Compositional Modeling and Minimization of Time-Inhomogeneous Markov Chains

Tingting Han[1,2], Joost-Pieter Katoen[1,2], and Alexandru Mereacre[1,3]

[1] RWTH Aachen University, Software Modeling and Verification Group, Germany
{tingting.han,katoen,mereacre}@cs.rwth-aachen.de
[2] University of Twente, Formal Methods and Tools Group, The Netherlands
[3] University of Trento, Dept. of Information and Communication Technology, Italy

Abstract. This paper presents a compositional framework for the modeling of interactive continuous-time Markov chains with time-dependent rates, a subclass of communicating piecewise deterministic Markov processes. A poly-time algorithm is presented for computing the coarsest quotient under strong bisimulation for rate functions that are either piecewise uniform or (piecewise) polynomial. Strong as well as weak bisimulation are shown to be congruence relations for the compositional framework, thus allowing component-wise minimization. In addition, a new characterization of transient probabilities in time-inhomogeneous Markov chains with piecewise uniform rates is provided.

1 Introduction

Modeling large stochastic discrete-event dynamic systems is a difficult task that typically requires human intelligence and ingenuity. To facilitate this process, formalisms are needed that allow for the modeling of such systems in a compositional manner. This allows to construct models of simpler components—usually from first principles—that can be combined by appropriate composition operators to yield complete system models. In concurrency theory, process algebra [20, 16] has emerged as an important framework to achieve compositionality: it provides a formal apparatus for compositional reasoning about structure and behavior of systems, and features abstraction mechanisms allowing system components to be treated as black boxes.

Although originally aimed at purely functional behavior, process algebras for stochastic systems have been investigated thoroughly, see e.g., [15, 14]. In all these approaches, the dynamics of the stochastic models is assumed to be *time-homogeneous*, i.e., the probabilistic nature of mode transitions as well as the time-driven behavior are independent of the global time. This is, however, a serious drawback to adequately model random phenomena that occur in practice such as failure rates of hardware components (a bath-tub curve), software reliability (which reduces due to memory leaks and increases after a restart), and battery depletion (where the power extraction rate non-linearly depends on the remaining amount of energy [5]), to mention a few. This paper attempts to overcome this deficiency by providing a process algebra for time-*inhomogeneous* continuous-time Markov chains (ICTMCs). This is a very versatile class of models and is a natural stepping-stone towards more full-fledged stochastic hybrid system models such as piecewise deterministic Markov processes (PDPs [6]). We show that ICTMCs can be compositionally modeled by using a time-dependent adaptation of the framework of interactive Markov chains (IMCs) [14]. To facilitate this, ICTMCs are equipped with the potential for interaction, i.e., synchronization. Instrumental to this approach is the memoryless property of ICTMCs.

More importantly though, notions of strong and weak bisimulation are defined and shown to be congruences. Together with efficient quotienting algorithms this allows for the component-wise minimization of hierarchical ICTMC models. Finally, we present an axiomatization for strong and weak bisimulation which allows to simplify models by pure syntactic manipulations as opposed to performing minimization on the model level. As a generalization of results on ordinary lumpability on Markov chains [3], we show that strong bisimulation preserves transient and long-run state probabilities in ICTMCs. This allows to minimize symbolically ICTMCs prior to their analysis.

We present a bisimulation minimization algorithm to obtain the *coarsest* (and thus smallest) *strong bisimulation quotient* of a large class of interactive ICTMCs, viz. those that have piecewise uniform—rate $\mathbf{R}_k(t)$ on piece k is of the form $f_k(t)\cdot\mathbf{R}$ for integrable function $\mathbf{R}$—polynomial, or piecewise polynomial—where each polynomial is of degree three—rate functions. The worst-case time and space complexity is $\mathcal{O}\left(m_a \lg(n) + M m_r \lg(n)\right)$ and $\mathcal{O}\left(m_a + m_r\right)$, respectively, where $M{+}1$ is the number of pieces (or degrees of the polynomial), m_a is the number of action-labeled transitions and m_r the number of rate-labeled transitions in the ICTMC under consideration. This algorithm is based on the partition-refinement bisimulation algorithm for Markov chains by Derisavi *et al.* [7] and Paige-Tarjan's algorithm for labeled transition systems (LTS) [21].

Related work. ICTMCs are related to piecewise deterministic Markov processes (PDPs), a more general class of continuous-time stochastic discrete-event dynamic systems proposed by Davis [6]. The probabilistic nature of mode transitions in PDPs is as for ICTMCs; in fact, ICTMCs are a subclass of PDPs when the global time t has a clock dynamics i.e., $\dot{t} = 1$. The notion of parallel composition of ICTMCs corresponds to that for communicating PDPs (CPDPs) as introduced by Strubbe and van der Schaft [24, 23]. Alternative modeling formalisms for PDPs are, e.g., variants of colored Petri nets [9] but they lack a clear notion of compositionality. Compositional modeling formalisms for hybrid systems have been considered by, e.g., [2, 1]. Strong bisimulation has been proposed for several classes of (stochastic) hybrid systems, see e.g., [4, 12, 25]. Our notion of bisimulation is closely related to that for CPDPs [25] but differs in the fact that the maximal progress assumption—a race between one or more rates and a transition that is not subject to interaction with the environment is resolved in favor of the internal transition—is not considered in [25]. Proofs of the major results are contained in [13].

2 Inhomogeneous Continuous Time Markov Chains

Definition 1 (ICTMC). *An inhomogeneous continuous-time Markov chain is a tuple* $\mathcal{C} = (\mathbb{S}, \mathbf{R})$ *where:* $\mathbb{S} = \{1, 2, \ldots, n\}$ *is a finite set of states, and* $\mathbf{R}(t) = [R_{i,j}(t) \geq$

$0] \in \mathbb{R}^{n \times n}$ *is a time-dependent rate matrix, where* $R_{i,j}(t)$ *is the rate between states* $i, j \in \mathbb{S}$ *at time* $t \in \mathbb{R}_{\geq 0}$.

Let diagonal matrix $\mathbf{E}(t) = \mathrm{diag}\,[E_i(t)] \in \mathbb{R}^{n \times n}$, where $E_i(t) = \sum_{j \in \mathbb{S}} R_{i,j}(t)$ for all $i, j \in \mathbb{S}$, $i \neq j$ i.e., $E_i(t)$ is the total exit rate of state i at time t. Consider a non-homogeneous Poisson process $\{Z(t) | t \geq 0\}$ with rate $R(t)$. The probability of k arrivals in the interval $[t, t + \Delta t]$ is:

$$\Pr\{Z(t + \Delta t) - Z(t) = k\} = \frac{\left[\int_t^{t+\Delta t} R(\ell) d\ell\right]^k}{k!} e^{-\int_t^{t+\Delta t} R(\ell) d\ell}, \quad k = 0, 1, \ldots.$$

The probability that there will be no arrivals in the interval $[t, t + \Delta t]$ is:

$$\Pr\{Z(t + \Delta t) - Z(t) = 0\} = e^{-\int_t^{t+\Delta t} R(\ell) d\ell} = e^{-\int_0^{\Delta t} R(t+\ell) d\ell}. \tag{1}$$

Let the random variable $W_{i,j}(t)$ be the firing time of transition $i \to j$ ($i, j \in \mathbb{S}$) with rate $R_{i,j}(t)$ at time t. From (1) we obtain the cumulative probability distribution of the firing time of transition $i \to j$:

$$\Pr\{W_{i,j}(t) \leq \Delta t\} = 1 - \Pr\{Z(t + \Delta t) - Z(t) = 0\} = 1 - e^{-\int_0^{\Delta t} R_{i,j}(t+\ell) d\ell}. \tag{2}$$

Probability measures. For every ICTMC one can specify measures of interest. These measures are either related to the states or to the transitions of an ICTMC. Consider a random variable $W_i(t)$ which denotes the waiting time in state i.

Property 1.

$$\Pr\{W_i(t) \leq \Delta t\} = 1 - e^{-\int_0^{\Delta t} E_i(t+\ell) d\ell}. \tag{3}$$

An intuitive explanation of (3) is that the waiting time $W_i(t)$ in state i is determined by the minimal firing time of all k outgoing transitions from state i, i.e., $W_i(t) = \min\{W_{i,1}(t), \ldots, W_{i,k}(t)\}$. When $R_{i,j}(t) = R_{i,j}$ and $E_i(t) = E_i$ for all $t \in \mathbb{R}_{\geq 0}$, i.e., the ICTMC is a CTMC, $W_i(t)$ has the distribution $1 - e^{-E_i \Delta t}$. An interesting property is that the waiting time in any state i is *memoryless*, i.e.:

$$\Pr\{W_i(t) \leq t' + \Delta t | W_i(t) > t'\} = \Pr\{W_i(t + t') \leq \Delta t\}. \tag{4}$$

This can be shown as follows:

$$\begin{aligned}
\Pr\{W_i(t) \leq t' + \Delta t | W_i(t) > t'\} &= \frac{e^{-\int_0^{t'} E_i(t+\ell) d\ell} - e^{-\int_0^{t'+\Delta t} E_i(t+\ell) d\ell}}{e^{-\int_0^{t'} E_i(t+\ell) d\ell}} \\
= 1 - e^{-\int_0^{t'+\Delta t} E_i(t+\ell) d\ell + \int_0^{t'} E_i(t+\ell) d\ell} &= \Pr\{W_i(t + t') \leq \Delta t\}.
\end{aligned}$$

Equation (4) will be of importance when we later define a calculus for ICTMCs.

Property 2. The probability $\Pr_{i,j}(t)$ to select transition $i \to j$ ($i \neq j, i, j \in \mathbb{S}$) with rate $R_{i,j}(t)$ at time t is:

$$\Pr_{i,j}(t) = \int_0^\infty R_{i,j}(t + \tau) e^{-\int_0^\tau E_i(t+\ell) d\ell} d\tau. \tag{5}$$

When rates are constant, the measure (5) takes the form $\Pr_{i,j} = \frac{R_{i,j}}{E_i}$ ($\Pr_{i,j}(t) = \Pr_{i,j}$ for all $t \in \mathbb{R}_{\geq 0}$), which corresponds to transition probability in CTMCs.

Property 3. The cumulative probability distribution $\Pr_{i,j}(t, \Delta t)$ to move from state i to state j ($i \neq j$) with rate $R_{i,j}(t)$ in Δt time units starting at time t:

$$\Pr_{i,j}(t, \Delta t) = \int_0^{\Delta t} R_{i,j}(t+\tau) e^{-\int_0^{\tau} E_i(t+\ell) d\ell} d\tau. \tag{6}$$

Notice that (6) is the same as (5) except that the range of the outer-most integral is $[0, \Delta t]$. For CTMCs ($\Pr_{i,j}(t, \Delta t) = \Pr_{i,j}(\Delta t)$ for all $t \in \mathbb{R}_{\geq 0}$), equation (6) results in $\Pr_{i,j}(\Delta t) = \frac{R_{i,j}}{E_i}\left(1 - e^{-E_i \Delta t}\right)$.

Transient probability distribution. One important measure which quantifies the probability to be in a specific state at some time point is the *transient probability distribution*. Consider an ICTMC described by the stochastic process $\{X(t) | t \geq 0\}$. The transient probability distribution $\Pr\{X(t+\Delta t) = j\}$, denoted by $\pi_j(t+\Delta t)$, is the probability to be in state j at time $t + \Delta t$, and is described by the equation:

$$\pi_j(t+\Delta t) = \sum_{i \in \mathbb{S}} \Pr\{X(t) = i\} \cdot \Pr\{X(t+\Delta t) = j | X(t) = i\}. \tag{7}$$

Equation (7) can be expressed in matrix form as: $\boldsymbol{\pi}(t+\Delta t) = \boldsymbol{\pi}(t)\boldsymbol{\Phi}(t+\Delta t, t)$, where $\boldsymbol{\pi}(t) = [\pi_1(t), \ldots, \pi_n(t)]$ and $\boldsymbol{\Phi}(t+\Delta t, t)$ represents the *transition probability matrix*. This equation represents the solution of a system of ODEs:

$$\frac{d\boldsymbol{\pi}(t)}{dt} = \lim_{\Delta t \to 0} \frac{\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)}{\Delta t} = \lim_{\Delta t \to 0} \boldsymbol{\pi}(t) \frac{[\boldsymbol{\Phi}(t+\Delta t, t) - \mathbf{I}]}{\Delta t}. \tag{8}$$

For the diagonal elements $q_{i,i}(t)$ of the matrix $\lim_{\Delta t \to 0} \frac{[\boldsymbol{\Phi}(t+\Delta t,t) - \mathbf{I}]}{\Delta t}$ from (8), we obtain $q_{i,i}(t) = \lim_{\Delta t \to 0} \frac{\Pr\{X(t+\Delta t)=i | X(t)=i\} - 1}{\Delta t}$. As $\Pr\{X(t+\Delta t) = i | X(t) = i\}$ denotes the probability to stay in state i for at least Δt units of time or the probability to return to state i in two or more steps, it follows:

$$q_{i,i}(t) = \lim_{\Delta t \to 0} \frac{e^{-\int_0^{\Delta t} E_i(t+\ell) d\ell} - 1 + o(\Delta t)}{\Delta t} = -E_i(t),$$

where $o(\Delta t)$ denotes the probability to make two or more transitions in Δt units of time. Notice that $\lim_{\Delta t \to 0} \frac{o(\Delta t)}{\Delta t} = 0$. For the off-diagonal elements $q_{i,j}(t)$ ($i \neq j$) of matrix $\lim_{\Delta t \to 0} \frac{[\boldsymbol{\Phi}(t+\Delta t,t) - \mathbf{I}]}{\Delta t}$, the relation is similar:

$$q_{i,j}(t) = \lim_{\Delta t \to 0} \frac{\Pr\{X(t+\Delta t) = j | X(t) = i\}}{\Delta t} = \lim_{\Delta t \to 0} \frac{\Pr_{i,j}(t, \Delta t) + o(\Delta t)}{\Delta t},$$

which can be reduced using (6) to:

$$q_{i,j}(t) = \lim_{\Delta t \to 0} \frac{\int_0^{\Delta t} R_{i,j}(t+\tau) e^{-\int_0^{\tau} E_i(t+\ell) d\ell} d\tau + o(\Delta t)}{\Delta t} = R_{i,j}(t).$$

The resulting *infinitesimal generator* matrix $\mathbf{Q}(t)$ has the form:

$$\mathbf{Q}(t) = \lim_{\Delta t \to 0} \frac{[\mathbf{\Phi}(t+\Delta t, t) - \mathbf{I}]}{\Delta t} = \mathbf{R'}(t) - \mathbf{E}(t),$$

where $\mathbf{R'}$ equals $\mathbf{R}$ except that $R^{'}_{i,i}(t) = 0$. Plugging $\mathbf{Q}(t)$ into equation (8) yields the system of ODEs which describe the evolution of transient probability distribution over time (Chapman-Kolmogorov equations):

$$\frac{d\boldsymbol{\pi}(t)}{dt} = \boldsymbol{\pi}(t)\mathbf{Q}(t), \qquad \sum_{i=1}^{n} \pi_i(t_0) = 1, \tag{9}$$

where $\boldsymbol{\pi}(t_0)$ is the initial condition. From the literature (see [17, pages 594–631]) it is known that the solution $\boldsymbol{\pi}(t)$ of (9), written as:

$$\boldsymbol{\pi}(t) = \boldsymbol{\pi}(t_0)\mathbf{\Phi}(t, t_0) \tag{10}$$

has the transition probability matrix given by the Peano-Baker series:

$$\mathbf{\Phi}(t, t_0) = \mathbf{I} + \int_{t_0}^{t} \mathbf{Q}(\tau_1) d\tau_1 + \int_{t_0}^{t} \mathbf{Q}(\tau_1) \int_{t_0}^{\tau_1} \mathbf{Q}(\tau_2) d\tau_2 d\tau_1 + \ldots. \tag{11}$$

Note that if $\mathbf{Q}(\tau_1) \int_{t_0}^{\tau_1} \mathbf{Q}(\tau_2) d\tau_2 = \int_{t_0}^{\tau_1} \mathbf{Q}(\tau_2) d\tau_2 \mathbf{Q}(\tau_1)$ then $\mathbf{\Phi}(t, t_0) = e^{\int_{t_0}^{t} \mathbf{Q}(\tau) d\tau}$. If the rate matrix $\mathbf{R}(t)$ is *piecewise constant* i.e., $\mathbf{R}(t) = \mathbf{R}_k$ or $\mathbf{Q}(t) = \mathbf{Q}_k$ for all $t \in [t_k, t_{k+1})$ and $k \leq M \in \mathbb{N}$ ($M+1$ is the total number of constant pieces), equation (10) can also be rewritten as (see [22]):

$$\boldsymbol{\pi}(t) = \begin{cases} \boldsymbol{\pi}(t_0) e^{\mathbf{Q}_0 (t - t_0)} & \text{if } t \in [t_0, t_1) \\ \vdots & \vdots \\ \boldsymbol{\pi}(t_M) e^{\mathbf{Q}_M (t - t_M)} & \text{if } t \in [t_M, \infty) \end{cases} \quad \text{and } \boldsymbol{\pi}(t_k) = \boldsymbol{\pi}(t_{k-1}) e^{\mathbf{Q}_{k-1}(t_k - t_{k-1})}.$$

The general case is when the rate matrix is *piecewise uniform* i.e., $\mathbf{R}(t) = \mathbf{R}_k(t) = f_k(t)\mathbf{R}_k$ or $\mathbf{Q}(t) = \mathbf{Q}_k(t) = f_k(t)\mathbf{Q}_k$ for any integrable function $f_k(t) : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ on time interval $[t_k, t_{k+1})$, constant matrices $\mathbf{R}_k$ and $\mathbf{Q}_k$.

Theorem 1. *The transient probability distribution* $\boldsymbol{\pi}(t)$ *of an* $\mathrm{ICTMC}\ \mathcal{C} = (\mathbb{S}, \mathbf{R})$ *with a piecewise uniform rate matrix* $\mathbf{R}(t)$ *and* $M+1$ *pieces is given by:*

$$\boldsymbol{\pi}(t) = \begin{cases} \boldsymbol{\pi}(t_0) e^{\mathbf{Q}_0 \int_{t_0}^{t} f_0(\tau) d\tau} & \textit{if } t \in [t_0, t_1) \\ \vdots & \vdots \\ \boldsymbol{\pi}(t_M) e^{\mathbf{Q}_M \int_{t_M}^{t} f_M(\tau) d\tau} & \textit{if } t \in [t_M, \infty) \end{cases}$$

where $\boldsymbol{\pi}(t_k) = \boldsymbol{\pi}(t_{k-1}) e^{\mathbf{Q}_{k-1} \int_{t_{k-1}}^{t_k} f_{k-1}(\tau) d\tau}$.

3 Inhomogeneous Interactive Markov Chains

In order to facilitate the compositional modeling of ICTMCs, we equip these processes with the capability to allow for their mutual interaction. This is established by adding actions to ICTMCs. Let Act be the countable universe of actions. The aim of these actions is that certain actions can only be performed together with other processes.

Definition 2 (I^2MC). *An inhomogeneous interactive Markov chain* $(\mathrm{I}^2\mathrm{MC})$ *is a tuple* $\mathcal{I} = (\mathbb{S}, Act, \rightarrow, \mathbf{R}, s^0)$ *where* $\mathbb{S}$ *and* $\mathbf{R}$ *are as before,* $\rightarrow \subseteq \mathbb{S} \times Act \times \mathbb{S}$ *is a transition relation and* $s^0 \in \mathbb{S}$ *is the initial state.*

The semantic model of I^2MC represents the time-dependent variant of IMC [14].

Process algebra for $\mathrm{I}^2\mathrm{MC}$. Originally developed by Hoare and Milner (see [20, 16]), process algebras have been developed as a compositional framework for describing the functional behavior of the system. It allows for modeling complex systems in a component-wise manner by offering a set of *operators* to combine component models. Actions are the most elementary notions. The combination of several actions using the operators forms a *process*. We extend this framework by stochastic timing facilities.

Definition 3. *Let* X *be a process variable,* $\lambda(t) \in \mathbb{R}_{\geq 0}$ *with* $t \in \mathbb{R}_{\geq 0}$, $A \subseteq Act$ *and* $a \in Act$. *The syntax of inhomogeneous interactive Markov language* $(\mathrm{I}^2\mathrm{ML})$ *for* $\mathrm{I}^2\mathrm{MCs}$ *is defined as follows:*

$$P ::= 0 \quad | \quad a.P \quad | \quad \lambda(t).P \quad | \quad P + P \quad | \quad P \|_A P \quad | \quad P \setminus A \quad | \quad X.$$

Process variables are assumed to be defined by recursive equations of the form $X := P$, where P is an I^2ML term. The *null process* 0 is the deadlock process and cannot perform any action. The prefix operators are $a.P$ and $\lambda(t).P$ for actions and rates, respectively. The *choice* operator $P + Q$ chooses between processes P or Q. Process $P\|_A Q$ denotes the *parallel composition* of processes P and Q where synchronization is required only for actions in A; actions not in A are performed autonomously. The process $P \setminus A$ behaves like P except that all actions in A become unobservable to other processes; this is established by relabeling a by the distinguished action $\tau \in Act$. The operational semantics of I^2ML terms is defined by the inference rules in Table 1 where for the sake of conciseness symmetric rules are not shown.

A few remarks concerning time-prefix and choice are in order. The process $\lambda(t).P$ evolves into P within Δt time units with probability:

$$\mathrm{Pr}_{\lambda(t).P,P}(t, \Delta t) = \int_0^{\Delta t} \lambda(t+\tau) e^{-\int_0^{\tau} \lambda(t+\ell)d\ell} d\tau = 1 - e^{-\int_0^{\Delta t} \lambda(t+\ell)d\ell},$$

given that $\lambda(t).P$ is enabled at the global time t. The above relation can be easily proven from (6) by taking $i = \lambda(t).P$, $j = P$, $R_{i,j}(t+\tau) = \lambda(t+\tau)$ and $E_s(t+\ell) = \lambda(t+\ell)$. The process $\lambda(t).P + \mu(t).Q$ can evolve into P if the time delay generated by a stochastic process with rate $\lambda(t)$ is smaller than that generated by a different stochastic process with rate $\mu(t)$. By a symmetric argument it may evolve into Q. Therefore, from (3) it follows that the distribution of time until a choice is made is

$\dfrac{}{a.P \xrightarrow{a} P}$	$\dfrac{P \xrightarrow{a} P' \text{ and } Q \xrightarrow{a} Q'}{P\|_A Q \xrightarrow{a} P'\|_A Q'} (a \in A)$	$\dfrac{}{\lambda(t).P \xrightarrow{\lambda(t)} P}$	$\dfrac{P \xrightarrow{\lambda(t)} P'}{P\backslash A \xrightarrow{\lambda(t)} P'\backslash A}$
$\dfrac{P \xrightarrow{a} P'}{P{+}Q \xrightarrow{a} P'}$	$\dfrac{P \xrightarrow{a} P'}{P\|_A Q \xrightarrow{a} P'\|_A Q} (a \notin A)$	$\dfrac{P \xrightarrow{\lambda(t)} P'}{P{+}Q \xrightarrow{\lambda(t)} P'}$	$\dfrac{E[X{:=}E/X] \xrightarrow{\lambda(t)} E'}{X{:=}E \xrightarrow{\lambda(t)} E'}$
$\dfrac{P \xrightarrow{a} P'}{P\backslash A \xrightarrow{a} P'\backslash A} (a \notin A)$	$\dfrac{P \xrightarrow{a} P'}{P\backslash A \xrightarrow{\tau} P'\backslash A} (a \in A)$	$\dfrac{P \xrightarrow{\lambda(t)} P'}{P\|_A Q \xrightarrow{\lambda(t)} P'\|_A Q}$	$\dfrac{E[X{:=}E/X] \xrightarrow{a} E'}{X{:=}E \xrightarrow{a} E'}$

Table 1. Inference rules for the operational semantics of I^{2}ML.

$\Pr\{W(t) \leq \Delta t\} = 1 - e^{-\int_0^{\Delta t} \lambda(t+\tau)+\mu(t+\tau) d\tau}$. For a choice between $|J|$ processes (J is a finite index set), the distribution of the waiting time becomes $\Pr\{W(t) \leq \Delta t\} = 1 - e^{-\int_0^{\Delta t} \sum_{i\in J} \lambda_i(t+\tau) d\tau}$. If the rates $\lambda_i(t)$ in the process $\sum_{i\in J} \lambda_i(t).P_i$ are constant ($\lambda_i(t) = \lambda_i$), then the waiting time is exponentially distributed with the sum of the rates λ_i i.e. $\Pr\{W(t) \leq \Delta t\} = 1 - e^{-\sum_{i\in J} \lambda_i \Delta t}$. This corresponds to the interpretation of choice in Markovian process algebras [15]. It is important to note that when $P_i = P$ for all $i \in J$, the process $\sum_{i\in J} \lambda_i(t).P$ will evolve into P with rate $\sum_{i\in J} \lambda_i(t)$.

Parallel composition. When considering just actions the asynchronous parallel composition has the same functionality as that from basic process calculi. On the other hand when considering stochastic delays the composition is more involved. Consider $P := \lambda(t).P'$ and $Q := \mu(t).Q'$. They can evolve into P' and Q' after a time delay governed by a distribution with rate $\lambda(t)$ and $\mu(t)$, respectively. Since the waiting time in any state is memoryless (4), we can show the way by which processes P and Q are composed (see diagram below).

$$\begin{tikzcd}
 & \boxed{P\|Q} \arrow[dl, "\lambda(t)"'] \arrow[dr, "\mu(t)"] & \\
\boxed{P'\|Q} \arrow[dr, "\mu(t)"'] & & \boxed{P\|Q'} \arrow[dl, "\lambda(t)"] \\
 & \boxed{P'\|Q'} &
\end{tikzcd}$$

First consider that when both processes start their execution in initial state $P\|Q$ (the shadowed state) they probabilistically select a time delay, say, Δt_λ for P and Δt_μ for Q. If $\Delta t_\lambda < \Delta t_\mu$ then P finishes its execution first and evolves into P'. The same applies to Q when $\Delta t_\mu < \Delta t_\lambda$. By intuition we could think that when it is already in $P'\|Q$, $\Delta t_\lambda = 0$ and the remaining delay for process Q until it finishes its execution is $\Delta t_\mu - \Delta t_\lambda$. What really happens is that on entering state $P'\|Q$ both delays are set to zero i.e., $\Delta t_\lambda = \Delta t_\mu = 0$. As P' has no transitions, Δt_λ remains 0 but for Q its delay is initialized to a new value which might be different from $\Delta t_\mu - \Delta t_\lambda$ due to a probabilistic selection. Due to the memoryless property, however, the remaining delay for Q is fully determined by μ only.

Example 1. Consider two hardware components described by the equations $P := \lambda_1(t).0{+}\lambda_2(t).use.P$ and $Q := \mu_1(t).0{+}\mu_2(t).use.Q$, respectively. Each of the components may fail with rate $\lambda_1(t)$ and $\mu_1(t)$, respectively. As a result of the failure they evolve into process 0. On the other hand, the components may move to the working state with the rate $\lambda_2(t)$ and $\mu_2(t)$, respectively, where they can *use* some resources. If one of them fails then the entire system fails. Both components can use the resources at the same time if the system is working properly. Figure 1 depicts the I^{2}MC of $P\|_{\{use\}}Q$.

4 Strong and Weak Bisimulation

In order to compare the behavior of ICTMCs (and their interactive variants) we exploit the well-studied and widely accepted notion of bisimulation [3, 20, 14]. A classical bisimulation relation requires equivalent states to be able to mutually mimic their stepwise behavior. In the probabilistic setting this is interpreted as requiring equivalent states to have equal cumulative rates to move to any equivalence class. Bisimulation is considered as a natural notion of equivalent behavior, is equipped with quotienting algorithms, and has a clear correspondence to equivalence in terms of logical behavioral specifications. In this section, we will define strong bisimulation for I^{2}MC starting from a similar notion on ICTMCs. Some algebraic and probabilistic properties of bisimulation are investigated. The same applies to *weak* bisimulation that allows for the abstraction of internal, i.e., τ actions.

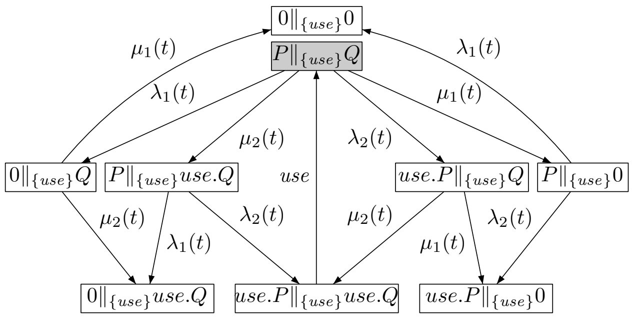

Fig. 1. $P\|_{\{use\}}Q$.

Bisimulation for ICTMCs.

Definition 4 (ICTMC strong bisimulation). *An equivalence $\mathcal{R} \subseteq \mathbb{S} \times \mathbb{S}$ is a strong bisimulation whenever for all $(P, Q) \in \mathcal{R}$, $t \in \mathbb{R}_{\geq 0}$ and $C \in \mathbb{S}/\mathcal{R}$:*

$$R(P, C, t) = R(Q, C, t),$$

where $R(P, C, t) = \sum_i \{|\lambda(t)|P \xrightarrow{\lambda(t)}_i P', P' \in C|\}$. P and Q are strongly bisimilar, denoted $P \sim Q$, if (P, Q) is contained in some strong bisimulation $\mathcal{R}$.

Here, $\{|\ldots|\}$ denotes a multiset. It follows that $\sim$ is the largest strong bisimulation, i.e., it contains any strong bisimulation. To be able to compare ICTMCs by bisimulation, let us equip an ICTMC with an initial state $s^0 \in \mathbb{S}$. Two ICTMCs $\mathcal{C}_P = (\mathbb{S}_P, \mathbf{R}_P, s^0_P)$ and $\mathcal{C}_Q = (\mathbb{S}_Q, \mathbf{R}_Q, s^0_Q)$ are bisimilar, denoted $\mathcal{C}_P \sim \mathcal{C}_Q$, iff their initial states are bisimilar, i.e., $s^0_P \sim s^0_Q$. The quotient of an ICTMC under $\sim$ is defined in the following way.

Definition 5 (Bisimulation quotient). *For the* ICTMC $\mathcal{C} = (\mathbb{S}, \mathbf{R}, s^0)$ *and* $\sim$*, the quotient $\mathcal{C}/\!\sim$ is defined by $\mathcal{C}/\!\sim = (\mathbb{S}/\!\sim, \mathbf{R}_\sim, s^0_\sim)$ where $s^0_\sim = [s^0]_\sim$ and $\mathbf{R}_\sim$ is defined by:* $R_\sim([P]_\sim, [P']_\sim, t) = R(P, [P']_\sim, t) \quad \text{for all } t \in \mathbb{R}_{\geq 0}$.

Note that $\mathcal{C}$ is strongly bisimilar to $\mathcal{C}/\!\sim$. An important property of strong bisimulation is that it preserves transient probabilities; in particular, this means that there is a strong relationship between the transient probabilities in an ICTMC and its quotient.

Theorem 2. *Let $\mathcal{C} = (\mathbb{S}, \mathbf{R}, s^0)$ be an* ICTMC*. For every $C \in \mathbb{S}/\sim$, the transient probability distribution $\pi_C(t)$ of the state C in the quotient chain $\mathcal{C}/\sim$ is:*

$$\pi_C(t) = \sum_{s \in C} \pi_s(t) \quad \text{for all } t \in \mathbb{R}_{\geq 0},$$

where $\pi_s(t)$ is the transient probability distribution of state $s \in \mathbb{S}$ in $\mathcal{C}$.

From Theorem 2 we may conclude that the steady state probability distribution (if it exists) is also preserved.

Corollary 1. *Let $\mathcal{C} = (\mathbb{S}, \mathbf{R}, s^0)$ be an* ICTMC*. For every $C \in \mathbb{S}/\sim$, the steady-state probability distribution π_C of the state C in the quotient chain $\mathcal{C}/\sim$ is:*

$$\pi_C = \lim_{t \to \infty} \pi_C(t) = \lim_{t \to \infty} \sum_{s \in C} \pi_s(t) = \sum_{s \in C} \pi_s,$$

where π_s is the steady-state probability distribution of state $s \in \mathbb{S}$.

In many cases it is reasonable to assume that two processes P and Q are equal up to time T. For this case we propose the *finite-horizon bisimulation*.

Definition 6. *An equivalence $\mathcal{R} \subseteq \mathbb{S} \times \mathbb{S}$ is a finite-horizon bisimulation whenever for all $(P, Q) \in \mathcal{R}$, $t \in [0, T]$ $(T \in \mathbb{R}_{\geq 0})$ and $C \in \mathbb{S}/\mathcal{R}$: $R(P, C, t) = R(Q, C, t)$. P and Q are finitely-horizon bisimilar, denoted $P \sim^T Q$, if (P, Q) is contained in some finite-horizon bisimulation $\mathcal{R}$.*

Notice that the definition of finite-horizon bisimulation $\sim^T$ is the same except that the time t lies in the interval $[0, T]$. It is easy to see that finite-horizon bisimulation preserves the transient distribution up to time T.

Proposition 1. *For $0 < \cdots < T < \cdots < \infty$ it holds: $\sim^0 \subseteq \; \cdots \subseteq \sim^T \cdots \subseteq \sim$.*

Thus , $P \sim^{t_i} Q$ implies $P \sim^{t_j} Q$ for every $t_j < t_i$. It follows that for $t_j < t_i$, the quotient under $\sim^{t_j}$ is coarser than under $\sim^{t_i}$.

Bisimulation for $\mathrm{I^2MCs}$*.* So far, we have presented bisimulation for ICTMCs. In order to define bisimulation for I^{2}MCs, unobservable actions (i.e., τ) require special care. Consider four states such that $P_1 \sim P_2 \sim Q_1 \sim Q_2$ (see diagram below).

$$\begin{tikzcd}
P_0 \arrow[r, "2\lambda(t)"] \arrow[dr, "\tau"'] & P_1 \arrow[d, dash, dashed, "\sim"] & Q_0 \arrow[r, "\lambda(t)"] \arrow[dr, "\lambda(t)"'] & Q_1 \arrow[d, dash, dashed, "\sim"] \\
& P_2 & & Q_2
\end{tikzcd}$$

At first sight, it seems natural that $P_0 \sim Q_0$ as $R(P_0, C, t) = R(Q_0, C, t) = 2\lambda(t)$. But, state P_0 can do something more. There is a transition $P_0 \xrightarrow{\tau} P_2$ which consumes no time since a τ-action is an internal one and is not prevented by the environment (maximal progress assumption). The probability that transition $P_0 \xrightarrow{2\lambda(t)} P_1$ will be taken in 0 time units is $\mathrm{Pr}_{P_0,P_1}(t, 0) = \int_0^0 2\lambda(t + \tau) e^{-\int_0^\tau 2\lambda(t+\ell) d\ell} d\tau = 0$. Thus, we may conclude that $P_0 \nsim Q_0$. When specifying the definition of bisimilarity we have to treat immediate actions (τ) in a special way. Let $\mathbb{S}$ be the state-space of an I^{2}MC.

$P + 0 = P$	$a.P + a.P = a.P$	$(P + Q) + R = P + (Q + R)$
$P + Q = Q + P$	$\lambda(t).P + \tau.Q = \tau.Q$	$\lambda(t).P + \mu(t).P = (\lambda(t) + \mu(t)).P$

Table 2. Sound and complete axioms for $\sim$ on the I^{2}ML sequential fragment.

Definition 7 (I^{2}MC strong bisimulation). *An equivalence $\mathcal{R} \subseteq \mathbb{S} \times \mathbb{S}$ is a strong bisimulation whenever for all $(P, Q) \in \mathcal{R}$, $t \in \mathbb{R}_{\geq 0}$, $a \in Act$ and $C \in \mathbb{S}/\mathcal{R}$:*

– *$P \xrightarrow{a} P'$ implies $Q \xrightarrow{a} Q'$ for some Q' and $(P', Q') \in \mathcal{R}$.*
– *$Q \xrightarrow{a} Q'$ implies $P \xrightarrow{a} P'$ for some P' and $(P', Q') \in \mathcal{R}$.*
– *$P \stackrel{\tau}{\not\rightarrow}$ (or $Q \stackrel{\tau}{\not\rightarrow}$) implies $R(P, C, t) = R(Q, C, t)$.*

P and Q are strongly bisimilar, denoted $P \sim Q$, if (P, Q) is contained in some strong bisimulation $\mathcal{R}$.

Example 2. Consider the I^{2}MC from Fig. 1 (c) and $\lambda_1(t) = \mu_1(t)$, $\lambda_2(t) = \mu_2(t)$. Its quotient under bisimulation is depicted in Fig. 2. The equivalence classes C_1, C_2 and C_3 contain the following states $C_1 = \{0\|_{\{use\}}Q, P\|_{\{use\}}0\}$, $C_2 = \{P\|_{\{use\}}use.Q, use.P\|_{\{use\}}Q\}$ and $C_3 = \{0\|_{\{use\}}use.Q, use.P\|_{\{use\}}0, 0\|_{\{use\}}0\}$.

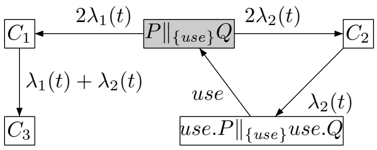

Fig. 2. Bisimulation quotient.

In a similar way as for ICTMCs, one can consider the quotient of an I^{2}MC. The compositional nature of I^{2}MC, however, allows in principle for obtaining such quotient in a component-wise manner, e.g., the quotient of $P\|_A Q$ can be obtained by first constructing the quotients of P and Q, then combine them, and quotienting the composition. The necessary requirement that $\sim$ needs to fulfill is that it is a *congruence* relation. The relation $\sim$ is a congruence whenever for processes P and Q it holds: $P \sim Q$ implies $\mathbf{C}[P] \sim \mathbf{C}[Q]$ where $\mathbf{C}[\cdot]$ is any context. (A context is basically a process term containing a hole that may be filled with any process.)

Theorem 3. *$\sim$ is a congruence with respect to all operators in* I^{2}ML.

Finite-horizon bisimulation is a congruence with one additional property.

Proposition 2. *For any processes P, P', Q, Q' and intervals $[0, T_1]$ and $[0, T_2]$ with $T_1, T_2 \in \mathbb{R}_{\geq 0}$ we have:*

$$P \sim^{T_1} P' \text{ and } Q \sim^{T_2} Q' \text{ implies } P\|_A Q \sim^{min(T_1,T_2)} P'\|_A Q' \text{ for all } A \subseteq Act.$$

As a next step, we consider the possibility to establish bisimulation symbolically, i.e., on the level of the syntax of the earlier introduced language I^{2}ML. This is facilitated by an axiomatization for $\sim$. The soundness of these axioms ensures that any two terms that are syntactically equal (denoted =) are bisimilar; formally, $P = Q \Rightarrow P \sim Q$. Whenever the axioms are complete, in addition, any strongly bisimilar processes can be represented by the same expressions in I^{2}ML, i.e., $P \sim Q \Rightarrow P = Q$. Summarizing, any bisimulation can be established syntactically, i.e., by just manipulating terms rather than I^{2}MCs, provided the axiom system is sound and complete. Let $\mathcal{A}_\sim$ be the set of axioms listed in Table 2 extended with the expansion law:

$$P\|_A Q = \sum_{i\in J_1} \lambda_i(t).\,(P_i\|_A Q) + \sum_{k\in J_3} \mu_k(t).\,(P\|_A Q_k) + \sum_{a_j=b_l\in A} a_j.\,(P_j\|_A Q_l) + \\ + \sum_{a_j\notin A\wedge a_j\in J_2} a_j.\,(P_j\|_A Q) + \sum_{b_l\notin A\wedge b_l\in J_4} b_l.\,(P\|_A Q_l)$$

where $P := \sum_{i\in J_1} \lambda_i(t).P_i + \sum_{j\in J_2} a_j.P_j$ and $Q := \sum_{k\in J_3} \mu_k(t).Q_k + \sum_{l\in J_4} b_l.Q_l$ with the finite index sets J_1, J_2, J_3 and J_4. Then the following holds:

Theorem 4. *For any $P, Q \in \mathbf{RG}$, $\mathcal{A}_\sim \vdash P = Q$ if and only if $P \sim Q$.*

The term $\mathbf{RG}$ denotes the set of all *regular* (no parallel composition inside recursion) and *guarded* (by actions or rates) expressions. While $\mathcal{A}_\sim \vdash P = Q$ means that $P = Q$ can be deduced from the set of sound and complete axiom system $\mathcal{A}_\sim$. The axiom $\lambda(t).P + \mu(t).P = (\lambda(t) + \mu(t)).P$ is due to the fact that the sum of two Poisson processes with rates $\lambda(t)$ and $\mu(t)$ is a Poisson process with the rate $\lambda(t)+\mu(t)$, whereas the axiom $\lambda(t).P + \tau.Q = \tau.Q$ is due to the maximal progress assumption. Notice that $\mathcal{A}_\sim$ also contains all standard axioms which involve hiding and recursion operators which are standard and omitted here.

Bisimulation minimization. The previous sections have set the stage for bisimulation minimization. Experiments have shown that in the traditional [11] as well as in the stochastic setting [19] exponential state space savings can be achieved. Given that $\sim$ is a congruence, individual processes can be replaced by their bisimilar quotient (under $\sim$) and the peak memory requirements can be reduced significantly. This all, however, requires an efficient bisimulation minimization algorithm. We adopt the *partition-refinement* paradigm to obtain a minimization algorithm for I^{2}MCs. As the problem for arbitrary rate functions is undecidable, we restrict to three classes of rate matrices $\mathbf{R}(t)$: piecewise uniform, polynomial ($\mathbf{R}(t) = t^{M+1}\mathbf{R}_{M+1} + \cdots + t\mathbf{R}_2 + \mathbf{R}_1$, where $\mathbf{R}_i$ with $i \leq M+1 \in \mathbb{N}$ are constant matrices) and piecewise polynomial (each piece is a polynomial of degree three). The same classes have been considered for the transient probability distribution, cf. Theorem 1. Rate comparisons and summations can easily be realized for these classes of functions. For rate matrix $\mathbf{R}$, let $M + 1$ denote the total number of intervals for piecewise uniform $\mathbf{R}(t)$, the polynomial degree when $\mathbf{R}(t)$ is polynomial, and the number of polynomial pieces when $\mathbf{R}(t)$ is piecewise polynomial.

Our bisimulation minimization algorithm for I^{2}MCs is based on a generalization of the algorithm for obtaining the coarsest quotient of a Markov chain under bisimulation by Derisavi *et al.* [7], and Paige-Tarjan's algorithm for LTS. The basic idea is to minimize iteratively over all pieces (or degrees of the polynomials). The bisimulation algorithm exploits an efficient data structure which groups all states with the same outgoing rate. This is in fact a binary tree where each *node* has four parameters: *node.left* and *node.right* - pointers to the left and right child, respectively, *node.sum* - stores the sum of the rates and *node.S* - stores all states with the same *node.sum*. Using such data structures, the time- and space complexity of bisimulation minimization for I^{2}MCs reduces to:

Theorem 5. *The coarsest quotient under $\sim$ of any I^{2}MC can be obtained in a worst-case time complexity $\mathcal{O}(m_a \lg(n) + M m_r \lg(n))$ and space complexity $\mathcal{O}(m_a + m_r)$, where m_a and m_r is the number of action-labeled and rate-labeled transitions, respectively.*

Recall that ICTMCs are I^{2}MCs that contain no action-labeled transitions. As a side result, the above theorem yields that the coarsest bisimulation quotient of a time-inhomogeneous CTMC can be obtained with time and space complexity $\mathcal{O}(M m_r \lg(n))$ and $\mathcal{O}(m_r)$, respectively. (The time complexity for homogeneous Markov chains is $\mathcal{O}(m_r \lg(n))$ [7]). Given the results in this paper that $\sim$ preserves transient and steady state distributions, our algorithm can be used to minimize prior to any such analysis.

Weak bisimulation for $\mathrm{I^2MCs}$. Strong bisimulation requires equivalent states to simulate their mutual stepwise behavior. While preserving the branching structure, strong bisimulation also requires mimicking of immediate actions (τ). As immediate actions consume no time it seems reasonable that two states will be equivalent regardless of the number of τ-steps in a sequence that they make. Therefore, the equivalence which will allow for the abstraction of sequences of immediate actions will be denoted as *weak bisimulation*. Let the transition $\stackrel{\tau}{\Longrightarrow}$ be the reflexive and transitive closure of $\stackrel{\tau}{\longrightarrow}^*$ and $\stackrel{a}{\Longrightarrow}$ a shorthand for $\stackrel{\tau}{\Longrightarrow}\stackrel{a}{\longrightarrow}\stackrel{\tau}{\Longrightarrow}$ $(a \neq \tau)$.

Definition 8 (I^{2}MC weak bisimulation). *An equivalence $\mathcal{R} \subseteq \mathbb{S} \times \mathbb{S}$ is a weak bisimulation whenever for all $(P, Q) \in \mathcal{R}$, $t \in \mathbb{R}_{\geq 0}$, $a \in Act$ and $C \in \mathbb{S}/\mathcal{R}$:*

- *$P \stackrel{a}{\longrightarrow} P'$ implies $Q \stackrel{a}{\Longrightarrow} Q'$ for some Q' and $(P', Q') \in \mathcal{R}$.*
- *$P \stackrel{\tau}{\not\rightarrow}$ implies $R(P, C, t) = R(Q'', C, t)$ for some $Q'' \stackrel{\tau}{\not\rightarrow}$ such that $Q \stackrel{\tau}{\Longrightarrow} Q''$ and $(P, Q'') \in \mathcal{R}$.*

For Q symmetric rules apply. P and Q are weakly bisimilar, denoted $P \approx Q$, if (P, Q) is contained in some weak bisimulation $\mathcal{R}$.

It seems intuitive that for the sequence $Q \stackrel{\tau}{\Longrightarrow} Q''$ the rates $R(P, C, t)$ and $R(Q'', C, t)$ have to be compared starting from time $t' = t + \Delta t$ where Δt is the time needed to make all τ in the sequence $Q \stackrel{\tau}{\Longrightarrow} Q''$. As τ transitions take no time the result will be the same even when the rates are compared from time t.

Example 3. Consider the $\mathrm{I^2MC}$ from Fig. 2 and its abstraction i.e. all actions are transformed into immediate ones (τ). The quotient under $\approx$ is depicted in Fig. 3, with C_1, C_2 and C_3 as in Fig. 2 and $C_0 = \{P\|_{\{use\}}Q,\ use.P\|_{\{use\}}use.Q\}$. It is important to note that after abstraction the transition labeled with use results in an immediate transition which gives the possibility to put the states $P\|_{\{use\}}Q$ and $use.P\|_{\{use\}}use.Q$ in the same equivalence class. Also note that the obtained I^{2}MC has no transitions labeled with actions, i.e., it is an ICTMC. This shows that weak bisimulation may be an effective mechanism to turn an $\mathrm{I^2MC}$ into an ICTMC, which may be subject to analysis as discussed in Section 2.

$a.\tau.P = a.P$	$P + \tau.P = \tau.P$	$\lambda(t).\tau.P = \lambda(t).P$	$a.(P + \tau.Q) + a.Q = a.(P + \tau.Q)$

Table 3. Sound and complete axioms for $\simeq$ on the $\mathrm{I^2ML}$ sequential fragment.

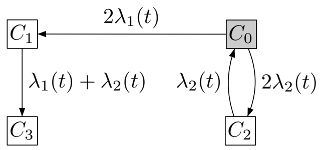

Fig. 3. Weak bisimulation quotient.

As in the case of strong bisimulation, weak bisimulation is also a congruence with respect to $\mathrm{I^2ML}$ operators. But there is an exception. Weak bisimulation is not a congruence with respect to the choice $(P + Q)$ operator [20]. This is due to the fact that weak bisimulation will equate two processes whenever one can do $\stackrel{\tau}{\Longrightarrow}$ and the other one can do nothing. In order to cope with the choice operator one has to differentiate between $\stackrel{a}{\Longrightarrow}$ and $\stackrel{\tau}{\Longrightarrow}\stackrel{a}{\longrightarrow}\stackrel{\tau}{\Longrightarrow}$ when $a = \tau$ as follows:

Definition 9 (Weak congruence). *$P$and Q are weakly congruent, denoted by $P \simeq Q$, whenever for all $a \in Act$, $t \in \mathbb{R}_{\geq 0}$ and $C \in \mathbf{RG}/\approx$:*

– *$P \xrightarrow{a} P'$ implies $Q \stackrel{\tau}{\Longrightarrow}\stackrel{a}{\longrightarrow}\stackrel{\tau}{\Longrightarrow} Q'$ for some Q' and $P' \approx Q'$.*
– *$Q \xrightarrow{a} Q'$ implies $P \stackrel{\tau}{\Longrightarrow}\stackrel{a}{\longrightarrow}\stackrel{\tau}{\Longrightarrow} P'$ for some P' and $P' \approx Q'$.*
– *$P \stackrel{\tau}{\not\rightarrow}$ (or $Q \stackrel{\tau}{\not\rightarrow}$) implies $R(P, C, t) = R(Q, C, t)$.*

Theorem 6. *$\simeq$ is a congruence with respect to all operators in $\mathrm{I^2ML}$.*

Consider the set of axioms from Table 2 and 3 together with axioms related to hiding and recursion operators as $\mathcal{A}_{\simeq}$. As for strong bisimulation the following also holds for weak congruence:

Theorem 7. *For any $P, Q \in \mathbf{RG}$, $\mathcal{A}_{\simeq} \vdash P = Q$ if and only if $P \simeq Q$.*

Recall that P and Q are regular and guarded process terms.

5 Concluding Remarks and Future Work

This paper presented a compositional formalism for time-inhomogeneous continuous-time Markov chains (ICTMCs), a subclass of piecewise deterministic Markov processes (PDPs). The main contributions are a full-fledged process algebra for interactive ICTMCs, congruence results for weak and strong bisimulation, and a polynomial-time quotienting algorithm. In addition, a new characterization of transient probabilities is provided for rate functions that are piecewise uniform. In contrast to works on communicating PDPs [24, 23, 25], this paper considers weak bisimulation, congruence results and axiomatization, and, more importantly a notion of bisimulation which respects maximal progress. Current work consists of investigating improvements to the quotienting algorithm akin to [8], model-checking algorithms [18], and simulation relations for ICTMCs.

Acknowledgment. This research has been performed as part of QUPES project that is financed by the Netherlands Organization for Scientific Research (NWO).

References

1. Alur, R., Grosu, R., Hur, Y., Kumar, V., Lee, I.: Modular specification of hybrid systems in CHARON. In HSCC, LNCS 1790: 6–19, 2000.
2. Alur, R., Grosu, R., Sokolsky, O., Lee, I.: Compositional modeling and refinement for hierarchical hybrid systems. *J. Log. Algebr. Program.*, **68**(1-2): 105–128, 2006.
3. Buchholz, P.: Exact and ordinary lumpability in finite Markov chains. *J. of Applied Probability*, **31**:59–75, 1994.
4. Bujorianu, M. L., Lygeros, J., Bujorianu, M. C.: Bisimulation for general stochastic hybrid systems. In HSCC, LNCS 3414: 198–214, Springer, 2005.
5. Cloth, L., Haverkort, B. R., Jongerden, M.: Computing battery lifetime distributions. In DSN, pp. 780–789, IEEE CS Press, 2007.
6. Davis, M. H. A.: *Markov Models and Optimization*. Chapman and Hall, 1993.
7. Derisavi, S., Hermanns, H., Sanders, W. H.: Optimal state-space lumping in Markov chains. *Inf. Processing Letters*, **87**(6):309–315, 2003.
8. Dovier, A., Piazza, C., Policriti, A.: An efficient algorithm for computing bisimulation equivalence. *Theor. Comput. Sci.*, **311**(1-3):221–256, 2004.
9. Everdij, M., Blom, H.: Piecewise deterministic Markov processes represented by dynamically coloured Petri nets. *Stochastics*, **77**(1):1–29, 2005.
10. Fernandez, J. C.: An implementation of an efficient algorithm for bisimulation equivalence. *Science of Computer Programming*, **13**:219–236, 1989.
11. Fisler, K., Vardi, M. Y.: Bisimulation minimization in an automata-theoretic verification framework. In FMCAD, LNCS 1522: 115–132, 1998.
12. Haghverdi, E., Tabuada P., Pappas, G. J.: Bisimulation relations for dynamical, control, and hybrid systems. *Theor. Comput. Sci.*, **342**(2-3): 229–261, 2005.
13. Han, T., Katoen, J.-P., Mereacre, A.: Compositional Modeling and Minimization of Time-Inhomogeneous Markov Chains. Technical report AIB200721, RWTH Aachen, Germany, 2007.
14. Hermanns, H.: *Interactive Markov Chains: The Quest for Quantified Quality*. LNCS 2428, Springer, 2002.
15. Hermanns, H., Herzog, U., Katoen, J.-P.: Process algebra for performance evaluation. *Theor. Comput. Sci.*, **274**(1-2):43–87, 2002.
16. Hoare, C. A. R.: *Communicating Sequential Processes*. Prentice-Hall, 1985.
17. Kailath, T.: *Linear Systems*. Prentice-Hall, 1980.
18. Katoen, J.-P.: Stochastic model checking. In *Stochastic Hybrid Systems*, pp. 79–106. CRC Press, 2006.
19. Katoen, J.-P., Kemna, T., Zapreev, I. S., Jansen, D. N.: Bisimulation minimisation mostly speeds up probabilistic model checking. In TACAS, LNCS 4424: 76–92, 2007.
20. Milner, R. J.: *Communication and Concurrency*. Prentice-Hall, 1989.
21. Paige, R., Tarjan, R.: Three partition refinement algorithms. *SIAM J. of Computing*, **16**(6):973–989, 1987.
22. Rindos, A., Woolet, S., Viniotis, I., Trivedi, K. S.: Exact methods for the transient analysis of non-homogeneous continuous-time Markov chains. In *Numerical Solution of Markov Chains (NSMC)*, pp. 121–134, Kluwer, 1995.
23. Strubbe, S. N., Julius, A. A., van der Schaft, A. J.: Communicating piecewise deterministic Markov processes. In ADHS, pp. 349–354, 2003.
24. Strubbe, S. N., van der Schaft, A. J.: Compositional modelling of stochastic hybrid systems. In *Stochastic Hybrid Systems*, pp. 47–77. CRC Press, 2006.
25. Strubbe, S. N., van der Schaft, A. J.: Bisimulation for communicating piecewise deterministic Markov processes (CPDPs). In HSCC, LNCS 3414: 623–639, 2005.