in K\right)=\\&\int_K\int_{\overline{\mathcal{M}}(b,k)}\prod_{t_j\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\mathrm{pdf}\left(\mathrm{M}_{t_j}\left(b,(i+1)\ \tau\right)=k_j\right)\,\mathrm{d}\left(k_1,k_2,\ldots,k_l\right)\,\mathrm{d}k\end{aligned}\tag{15}$$

Theorem 1. *For any $q\in Q$, $b\in P\cup\mathcal{V}'(q)$, and $k\in\mathcal{M}(b)$, the integral*

$$\int_{\overline{\mathcal{M}}(b,k)}\prod_{t_j\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}pdf(\mathrm{M}_{t_i}\left(b,(i+1)\ \tau\right)=k_i)\,\mathrm{d}\left(k_1,k_2,\ldots,k_l\right)$$

from (15) converges to pdf$(\mathrm{M}\left(b,(i+1)\ {}_{\Delta}\tau\right)=k)$ from (10) for ${}_{\Delta}\tau\xrightarrow{{}_{\Delta}\tau>0}0$.

Note that Thm. 1 implies that our above decoupling approach produces a set of SDEs (one for every decoupled flow) for the different modes of our SHA model; this set approximates the original coupled passenger flow dynamics in the limiting case of vanishing discrete simulation time steps, when we let the decoupled slices communicate their results.

Proof of Theorem 1. Common Initial State: From (7), note that

$$\begin{aligned}\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\mathrm{M}_{q,t}\left(b,i\ {}_{\Delta}\tau\right)&=\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\lambda\left(b,t,q\right)\mathrm{M}\left(b,i\ {}_{\Delta}\tau\right)\\&=\mathrm{M}\left(b,i\ {}_{\Delta}\tau\right)\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\lambda\left(b,t,q\right).\end{aligned}\tag{16}$$

From (9) follows

$$\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\lambda\left(b,t,q\right)=1,\tag{17}$$

which in turn implies

$$\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\mathrm{M}_{q,t}\left(b,i\ {}_{\Delta}\tau\right)=\mathrm{M}\left(b,i\ {}_{\Delta}\tau\right).\tag{18}$$

Common Differential Dynamics: The continuous time evolution of

$$\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\mathrm{M}_{q,t}\left(b,\tau\right)$$

in the time interval $\tau\in\left[i\ {}_{\Delta}\tau,(i+1)\ {}_{\Delta}\tau\right)$ is defined by

$$\mathrm{d}\left(\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\mathrm{M}_{q,t}\left(b,\tau\right)\right)=\sum_{t\in({}^{\bullet}b\cup b^{\bullet})\cap T'(q)}\mathrm{dM}_{q,t}\left(b,\tau\right),\tag{19}$$

with initial state

$$\mathrm{M}_{q,t}\left(b, i\,{}_{\Delta}\tau\right),$$

for some $i \in \mathbb{N}_{\geq 0}$ and $\tau_{\Delta}\tau > 0$, which is identical to (5) for ${}_{\Delta}\tau \to 0$ given the specification of (6) acc. to Tab. 1 and Tab. 2, q.e.d.

3.4 Consequence of Our Decoupling Approach

In the original approximate computation of our SHA model's state space, we were confronted with one system of coupled SDEs for every station $s \in \mathcal{S}$ in every mode. The dimension of this system is $n := (n_{s,1} + n_{s,2})\, n_{\mathrm{y}}$ iff $n_{s,1}$ corresponds to the number of different gathering points in s, $n_{s,2}$ corresponds to the number of vehicles docked to s, and $n_{\mathrm{y}} := |\mathcal{Y}|$ corresponds to the number of the passengers' different trip profiles in the TN at hand. Now our decoupling approach replaces this n-dimensional system of SDEs by a set of probably much smaller systems of ODEs (with uncertain initial states) and SDEs: Every of this new/replacing system of equations has $2\, n_{\mathrm{y}}$ dimensions if it captures a transfer flow, and n_{y} dimensions otherwise.

4 Summary & Outlook

In this paper, we have considered one major bottleneck that may arise in the approximate computation of our SHA model from [5]: the numerical computation of the many high-dimensional SDEs, which define the passenger flow dynamics in its different modes. More specifically, we have shown how all passenger flows can be systematically decoupled in the different modes of our SHA model, which produces a set of lower-dimensional ODEs and SDEs replacing the original SDEs. We proved correctness of this decoupling approach. Numerical experiments are under way. We want to share our insights obtained from them in future publications, where we also intend to (i) discuss improvements targeting the computation of the SHA model's discrete state, and (ii) show how our model and algorithms for its approximate computation can be applied to the perturbation analysis of a multimodal TN.

Acknowledgement. The final publication will be available at Springer.

References

1. Brooks, S.e.a.: Handbook of Markov Chain Monte Carlo (Chapman & Hall/CRC Handbooks of Modern Statistical Methods). Chapman and Hall/CRC (2011)
2. Causon, D.M., Mingham, C.G.: Introductory Finite Volume Methods for PDEs. Ventus Publishing ApS (2011)
3. Ciardo, G., Nicol, D., Trivedi, K.S.: Discrete-event simulation of fluid stochastic Petri nets. IEEE Transactions on Software Engineering 25, 207 – 217 (1997)

4. Haar, S., Theissing, S.: A hybrid-dynamical model for passenger-flow in transportation systems. In: 5th IFAC Conference on Analysis and Design of Hybrid Systems (2015)
5. Haar, S., Theissing, S.: Forecasting Passenger Loads in Transportation Networks (2016), https://hal.inria.fr/hal-01259585, working paper
6. Haar, S., Theissing, S.: Predicting traffic load in public transportation networks (2016), https://hal.archives-ouvertes.fr/hal-01286476, working paper
7. Hoogerheide, L., Kaashoek, J., van Dijk, H.: Functional approximations to posterior densities: a neural network approach to efficient sampling (Dec 2002), http://hdl.handle.net/1765/1727
8. Hu, J., Lygeros, J., Sastry, S.: Towards a theory of stochastic hybrid systems. In: Hybrid Systems: Computation and Control. Springer Berlin Heidelberg (2000)
9. MacKay, D.J.C.: Introduction to monte carlo methods. In: Proceedings of the NATO Advanced Study Institute on Learning in Graphical Models (1998)
10. Wolter, K.: Modelling hybrid systems with fluid stochastic petri nets. In: Proceedings of the 4th International Conference on Automation of Mixed Processes: Hybrid Dynamic Systems (2000)