

Compositionality for Markov Reward Chains with Fast Transitions

J. Markovski^{1*}, A. Sokolova^{2**}, N. Trčka^{1***}, and E.P. de Vink¹

¹ Technische Universiteit Eindhoven, Formal Methods Group
Den Dolech 2, 5612 AZ, Eindhoven, The Netherlands

² University of Salzburg, Computational Systems Group
Jakob-Haringer-Straße 2, 5020 Salzburg, Austria

Abstract. A parallel composition is defined for Markov reward chains with fast transitions and for discontinuous Markov reward chains. In this setting, compositionality with respect to the relevant aggregation preorders is established. For Markov reward chains with fast transitions the preorders are τ -lumping and τ -reduction. Discontinuous Markov reward chains are ‘limits’ of Markov reward chains with fast transitions, and have related notions of lumping and reduction. In total, four compositionality results are shown. In addition, the two parallel operators are related by a continuity property.

Keywords: discontinuous Markov reward chains, Markov reward chains with fast transitions, parallel composition, compositionality, lumpability, reduction, Kronecker product and sum

1 Introduction

Compositionality is a central issue in the theory of concurrent processes. Discussing compositionality requires three ingredients: (1) a class of processes or models; (2) a composition operation on the processes; and (3) a notion of behaviour, usually given by a semantic preorder or equivalence relation on the class of processes. For the purpose of this paper, we will have semantic preorders and the parallel composition as operation. Therefore, the compositionality result can be stated as:

$$P_1 \geq \bar{P}_1, P_2 \geq \bar{P}_2 \implies P_1 \parallel P_2 \geq \bar{P}_1 \parallel \bar{P}_2$$

where P_1, P_2, \bar{P}_1 and \bar{P}_2 are arbitrary processes, \parallel and \geq denote their parallel composition and the semantic preorder relation. Hence, compositionality enables the narrowing of a parallel composition by composing simplifications of its components, thus avoiding the construction of the actual parallel system. In this paper, we study compositionality for augmented types of Markov chains.

* Corresponding author: j.markovski@tue.nl, Supported by Bsik-project BRICKS AFM 3.2.

** Supported by the Austrian Science Fund (FWF) project P18913.

*** Supported by the Netherlands Organization for Scientific Research (NWO) project 612.064.205.

Homogeneous continuous-time Markov chains, Markov chains for short, are among the most important and wide-spread analytical performance models. A Markov chain is given by a graph with nodes representing states and outgoing arrows determining the stochastic behavior of each state. An initial probability vector indicates which states may act as starting ones. Markov chains often come equipped with rewards that are used to measure their performance (e.g., throughput, utilization, etc.) [1]. In this paper, we focus on state rewards only, and we refer to a Markov chain with rewards as a Markov reward chain.


To cope with the ever growing complexity of the systems, several performance modeling techniques have been developed to support the compositional generation of Markov reward chains. Such are stochastic process algebras [2, 3], (generalized) stochastic Petri nets [4, 5], probabilistic I/O automata [6], stochastic automata networks [7], etc. The compositional modeling enables composing a bigger system from several smaller components. The size of the state space of the resulting system is in the range of the product of the sizes of the constituent state spaces. Hence, compositional modeling usually suffers from state space explosion.

In the process of compositional modeling, performance evaluation techniques produce intermediate constructs that are typically extensions of Markov chains featuring transitions with communication labels. In the final modeling phase, all labels are discarded and communication transitions are assigned instantaneous behavior. Previous work [8–10] gave an account of handling these models by using Markov chains with fast transitions, which present extension of the standard Markov reward chains with transitions decorated with a real-valued linear parameter. To capture the intuition that the labeled transitions are instantaneous, a limit for the parameter to infinity is taken. The resulting process is a generalization of the standard Markov chain that can perform infinitely many transitions in a finite amount of time. This model was initially studied in [11, 12] without rewards, and is called a discontinuous Markov reward chain. The process exhibits stochastic discontinuity and it is often considered as pathological. However, as shown in [12, 13, 5], it proves very useful for explanation of results. Here, we consider discontinuous Markov reward chains and Markov reward chains with fast transitions. These two models are intimately related: Markov reward chains with fast transitions are used for modeling, but the notions for these processes are expressed asymptotically in terms of discontinuous Markov reward chains. We define parallel composition of both models in vein of standard Markov reward chains [14] using Kronecker products and sums.

As already mentioned, compositional modeling may lead to state space explosion. Current analytical and numerical methods can handle Markov reward chains with millions of states efficiently. However, they only alleviate the problem and many real world problems still cannot be feasibly solved. Several aggregation techniques have been proposed to reduce the state space of Markov reward chains. Ordinary lumping is the most prominent one [15, 14]. The method partitions the state space into partition classes. In each class, the states exhibit equivalent behavior for transiting to other classes, i.e. the cumulative proba-


bility of transiting to another class is the same for every state of the class. If non-trivial lumping exists, i.e. at least one partition contains more than one state, then the method produces a smaller Markov chain that retains the performance characteristics of the original one. For example, the total reward gained in a given amount of time is the same for the original as for the reduced, so-called lumped, process. Another lumping-based method is exact lumping [14]. This method requires that each partition class of states has the same cumulative probability of transiting to every state of another class and also each state in the class has the same initial probability. The gain of exact lumping is that the probabilities of the original process can be computed for a special class of initial probability vectors by using the lumped Markov reward chain only.

A preliminary treatment of relational properties of lumping-based aggregations of Markov chains has been given in [16]. It has been shown that the notion of exact lumping is not transitive, i.e., there are processes which have exactly lumped versions that can be non-trivially exactly lumped again, but the original process cannot be exactly lumped directly to the resulting process. On the other hand, ordinary lumping of Markov reward chains is transitive and, moreover, it has a property of strict confluence. Strict confluence means that whenever a process can be lumped using two different partitions, there is always a smaller process to which the lumped processes can lump to. Coming back to our models of interest, ordinary lumping is defined for discontinuous Markov reward chains in [8–10]. Also, so-called τ -lumping is proposed for Markov reward chains with fast transitions in [8–10]. The situation can be pictured as follows:



In addition, the same papers [9, 10] provide an aggregation method by reduction that eliminates the stochastic discontinuity and reduces a discontinuous Markov reward chain to a Markov reward chain. The reduction method is an extension of the method described in [17]. It is based on the elimination of stochastic discontinuity that arises in the context of instantaneous probabilistic transitions. The method is well-known in perturbation theory. Its advantage is the ability to split states. The lumping method, in contrast, provides more flexibility: also states that do not exhibit discontinuous behavior can be aggregated. The reduction-based aggregation straightforwardly extends to τ -reduction of Markov reward

chains with fast transitions. Therefore, we have the following situation.



Both the lumping aggregation method and the reduction method induce semantic preorders. Namely, for two processes P and \bar{P} we say that $P \geq \bar{P}$ if \bar{P} is an aggregated version of P . The compositionality is very important as it allows us to aggregate the smaller parallel components first, and then combine them into the aggregated complete system. We show that the relations induced by the lumping and reduction methods indeed define preorders, i.e., reflexive and transitive relations. Having all the ingredients in place, we show the compositionality of the aggregation preorders with respect to the defined parallel composition(s). We also show continuity of the parallel composition(s). In short, the parallel operators preserve the two diagrams above.

The structure of the rest of the paper is as follows. We start by defining the first model, discontinuous Markov reward chains, in Section 2, together with its notions of lumping and reduction. Section 3 focuses on the second model, Markov reward chains with fast transitions, and introduces τ -lumping and τ -reduction. In Section 4, we show that the aggregation methods define preorders on the models. Section 5 contains the main results of the paper, compositionality of the new parallel operator for each type of Markov chains with respect to both aggregation preorders. Section 6 wraps-up with conclusions.

Notation All vectors are column vectors if not indicated otherwise. By $\mathbf{1}^n$ we denote the vector of n 1's; by $\mathbf{0}^{n \times m}$ the $n \times m$ zero matrix; by I^n the $n \times n$ identity matrix. We omit the dimensions n and m when they are clear from the context. By $A[i, j]$ we denote an element of the matrix $A \in \mathbb{R}^{m \times n}$ assuming $1 \leq i \leq m$ and $1 \leq j \leq n$. We write $A \geq 0$ when all elements of A are non-negative. The matrix A is called *stochastic* if $A \geq 0$ and $A \cdot \mathbf{1} = \mathbf{1}$. By A^T we denote the transpose of A .

Let \mathcal{S} be a set. A set $\mathcal{P} = \{S_1, \dots, S_N\}$ of N subsets of \mathcal{S} is called a *partition* of \mathcal{S} if $\mathcal{S} = S_1 \cup \dots \cup S_N$, $S_i \neq \emptyset$ and $S_i \cap S_j = \emptyset$ for all i, j , with $i \neq j$. The partitions $\mathcal{P} = \{\mathcal{S}\}$ and $\mathcal{P} = \{\{i\} \mid i \in \mathcal{S}\}$ are the trivial partitions. Let $\mathcal{P}_1 = \{S_1, \dots, S_N\}$ be a partition of \mathcal{S} and $\mathcal{P}_2 = \{T_1, \dots, T_M\}$, in turn, a partition of \mathcal{P}_1 . The *composition* $\mathcal{P}_1 \circ \mathcal{P}_2$ of the partitions \mathcal{P}_1 and \mathcal{P}_2 is a partition of \mathcal{S} given by $\mathcal{P}_1 \circ \mathcal{P}_2 = \{U_1, \dots, U_M\}$, where $U_i = \bigcup_{C \in T_i} C$.

2 Discontinuous Markov reward chains

In the standard theory (cf. [18, 19, 1]) Markov chains are assumed to be stochastically continuous. This means that when $t \rightarrow 0$, the probability of the process

occupying at time t the same state as at time 0 is 1. As we include instantaneous transitions in our theory [12], this requirement must be dropped. Therefore, we work in the more general setting of discontinuous Markov chains [11].

A discontinuous Markov reward chain is a time-homogeneous finite-state stochastic process with an associated (state) reward structure that satisfies the Markov property. It is completely determined by: (1) a stochastic row initial probability vector that gives the starting probabilities of the process for each state, (2) a transition matrix function $P(t)$ that defines the stochastic behavior of the transitions, at time $t > 0$, and (3) a state reward vector that associates a number to each state representing the gain of the process while spending time in the state. The transition matrix function gives a stochastic matrix $P(t)$ at any time $t > 0$, and has the property $P(t + s) = P(t) \cdot P(s)$ [18, 19]. It has a convenient characterization independent of time [12, 20], which allows for the following equivalent definition.

Definition 1. A discontinuous Markov reward chain \mathbf{P} is a quadruple $\mathbf{P} = (\sigma, \Pi, Q, \rho)$ where σ is a stochastic row initial probability vector, ρ is a state reward vector and $\Pi \in \mathbb{R}^{n \times n}$ and $Q \in \mathbb{R}^{n \times n}$ satisfy the following six conditions: (1) $\Pi \geq 0$, (2) $\Pi \cdot \mathbf{1} = \mathbf{1}$, (3) $\Pi^2 = \Pi$, (4) $\Pi Q = Q \Pi = Q$, (5) $Q \cdot \mathbf{1} = \mathbf{0}$, and (6) $Q + c\Pi \geq 0$, for some $c \geq 0$.

The matrix function $P(t) = \Pi e^{Qt}$ is the transition matrix of a discontinuous Markov chain $\mathbf{P} = (\sigma, \Pi, Q, \rho)$. It is continuous at zero if and only if $\Pi = I$. In this case, Q is a standard generator matrix [12, 8]. Otherwise, the matrix Q might contain negative non-diagonal entries. We note that, unlike for standard Markov reward chains, a meaningful graphical representation of discontinuous Markov reward chains when $\Pi \neq I$ is not possible. The intuition behind the matrix Π is that $\Pi[i, j]$ denotes the probability that a process occupies two states via an instantaneous transition. Therefore, in case of no instantaneous transitions, i.e., when $\Pi = I$, we get a standard (continuous) Markov chain.

For every discontinuous Markov chain $\mathbf{P} = (\sigma, \Pi, Q, \rho)$, Π gets the following ‘ergodic’ form after a suitable renumbering of the states [12]:

$$\Pi = \begin{pmatrix} \Pi_1 & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Pi_2 & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \Pi_M & \mathbf{0} \\ \overline{\Pi}_1 & \overline{\Pi}_2 & \dots & \overline{\Pi}_M & \mathbf{0} \end{pmatrix} \quad L = \begin{pmatrix} \mu_1 & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mu_2 & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mu_M & \mathbf{0} \end{pmatrix} \quad R = \begin{pmatrix} \mathbf{1} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{1} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{1} \\ \delta_1 & \delta_2 & \dots & \delta_M \end{pmatrix}$$

where for all $1 \leq k \leq M$, $\Pi_k = \mathbf{1} \cdot \mu_k$ and $\overline{\Pi}_k = \delta_k \cdot \mu_k$ for a row vector $\mu_k > 0$ such that $\mu_k \cdot \mathbf{1} = 1$ and a vector $\delta_k \geq 0$ such that $\sum_{k=1}^m \delta_k = \mathbf{1}$. Then the pair of matrices (L, R) depicted above forms a canonical product decomposition of Π (cf. Section 2.1 below), needed for the definition of the reduction-based method of aggregation.

The new numbering induces a partition $\mathcal{E} = \{E_1, \dots, E_M, T\}$ of the state set $\mathcal{S} = \{1, \dots, n\}$, where E_1, \dots, E_M are the ergodic classes, determined by

Π_1, \dots, Π_M , respectively, and T is the class of transient states, determined by any $\overline{\Pi}_i$, $1 \leq i \leq M$. The partition \mathcal{E} is called the ergodic partition. For every ergodic class E_k , the vector μ_k is the vector of ergodic probabilities. If an ergodic class E_k contains exactly one state, then $\mu_k = (1)$ and the state is called regular. The vector δ_k contains the trapping probabilities from transient states to the ergodic class E_k .

We are now able to explain the behavior of a discontinuous Markov reward chain $P = (\sigma, \Pi, Q, \rho)$. It starts in a state with a probability given by the initial probability vector σ . In an ergodic class with multiple states the process spends a non-zero amount of time switching rapidly (infinitely many times) among the states. The probability that it is found in a specific state of the class is given by the vector of ergodic probabilities. The time the process spends in the class is exponentially distributed and determined by the matrix Q . In an ergodic class with a single state the row of Q corresponding to that state has the form of a row in a generator matrix, and $Q[i, j]$ for $i \neq j$ is interpreted as the rate from i to j . In a transient state the process spends no time (with probability one) and goes to some ergodic class, where it is trapped for some amount of time. Note that $\delta_k[i] > 0$ iff $i \in T$ can be trapped in the ergodic class E_k .

The total reward gained by the process up to time $t > 0$, notation $R(t)$, is calculated as $R(t) = \sigma P(t)\rho$. We have that the total reward remains unchanged if the reward vector ρ is replaced by $\Pi\rho$. To see this, note that $P(t) = P(t)\Pi$ (cf. [12]), so $\sigma P(t)\Pi\rho = \sigma P(t)\rho = R(t)$. Intuitively, the reward in a transient state can be replaced by the sum of the rewards of the ergodic states that it can get trapped in as the process gains no reward while ‘residing’ in transient states. The reward of an ergodic state is the sum of the rewards of all states inside its ergodic class weighted according to their ergodic probabilities.

2.1 Aggregation methods

In this section we recall the definitions and the main properties of the aggregation methods for discontinuous Markov reward chains [8–10].

Ordinary lumping We define ordinary lumping in terms of matrices. Every partition $\mathcal{P} = \{C_1, \dots, C_N\}$ of $\mathcal{S} = \{1, \dots, n\}$ can be associated with a so-called collector matrix $V \in \mathbb{R}^{n \times N}$ defined as $V[i, k] = 0$ if $i \notin C_k$, $V[i, k] = 1$ if $i \in C_k$, and vice versa. The k -th column of V has 1’s for elements corresponding to states in C_k and has 0’s otherwise. Note that $V \cdot \mathbf{1} = \mathbf{1}$. A distributor matrix $U \in \mathbb{R}^{N \times n}$ for \mathcal{P} is defined as a matrix $U \geq 0$, such that $UV = I^N$. To satisfy these conditions, the elements of the k -th row of U , which correspond to states in the class C_k , sum up to one, whereas the other elements of the row are 0.

An ordinary lumping is a partition of the state space into classes such that the states that are lumped together have equivalent behavior for transiting to other classes and additionally they have the same reward.

Definition 2. A partition \mathcal{L} of $\{1, \dots, n\}$ is an ordinary lumping, or lumping for short, of a discontinuous Markov reward chain $\mathbf{P} = (\sigma, \Pi, Q, \rho)$ iff the following holds: (1) $VU\Pi V = \Pi V$, (2) $VUQV = QV$ and (3) $VU\rho = \rho$, where V is the collector matrix and U is any distributor matrix for \mathcal{L} .

The lumping conditions only require that the rows of ΠV (resp. QV and ρ) that correspond to the states of the same partition class are equal. The following property [8–10] holds.

Proposition 1. Let $\mathbf{P} = (\sigma, \Pi, Q, \rho)$ be a discontinuous Markov reward chain and let $\mathcal{L} = \{C_1, \dots, C_N\}$ be an ordinary lumping. Define (1) $\bar{\sigma} = \sigma V$, (2) $\bar{\Pi} = U\Pi V$, (3) $\bar{Q} = UQV$ and (4) $\bar{\rho} = U\rho$, for the collector matrix V of \mathcal{L} and any distributor U . Then $\bar{\mathbf{P}} = (\bar{\sigma}, \bar{\Pi}, \bar{Q}, \bar{\rho})$ is a discontinuous Markov reward chain. \square

Definition 3. If the conditions of Proposition 1 hold, then $\mathbf{P} = (\sigma, \Pi, Q, \rho)$ lumps to $\bar{\mathbf{P}} = (\bar{\sigma}, \bar{\Pi}, \bar{Q}, \bar{\rho})$, called the lumped discontinuous Markov reward chain, with respect to \mathcal{L} . We write $\mathbf{P} \xrightarrow{\mathcal{L}} \bar{\mathbf{P}}$.

It can readily be seen that neither the definition of a lumping, nor the definition of the lumped process depends on the choice of a distributor matrix U . In the continuous case when $\Pi = I$ we have $\bar{\Pi} = I$, so \bar{Q} is a generator matrix and our notion of ordinary lumping coincides with the standard definition [15, 21]. The total reward is preserved by ordinary lumping: The lumped process has the same reward $\bar{R}(t)$ as the one of the original process $R(t)$, i.e., $\bar{R}(t) = \sigma V U P(t) V U \rho = \sigma P(t) V U \rho = \sigma P(t) \rho = R(t)$.

Reduction The reduction-based aggregation method masks the stochastic discontinuity of a discontinuous Markov reward chain $\mathbf{P} = (\sigma, \Pi, Q, \rho)$ and transforms it into a Markov reward chain [17, 12, 9, 10]. The idea of the method is to abstract away from the behavior of individual states in an ergodic class. It is based on the notion of a canonical product decomposition.

Definition 4. Let $\mathbf{P} = (\sigma, \Pi, Q, \rho)$ and assume that $\text{rank}(\Pi) = M$, i.e., that there are M ergodic classes. A canonical product decomposition of Π is a pair of matrices (L, R) with $L \in \mathbb{R}^{M \times n}$ and $R \in \mathbb{R}^{n \times M}$ such that $L \geq 0$, $R \geq 0$, $\text{rank}(L) = \text{rank}(R) = M$, $L \cdot \mathbf{1} = \mathbf{1}$, and $\Pi = RL$.

A canonical product decomposition always exists and can be constructed from the ergodic form of Π (see page 5). Moreover, it can be shown that any other canonical product decomposition is permutation equivalent to this one. Since a canonical product decomposition (L, R) of Π is a full-rank decomposition, and since Π is idempotent, we also have that $LR = I^M$. Note that $R \cdot \mathbf{1} = \mathbf{1}$. Also, we have $L\Pi = LRL = L$ and $\Pi R = RLR = R$. Now we can define the reduction method.

Definition 5. For a discontinuous Markov reward chain $\mathbf{P} = (\sigma, \Pi, Q, \rho)$, the reduced discontinuous Markov reward chain $\bar{\mathbf{P}} = (\bar{\sigma}, I, \bar{Q}, \bar{\rho})$ is given by $\bar{\sigma} = \sigma R$, $\bar{Q} = LQR$ and $\bar{\rho} = L\rho$, where (L, R) is a canonical product decomposition. We write $\mathbf{P} \rightarrow_r \bar{\mathbf{P}}$.

The reduced process is unique up to a permutation of the states, since so is the canonical product decomposition. The states of the reduced process are given by the ergodic classes of the original process, the transient states are ‘ignored’. Intuitively, the transient states are split probabilistically between the ergodic classes according to their trapping probabilities. In case a transient state is also an initial state, its initial probability is split according to its trapping probabilities. The reward is calculated as the sum of the individual rewards of the states of the ergodic class weighted by their ergodic probabilities. Like lumping, the reduction also preserves the total reward: $\bar{R}(t) = \sigma RLP(t)RL\rho = \sigma \Pi P(t)\Pi\rho = R(t)$. In case the original process has no stochastic discontinuity, i.e., $\Pi = I$, the reduced process is equal to the original.

3 Markov reward chains with fast transitions

A Markov reward chain with fast transitions is obtained by adding parameterized, so-called fast, transitions to a standard Markov reward chain. The remaining standard transitions are referred to as slow. The behavior of a Markov reward chain with fast transitions is determined by two generator matrices Q_s and Q_f , which represent the rates of the slow transitions and the speeds of the fast transitions, respectively.

Definition 6. *A Markov reward chain with fast transitions $P = (\sigma, Q_s, Q_f, \rho)$ is a function assigning to each $\tau > 0$, the Markov reward chain*

$$P_\tau = (\sigma, I, Q_s + \tau Q_f, \rho)$$

where $\sigma \in \mathbb{R}^{1 \times n}$ is an initial probability vector, $Q_s, Q_f \in \mathbb{R}^{n \times n}$ are two generator matrices, and $\rho \in \mathbb{R}^{n \times 1}$ is the reward vector.

By taking the limit when $\tau \rightarrow \infty$, fast transitions become instantaneous. Then, a Markov reward chain with fast transitions behaves as a discontinuous Markov reward chain [12].

Definition 7. *Let $P = (\sigma, Q_s, Q_f, \rho)$ be a Markov reward chain with fast transitions. The discontinuous Markov chain $Q = (\sigma, \Pi, Q, \Pi\rho)$ is the limit of P , where the matrix Π is the so-called ergodic projection at zero of Q_f , that is $\Pi = \lim_{t \rightarrow \infty} e^{Q_f t}$, and $Q = \Pi Q_s \Pi$. If Q is the limit of P , we write $P \rightarrow_\infty Q$.*

The initial probability vector and the reward vector are not affected by the limit construction. Below we motivate the choice of using the reward vector $\Pi\rho$ instead of just ρ .

3.1 Aggregation methods

Next, we recall the aggregation methods for Markov reward chains with fast transitions.

τ -lumping The notion of τ -lumping is based on ordinary lumping for discontinuous Markov reward chains.

Definition 8. A partition \mathcal{L} of the state space of a Markov reward chain with fast transitions \mathbf{P} is called a τ -lumping, if it is an ordinary lumping of its limiting discontinuous Markov reward chain \mathbf{Q} , i.e. $\mathbf{P} \rightarrow_{\infty} \mathbf{Q}$.

Note that since we defined the reward of the limit by $\Pi\rho$, a τ -lumping may identify states with different rewards.

Like for ordinary lumping, we define the τ -lumped process by multiplying σ , Q_s , Q_f and ρ with a collector matrix and a distributor matrix. However, unlike for ordinary lumping, not all distributors are allowed. Following [8–10], we provide a class of special distributors, called τ -distributors, that yield a τ -lumped process.

Definition 9. Let $\mathbf{P} = (\sigma, \Pi, Q, \rho)$ be a discontinuous Markov reward chain. Let V be a collector matrix for this chain. A matrix W is a τ -distributor for V iff it is a distributor for V , i.e. $\Pi V W \Pi = \Pi V W$, and the entries of W for the transient states that lump only with other transient states are positive.

An alternative, explicit definition of the τ -distributors can be found in [8–10]. Having defined τ -distributors, we can define a τ -lumped process.

Definition 10. Let $\mathbf{P} = (\sigma, Q_s, Q_f, \rho)$ and let \mathcal{L} be a lumping with a collector matrix V , and a corresponding τ -distributor W . The τ -lumped Markov reward chain with fast transitions $\bar{\mathbf{P}} = (\bar{\sigma}, \bar{Q}_s, \bar{Q}_f, \bar{\rho})$ is defined as $\bar{\sigma} = \sigma V$, $\bar{Q}_s = W Q_s V$, $\bar{Q}_f = W Q_f V$, $\bar{\rho} = W \rho$. We say that \mathbf{P} τ -lumps to $\bar{\mathbf{P}}$ and write $\mathbf{P} \xrightarrow{\mathcal{L}} \bar{\mathbf{P}}$.

In general, for a lumping with collector V and distributor U , $U Q_s V$ and $U Q_f V$ are not uniquely determined, i.e., they depend on the choice of the distributor. The restriction to τ -distributors does not change this. Subsequently, the τ -lumped process depends on the choice of the τ -distributor. The motivation for restricting to τ -distributors is that all τ -lumped processes are then equivalent in the limit. This is shown in the following proposition that, in addition, gives the exact connection between lumping and τ -lumping [8].

Proposition 2. The following diagram commutes

$$\begin{array}{ccc} \mathbf{P} & \xrightarrow{\mathcal{L}} & \bar{\mathbf{P}} \\ \infty \downarrow & & \downarrow \infty \\ \mathbf{Q} & \xrightarrow{\mathcal{L}} & \bar{\mathbf{Q}} \end{array}$$

that is, if $\mathbf{P} \xrightarrow{\mathcal{L}} \bar{\mathbf{P}} \rightarrow_{\infty} \bar{\mathbf{Q}}$ and if $\mathbf{P} \rightarrow_{\infty} \mathbf{Q} \xrightarrow{\mathcal{L}} \bar{\mathbf{Q}}'$, then $\bar{\mathbf{Q}} = \bar{\mathbf{Q}}'$, for \mathbf{P} and $\bar{\mathbf{P}}$ Markov reward chains with fast transitions, and \mathbf{Q} , $\bar{\mathbf{Q}}$, $\bar{\mathbf{Q}}'$ the respective limiting discontinuous Markov reward chains. \square

Moreover, the τ -lumped processes that originate from the same Markov reward chain with fast transitions become exactly the same once all fast transitions are eliminated [9, 10].


Proposition 3. Let P be a Markov reward chain with fast transitions. Suppose $P \xrightarrow{\mathcal{L}} \bar{P}$ and \bar{P} has no fast transitions, i.e., its speed matrix is the zero matrix. Then, whenever $P \xrightarrow{\mathcal{L}} \bar{P}'$ for any (other) τ -distributor, it holds that $\bar{P} = \bar{P}'$. \square

τ -reduction We now define a reduction-based aggregation method called τ -reduction. It aggregates a Markov reward chain with fast transitions to an asymptotically equivalent Markov reward chain.

Definition 11. A Markov reward chain with fast transitions $P = (\sigma, Q_s, Q_f, \rho)$ τ -reduces to the Markov reward chain $R = (\bar{\sigma}, I, \bar{Q}, \bar{\rho})$, given by (1) $\bar{\sigma} = \sigma R$, (2) $\bar{Q} = LQ_s R$, and (3) $\bar{\rho} = L\rho$, where $P \rightarrow_{\infty} (\sigma, \Pi, Q, \Pi\rho)$ and (L, R) is a canonical product decomposition of Π . When P τ -reduces to R , we write $P \rightsquigarrow_{\tau} R$.

The following simple property relates τ -reduction to reduction. It holds since $LQR = LQ_s R$ and $L\Pi\rho = L\rho$.

Proposition 4. The following diagram commutes



that is, if $P \rightsquigarrow_{\tau} R$ and $P \rightarrow_{\infty} Q \rightarrow_{\tau} R'$, then $R = R'$, for P a Markov reward chain with fast transitions, Q a discontinuous Markov reward chain and R, R' (continuous) Markov reward chains. \square

4 Relational properties

We investigate the relational properties of ordinary lumping for discontinuous Markov reward chains and τ -lumping for Markov reward chains with fast transitions. The combination of transitivity and strong confluence ensures that iterative application of ordinary lumping yields a uniquely determined process. In the case of τ -lumping, by Proposition 2 and Proposition 3, only the limit of the final reduced process is uniquely determined, unless the final process contains no fast transitions.

There is no need to investigate the relational properties of reduction and τ -reduction, since they act in one step (no iteration is possible), in a unique way, between different types of models. Proofs of the results of this section can be found in [22].

The following result gives the transitivity of ordinary lumping. Actually, we show the transitivity of the relation \geq on discontinuous Markov reward chains defined by

$$P_1 \geq P_2 \iff (\exists \mathcal{L}) P_1 \xrightarrow{\mathcal{L}} P_2.$$

Transitivity enables replacement of repeated application of ordinary lumping by a single application using an ordinary lumping that is a composition of the individual lumpings.

Theorem 1. *Let $P \xrightarrow{\mathcal{L}} \bar{P}$ and let $\bar{P} \xrightarrow{\bar{\mathcal{L}}} \bar{\bar{P}}$. Then $P \xrightarrow{\mathcal{L} \circ \bar{\mathcal{L}}} \bar{\bar{P}}$. \square*

The above relation is clearly reflexive, since the trivial partition is always a lumping, i.e., we have $P \xrightarrow{\Delta} P$ where Δ is the trivial partition in which every class is a singleton. Transitivity of τ -lumping also holds, i.e. the relation \geq defined by

$$P_1 \geq P_2 \iff (\exists \mathcal{L}) P_1 \xrightarrow{\mathcal{L}} P_2$$

is transitive. This relation is reflexive as well, via the trivial lumping Δ .

Theorem 2. *Let $P \xrightarrow{\mathcal{L}} \bar{P}$ and let $\bar{P} \xrightarrow{\bar{\mathcal{L}}} \bar{\bar{P}}$. Then $P \xrightarrow{\mathcal{L} \circ \bar{\mathcal{L}}} \bar{\bar{P}}$. \square*

Lumping and τ -lumping also have the strict confluence property. In case of lumping this means that if $P \xrightarrow{\mathcal{L}_1} P_1$ and $P \xrightarrow{\mathcal{L}_2} P_2$, then there exist two partitions $\bar{\mathcal{L}}_1$ and $\bar{\mathcal{L}}_2$ such that $P_1 \xrightarrow{\mathcal{L}_1 \circ \bar{\mathcal{L}}_1} \bar{P}$ and $P_2 \xrightarrow{\mathcal{L}_2 \circ \bar{\mathcal{L}}_2} \bar{P}$. One can prove the strict confluence property by adapting the proof for Markov reward chains, from [16] for example.

5 Parallel composition and compositionality

In this section we define the parallel composition for each of the models, and prove the compositionality results. The definitions of parallel composition are based on Kronecker products and sums, as for standard Markov reward chains [14]. The intuition behind this is that the Kronecker sum represents interleaving whereas the Kronecker product represents synchronization. Let us first recall the definition of Kronecker product and sum.

Definition 12. *Let $A \in \mathbb{R}^{n_1 \times n_2}$ and $B \in \mathbb{R}^{m_1 \times m_2}$. The Kronecker product of A and B is a matrix $(A \otimes B) \in \mathbb{R}^{n_1 m_1 \times n_2 m_2}$ defined by*

$$(A \otimes B)[(i-1)m_1 + k, (j-1)m_2 + \ell] = A[i, j]B[k, \ell]$$

for $1 \leq i \leq n_1$, $1 \leq j \leq n_2$, $1 \leq k \leq m_1$ and $1 \leq \ell \leq m_2$.

The Kronecker sum of two square matrices $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{m \times m}$ is a matrix $(A \oplus B) \in \mathbb{R}^{nm \times nm}$ defined as $A \oplus B = A \otimes I^m + I^n \otimes B$.

We also need the notion of a Kronecker product of two partitions. Let \mathcal{L}_1 and \mathcal{L}_2 be two partitions with corresponding collector matrices V_1 and V_2 , respectively. Then $\mathcal{L}_1 \otimes \mathcal{L}_2$ denotes the partition corresponding to the collector matrix $V_1 \otimes V_2$. In this section we present our results without proofs. Proof outlines are given in [22].

5.1 Composing discontinuous Markov reward chains

First, we present the definition of parallel composition of discontinuous Markov reward chains. The intuition is that ‘rates’ interleave, and the probabilities of the instantaneous transitions synchronize, i.e., are independent.

Definition 13. Let $P_1 = (\sigma_1, \Pi_1, Q_1, \rho_1)$ and $P_2 = (\sigma_2, \Pi_2, Q_2, \rho_2)$ be two discontinuous Markov reward chains. Their parallel composition is defined as:

$$P_1 \parallel P_2 = (\sigma_1 \otimes \sigma_2, \Pi_1 \otimes \Pi_2, Q_1 \otimes \Pi_2 + \Pi_1 \otimes Q_2, \rho_1 \otimes \mathbf{1}^{|\rho_2|} + \mathbf{1}^{|\rho_1|} \otimes \rho_2).$$

The following theorem shows that the parallel composition of two discontinuous Markov reward chains is well defined.

Theorem 3. Let P_1 and P_2 be two discontinuous Markov reward chains. Then $P_1 \parallel P_2$ is a discontinuous Markov reward chain. \square

In the special case, when both discontinuous Markov reward chains are continuous, their parallel composition is again a Markov reward chain as defined in [14]. Moreover, the following property shows that the parallel composition of two discontinuous Markov reward chains has a transition matrix that is the Kronecker product of the individual transition matrices, corresponding to the intuition that the Kronecker product represents synchronization. It justifies the definition of the parallel composition.

Theorem 4. Let P_1 and P_2 be discontinuous Markov reward chains. If $P_1(t)$ is the transition matrix of P_1 and $P_2(t)$ is the transition matrix of P_2 , then the transition matrix of $P_1 \parallel P_2$ is given by $P_1(t) \otimes P_2(t)$. \square

It is easy to see that the total reward of the parallel composition is the sum of the total rewards of the components. The following theorem shows that both lumping and reduction are compositional with respect to the parallel composition of discontinuous Markov reward chains.

Theorem 5. If $P_1 \xrightarrow{\mathcal{L}_1} \bar{P}_1$ and $P_2 \xrightarrow{\mathcal{L}_2} \bar{P}_2$, then $P_1 \parallel P_2 \xrightarrow{\mathcal{L}_1 \otimes \mathcal{L}_2} \bar{P}_1 \parallel \bar{P}_2$. Also, if $P_1 \rightarrow_r \bar{P}_1$ and $P_2 \rightarrow_r \bar{P}_2$, then $P_1 \parallel P_2 \rightarrow_r \bar{P}_1 \parallel \bar{P}_2$. \square

5.2 Composing Markov reward chains with fast transitions

We now present the definition of the parallel composition of Markov reward chains with fast transitions. It comprises a Kronecker sum of the generator matrices, i.e. interleaving of the rates for both slow and fast transitions.

Definition 14. Let $P_1 = (\sigma_1, Q_{s,1}, Q_{f,1}, \rho_1)$ and $P_2 = (\sigma_2, Q_{s,2}, Q_{f,2}, \rho_2)$ be two Markov reward chains with fast transitions. Then their parallel composition is defined as:

$$P_1 \parallel P_2 = (\sigma_1 \otimes \sigma_2, Q_{s,1} \oplus Q_{s,2}, Q_{f,1} \oplus Q_{f,2}, \rho_1 \otimes \mathbf{1} + \mathbf{1} \otimes \rho_2).$$

It is not difficult to see that the parallel composition of Markov reward chains with fast transitions is well defined. In Fig. 1 we present an example of parallel composition of two Markov reward chains with fast transitions: 1c) is the parallel composition of 1a) and 1b). The initial probabilities are depicted left above each




Fig. 1. Parallel composition of Markov reward chains with fast transitions

state, and the reward values right above. An exception is 1c) where for readability the rewards are omitted. They are given by the vector

$$(r_0 + r_3, r_0 + r_4, r_0, r_1 + r_3, r_1 + r_4, r_1, r_2 + r_3, r_2 + r_4, r_2).$$

Next we show that τ -lumping and τ -reduction are also compositional, with respect to the parallel composition of Markov reward chains with fast transitions.

Theorem 6. *If $P_1 \xrightarrow{\mathcal{L}_1} \bar{P}_1$ and $P_2 \xrightarrow{\mathcal{L}_2} \bar{P}_2$, then $P_1 \parallel P_2 \xrightarrow{\mathcal{L}_1 \otimes \mathcal{L}_2} \bar{P}_1 \parallel \bar{P}_2$. Also, if $P_1 \rightarrow_r \bar{P}_1$ and $P_2 \rightarrow_r \bar{P}_2$, then $P_1 \parallel P_2 \rightarrow_r \bar{P}_1 \parallel \bar{P}_2$. \square*

Fig. 2 presents the aggregated versions of the Markov reward chains with fast transitions from Fig. 1. The Markov reward chain with fast transitions in 2c) is the parallel composition of the Markov reward chains with fast transitions in 2a) and 2b). Remarkably, the aggregated versions 2a), 2b) and 2c) can be obtained from 1a), 1b), 1c), respectively, by either applying τ -reduction or τ -lumping. The τ -lumpings used are $\{\{1, 2\}, \{3\}\}$ for 1a) and 1b), and $\{\{1, 2, 4, 5\}, \{3, 6\}, \{7, 8\}, \{9\}\}$ for 1c). By Theorem 6, we have that the chain in 1c), in fact, is the parallel composition of the chains in 1a) and 1b).

Having defined parallel composition for both models, we show how they are related: the limit of the parallel composition of two Markov reward chains with fast transitions is the parallel composition of the limits of the components (that are discontinuous Markov reward chains). Hence, a continuity property of the parallel composition holds as stated in the next theorem.

Theorem 7. *Let $P_1 \rightarrow_\infty Q_1$ and $P_2 \rightarrow_\infty Q_2$. Then $P_1 \parallel P_2 \rightarrow_\infty Q_1 \parallel Q_2$. \square*

6 Conclusion

We considered two types of performance models: discontinuous Markov reward chains and Markov reward chains with fast transitions. The former models represent the limit behavior of the later ones. For both types of models, we presented

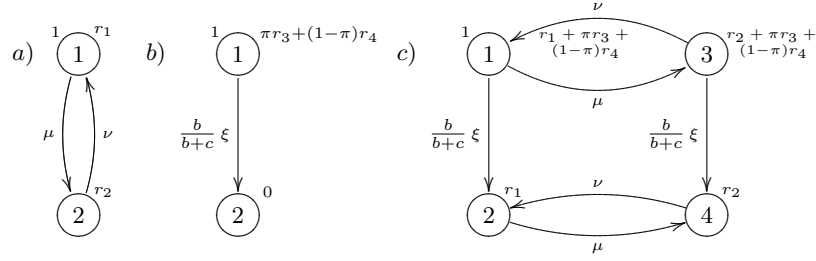


Fig. 2. Aggregated Markov reward chains with fast transitions

two aggregation methods: lumping and reduction for discontinuous Markov reward chains, respectively, τ -lumping and τ -reduction for Markov reward chains with fast transitions. In short, the contributions of the paper are:

- A definition of parallel composition of discontinuous Markov reward chains and of Markov reward chains with fast transitions, allowing for compositional modeling.
- Identification of preorder properties of the aggregation methods for both types of models.
- Compositionality theorems for each type of models and each corresponding aggregation preorder, and continuity property of the parallel compositions.




Fig. 3. Summary compositionality results

The results on compositionality are summarized by Fig. 3. In words, the parallel composition for Markov reward chains with fast transitions and the parallel composition for discontinuous Markov reward chains preserve the diagrams from Proposition 2 and Proposition 4. Fig. 3 is justified by the Theorems 1–6, as well as by Proposition 2 and Proposition 4.

For future work we schedule the analysis of extensions of Markov reward chains with fast and silent transitions to model both probabilistic and non-deterministic behavior. Our aim is to extend the compositionality results to that

setting, as well as to add action labeled transitions so that in addition to interleaving, synchronization can also be expressed.

References

1. Howard, R.: Semi-Markov and Decision Processes. Wiley (1971)
2. Hermanns, H.: Interactive Markov Chains: The Quest for Quantified Quality. Volume 2428 of LNCS. Springer (2002)
3. Hillston, J.: A Compositional Approach to Performance Modelling. Cambridge University Press (1996)
4. Ajmone Marsan, M., Balbo, G., Conte, G., Donatelli, S., Franceschinis, G.: Modelling with Generalized Stochastic Petri Nets. Wiley (1995)
5. Ciardo, G., Muppala, J., Trivedi, K.: On the solution of GSPN reward models. *Performance Evaluation* **12** (1991) 237–253
6. Wu, S.H., Smolka, S., Stark, E.: Composition and behaviors of probabilistic I/O automata. *Theoretical Computer Science* **176**(1–2) (1997) 1–38
7. Plateau, B., Atif, K.: Stochastic automata network of modeling parallel systems. *IEEE Transactions on Software Engineering* **17**(10) (1991) 1093–1108
8. Markovski, J., Trčka, N.: Lumping Markov chains with silent steps. In: QEST’06, Riverside, IEEE Computer Society (2006) 221–230
9. Markovski, J., Trčka, N.: Aggregation methods for Markov reward chains with fast and silent transitions. Technical Report CS 07/08, Technische Universiteit Eindhoven (2007)
10. Trčka, N.: Silent Steps in Transition Systems and Markov Chains. PhD thesis, Eindhoven University of Technology (2007)
11. Doeblin, W.: Sur l’équation matricielle $A(t+s) = A(t) \cdot A(s)$ et ses applications aux probabilités en chaîne. *Bull. Sci. Math.* **62** (1938) 21–32
12. Coderch, M., Willsky, A., Sastry, S., Castanon, D.: Hierarchical aggregation of singularly perturbed finite state Markov processes. *Stochastics* **8** (1983) 259–289
13. Ammar, H., Huang, Y., Liu, R.: Hierarchical models for systems reliability, maintainability, and availability. *IEEE Transactions on Circuits and Systems* **34**(6) (1987) 629–638
14. Buchholz, P.: Markovian process algebra: composition and equivalence. In: Proc. PAPM 94, Erlangen, Universität Erlangen-Nürnberg (1994) 11–30
15. Kemeny, J., Snell, J.: Finite Markov chains. Springer (1976)
16. Sokolova, A., de Vink, E.: On relational properties of lumpability. In: Proc. 4th PROGRESS symposium on Embedded Systems, Utrecht (2003)
17. Delebecque, F., Quadrat, J.: Optimal control of Markov chains admitting strong and weak interactions. *Automatica* **17** (1981) 281–296
18. Doob, J.: Stochastic Processes. Wiley (1953)
19. Chung, K.: Markov Chains with Stationary Probabilities. Springer (1967)
20. Hille, E., Phillips, R.: Functional Analysis and Semi-Groups. AMS (1957)
21. Nicola, V.: Lumping in Markov reward processes. IBM Research Report RC 14719, IBM (1989)
22. Markovski, J., Sokolova, A., Trčka, N., de Vink, E.: Compositionality for Markov chains with fast transitions. Technical Report CS 07/17, Technische Universiteit Eindhoven (2007)