

Compositional Model Checking and Model Repair for a Class of Product Form Models

Amin Soltanieh

*Department of Computer Science
Bundeswehr University Munich
Munich, Germany*

Markus Siegle

*Department of Computer Science
Bundeswehr University Munich
Munich, Germany*

Abstract

In the area of Markovian quantitative modelling, compositional model specification techniques such as Stochastic Process Algebra are widely used. However, exploiting a model's compositional structure for efficient analysis is still a difficult problem and mostly limited to special cases. This paper addresses some important issues in the area of compositional model checking of Markovian models for models with Boucherie-type product form. It closes a long-standing gap concerning the question whether compositional model checking of so-called global time-unbounded Until formulas is possible. The answer to this turns out to be negative. The paper then turns to the area of model repair, i.e. the question of how to fix a model in case it violates a given requirement. Here another general result and a useful proposition for compositional model repair are provided.

Keywords: Markov Chain, Product Form, Compositional Analysis, Probabilistic Model Checking, Continuous Stochastic Logic (CSL), Model Repair

1 Introduction

Markovian modelling and analysis is a versatile tool for quantitative system analysis, such as performance and dependability analysis, including probabilistic model checking. However, when applying Markovian techniques to large real-world case studies, users are often faced with the notorious problem of state space explosion, which can make practical application of the techniques intractable. In order to overcome or at least alleviate these difficulties, various strategies have been developed, such as Kronecker-based approaches [4], aggregation/disaggregation [18,5], abstraction [16], approximation [25], symbolic approaches [12], and, of course, combinations thereof. One particularly attractive approach is known as compositional analysis. Here the general idea is not to analyse the model as a whole, but to analyse parts (components) of the model individually and then in some way combine the partial results in order to obtain the desired overall result.

For special types of Markovian models, such as the well-known class of Product Form (PF) models, compositional computation of the system's steady-state probabilities is possible. It suffices to compute the steady-state probabilities of all the components in isolation and then combine them by multiplication, possibly with the need to calculate a normalising constant (in case some states of the product space are unreachable). Such an exploitation of a model's product form can lead to huge savings in runtime and memory used during analysis. For more complex types of analysis, such as probabilistic model checking, where the central problem is the

¹ Email: amin.soltanieh@unibw.de

² Email: markus.siegle@unibw.de

computation of path-based probabilities, it is difficult to exploit a model’s product form. Some results exist in the literature, but so far the open problems prevail.

The first contribution of this paper is to prove an impossibility result about the compositional model checking of so-called global time-unbounded CSL Until formulas. (CSL refers to the widely used Continuous Stochastic Logic [1], and the Until operator and its derivatives such as Eventually and Generally are the most important operators of this logic.) It is shown that this impossibility result holds not only for general compositional Markov models, but also for models with product form.

From model checking we then slightly change perspective and move to the problem of model repair, which is the problem of how to modify a model in case it violates a given requirement. We follow the idea of [22,23,10] to perform model repair by means of rate adaptation, and for this scenario we study the question whether model repair can be carried out in a compositional fashion. As our second contribution, we show with the help of a counterexample that in general this is not the case. However, as the third substantial contribution of the paper we show that for some important special type of Until operator, compositional model repair is indeed possible.

The further content of this paper is as follows: In Section 2, we introduce the Markovian modelling framework with a focus on product-form models, in particular models of Boucherie type. We also recapitulate the state-of-the-art of compositional model checking for such models and point out their current limitations (which are severe). In Section 3, we show that it is in general not possible to perform compositional model checking for the CSL Until operator, if the subformulas refer to state properties of more than one component of the overall model, even for the time-unbounded case, and even for models of Boucherie product form. Section 4 is devoted to the problem of model repair by rate adaptation, which is closely related to model checking. Here we provide another impossibility result, but also state a proposition which allows compositional model repair for quite a general class of CSL global Until formulas. Section 5 concludes the paper with a short summary of its main findings and a discussion of possible future work.

2 Modelling Framework and State-of-the-Art

This paper works with models which are Continuous-Time Markov Chains (CTMC), labelled with state and/or transition labels. A CTMC with state labels is referred to as a State-labelled Markov Chain (SMC). Each transition carries a transition rate, there are no immediate transitions, which means that our models are fully probabilistic. We consider a rather general compositional setting in which complex models are constructed from smaller components. The precise nature of composition may vary, depending on the concrete modelling formalism used. For instance, if we work with a Stochastic Process Algebra (SPA) such as PEPA [13] or CASPA [17], components are specified as sequential processes which cooperate through action synchronisation. If one works with Stochastic Petri Nets (SPN) [6,11] or Stochastic Activity Networks (SAN) [19], interaction of components may be realised by shared transitions and/or superposition of places. If one considers Boucherie’s framework for compositional Markov chains, which is of particular interest for this paper since its models are of product form (see Sec. 2.2), the interaction between components is realised indirectly via shared resources.

2.1 Product Form Models

At the outset, product form was discovered in queueing networks like Jackson’s theorem [15] which uses the product of equilibrium distributions of separate queues in order to calculate the joint distribution. Since then, there has been a lot of research towards product form solution for Stochastic Process Algebra (SPA) [20], stochastic Petri nets [7] and different types of Markovian models.

Boucherie [3] introduced a framework for the product process of a collection of Markov chains, to model the competition over resources by the exclusion of a part of the state space and blocking conditions. In this paper we mainly focus on the Boucherie framework regarding model checking and model repair (see section 2.2). There are some papers building on Boucherie’s framework: Hillston and Thomas [14] characterised the class of competing Markov processes identified by Boucherie using Performance Evaluation Process Algebra (PEPA). Their models are slightly different from Boucherie’s, while not really extending the Boucherie framework, but thanks to PEPA, the results are easier to understand and the components of the system and resources are explicitly represented in the model. Sereno [20] also used PEPA to define the product form equilibrium distribution for a certain class of SPA models satisfying some conditions. In [20] the components can be synchronised over some actions, whereas in [14] components must be independent and mutually exclusive and also resources must be independent: $(P_1 || P_2 \dots) \bowtie_S (R_1 || R_2 \dots)$, where P_i is a component, R_i is a resource and S is the set of synchronising actions. Finding an efficient way to calculate the normalizing constant is challenging in both [14] and [20].

Fourneau, Plateau and Stewart [9] assume a synchronisation-free system but with functional rates. So transitions of one component may depend on states of other components. Since the rates are functional, there is a set of transition matrices (instead of a single matrix) and using the kernel of these matrices and

under certain conditions, it is proved that product form holds. By using functional rates, [9] generalises some previously published papers and it is proved that the Boucherie framework of competing Markov chains is one of them, where the rate function is an indicator function.

Later Fourneau [8] considered a discrete-time generalisation of Boucherie’s strong blocking model [3] based on the generalised tensor (or Kronecker) product, and obtained product form steady-state solutions for DTMCs described as a generalised product of functional matrices. The paper [8] extends Boucherie’s theory and is one of the few papers about discrete-time networks of queues with product form steady-state solutions.

Thomas and Harrison [24] have used functional rates, system reversed model and minimal cycles in order to find a semi-product solution for the equilibrium state probabilities, while the functions satisfy specific forms. It is semi-product form because product form holds when some necessary conditions are satisfied.

2.2 Boucherie’s framework

In [3], Boucherie introduced a framework for compositional Markov chains, named “competing Markov chains”, which we summarise in this section. The model consists of a collection of independent processes competing over some resources, and for this collection the product process is introduced. In each transition of the product process, only one of the underlying Markov chains changes its state. Another condition characterizing this framework is that resources are mutually exclusive and there is so-called strong blocking. This means that while one component is using a resource, other components who compete over this resource are completely blocked and cannot move at all. Therefore, the competition is modelled as an exclusion of a part of the product state space and removal of some transitions.

Assume that C is a compositional CTMC with state space S , consisting of K finite CTMCs, C_k , $1 \leq k \leq K$, with state space S_k . Let C_{ki} be the CTMCs which compete with component k over resource i , and $A_{ki} \subseteq S_k$ denote those states of S_k where component k uses resource i . The state space of the product process is a subset of the Cartesian product of the state space of all components i.e. $S_1 \times \dots \times S_K$, excluding the states which violate the Boucherie condition of strong blocking:

$$S = \prod_{k=1}^K S_k \setminus \prod_{k=1}^K \prod_{i \in I} \prod_{j \in C_{ki}} A_{ki} \times A_{ji}$$

Combined states are thus vectors of components’ states $s = (s_1, \dots, s_K)$. The transition rates of the product process, considering the strong blocking condition, are defined as:

$$q(s, s') = \sum_{k=1}^K \left(q_k(s_k, s'_k) \cdot \prod_{r=1, r \neq k}^K \mathbf{1}(s_r = s'_r) \cdot \mathbf{1}(\forall i : (s_r \in A_{ri} \rightarrow k \notin C_{ri})) \right)$$


Here, $q_k(s_k, s'_k)$ is the local transition rate in the k ’th component, $\mathbf{1}(\cdot)$ is the indicator function and $s, s' \in S$ and $s_k, s'_k \in S_k$. The term $\mathbf{1}(s_r = s'_r)$ expresses that in each transition exactly one component changes its state, and the term $\mathbf{1}(\forall i : (s_r \in A_{ri} \rightarrow k \notin C_{ri}))$ guarantees that if resource i is being used by component r then all other components competing over resource i with component r are blocked until component r stops using resource i . This is referred to as the strong blocking property of Boucherie product form models.

Under the above mentioned conditions of the Boucherie framework, it is proved that the steady state probability of a state $s = (s_1, \dots, s_K)$ of the process C is of product form:

$$\pi(s) = B \prod_{k=1}^K \pi_k(s_k), s \in S$$


where B is the normalizing constant and $\pi_k(s_k)$ denotes the steady state probability of state s_k in component k when that component is analysed in isolation.

Example 2.1 As an example, in Fig. 1, we consider $K = 2$ components, named C_1 and C_2 , each having three states. The Greek letters denote the transition rates. In this example we extend the Boucherie framework by assigning atomic propositions to each state which are true in this state. The set of all atomic propositions is $AP = \bigcup_{k=1}^K AP_k$, where AP_k is the set of atomic propositions used in component k . In this example $AP_1 = \{p, q, r\}$, $AP_2 = \{s, t, u\}$, and thus $AP = \{p, q, r, s, t, u\}$. We will use atomic propositions for the purpose of model checking later in this paper.


 Fig. 1: CTMCs C_1 and C_2 with one resource

There is only one resource in this system (not modelled explicitly) and its use is highlighted in Fig. 1 by the bold circles. So, if component C_1 is in state 2, component C_2 is blocked and if component C_2 is in state 3, component C_1 is blocked. Therefore $A_{11} = \{2\}$ and $A_{21} = \{3\}$. Similarly, $A_{10} = \{1, 3\}$ and $A_{20} = \{1, 2\}$, which define the states in S_1 resp. S_2 with no resource. $C_{11} = \{2\}$ and $C_{21} = \{1\}$ represent the set of components which are competing over the resource with the first resp. second component.

The flat model of this system is provided in Fig. 2. This is the result of the parallel composition of the two components C_1 and C_2 by considering the effect of Boucherie conditions. Remember that in the Boucherie framework, there is no synchronisation between components, they are just related with each other by the means of resources.


 Fig. 2: Parallel composition of C_1 and C_2 and the effect of Boucherie's strong blocking condition

According to equation 2.2, states in $A_{11} \times A_{21} = \{(2, 3)\}$ must be excluded from the state space, so in Fig. 2, state (2, 3) and all the dashed transitions are removed and excluded due to the strong blocking condition. State (2, 3) is not reachable because the resource cannot be used by two components at the same time. Also, we see that in states (2, 1) and (2, 2), where the resource is occupied by component C_1 , the other component C_2 is blocked, so the corresponding transitions, depicted by dashed arcs in the figure, are removed. Likewise, in states (1, 3) and (3, 3) the resource is occupied by component C_2 , so component C_1 is blocked.

It can easily be shown that product form holds for this example. For instance, we obtain the steady state distribution of state (1, 1) of the flat model compositionally by

$$\pi(1, 1) = B \cdot \pi_1(1) \cdot \pi_2(1)$$

where $\pi_1(1)$ is the steady state probability for state 1 in component C_1 and $\pi_2(1)$ is the steady state probability for state 1 in component C_2 . Considering each component in isolation we find that:

$$\pi_1(1) = \frac{\delta\gamma}{\alpha\beta + \alpha\gamma + \delta\gamma} \quad , \quad \pi_2(1) = \frac{\epsilon\mu + \epsilon\nu}{\epsilon\lambda + \epsilon\mu + \epsilon\nu + \lambda\nu}$$

The normalizing constant B is determined by the flat model and in this example:

$$B = \frac{1}{[\pi_1(1) + \pi_1(3)] \sum_{i=1}^3 \pi_2(i) + \pi_1(2) \sum_{i=1}^2 \pi_2(i)}$$

2.3 Compositional Model Checking

Compositional approaches to model checking are very valuable, since they can help to deal with the problem of state space explosion. “Compositional” in this context means that the behaviour of a composed system relies on the behaviour of its components. Ideally, in a concurrent system composed of K components, one would like to model check a given overall requirement Φ by verifying some derived subrequirements Φ_k ($k = 1, \dots, K$), one for each component, and then combining their results in order to obtain the overall result. However, it is known that this ideal is very difficult to realise. Up to now, no generally applicable compositional model checking techniques for probabilistic systems have been found, all published approaches are only applicable to special cases.

In [2], Ballarini and Horvath present techniques for compositional model checking of Boucherie-type models against requirements specified by the popular logic CSL (Continuous Stochastic Logic [1]). They offer solutions for model checking the time-bounded and time-unbounded Next operator and the single component time-unbounded Until operator in a compositional fashion. In that context, “single component” means that both subformulas Φ' and Φ'' which are part of the Until formula $Pr_{\sim b}(\Phi'U\Phi'')$ (where \sim is a comparison operator and b is a probability bound) refer to state properties of only one single component of the model, which of course is a severe limitation. Conversely, a “global” formula would refer to state properties of more than one (possibly all) components.

For instance, consider the time-bounded global Next formula $Pr_{\sim b}(X^I\Phi)$, where $I = [a, b] \subseteq \mathbb{R}$ is a real time interval. It is shown in [2] that if the model is of Boucherie type and if Φ is of a special form, then compositional model checking of such a Next formula is possible. To be precise, Φ must be written in disjunctive normal form as $\Phi = \bigvee_i \bigwedge_j a_{ij}$, where a_{ij} is an atomic proposition referring to some component k (i.e. $a_{ij} \in AP_k$ for some k). Given a system of K components $C_k, k \in \{1, \dots, K\}$ and a state $s = (s_1, \dots, s_K)$ to be checked, a state-dependent formula $\xi_k(s)$ is defined for every component:

$$\xi_k(s) = \begin{cases} \neg tt & \text{if } \forall i \exists j a_{ij} \in AP_l, l \neq k \wedge s_l \neq a_{ij} \\ tt & \text{if } \exists i \forall j a_{ij} \in AP_l, l \neq k \wedge s_l = a_{ij} \\ \bigwedge_{i: \forall j', a_{ij'} \in AP_l, l \neq k, s_l = a_{ij'}} \bigvee_{j: a_{ij} \in AP_k} a_{ij} & \text{otherwise} \end{cases}$$

The probability of this formula can then be composed as follows [2]:

$$Pr(s, X^I\Phi) = \sum_{k \notin B(s)} p^k(s) Pr_k(s_k, X_{\frac{1}{p^k(s)}}^I \xi_k(s_1, \dots, s_K))$$

where $B(s)$ is the set of components which are blocked in state s , $p^k(s)$ is the probability that in state s component k is the first component to make a transition, and the probabilities $Pr_k(\cdot)$ can be obtained from model checking the individual submodels (note the stretched time intervals which reflect the fact that component k was the winner of the race among components in state s).

The paper [2] also addresses the question of compositional model checking of time-unbounded global Until formulas, but this part remains very vague and does not lead to a solution. As we will show in the next section, there is a good reason for this, and it is exactly at this point where we come in with the present paper and make our contributions. It actually turns out to be impossible to perform compositional model checking of global Until formulas, even in a Boucherie product form setting.

3 Compositional model checking for Boucherie product form

In this section we focus on global time-unbounded CSL Until formulas for a Boucherie-type model and prove by counterexample that, in general, compositional model checking for this specific type of formula is not possible. We start with a state classification inspired by [22,23].

Definition 3.1 (Partitioning of SMC) When an SMC M and the CSL requirement $P_{\sim b}(\varphi)$, where φ is an time-unbounded Until formula and ($\sim \in \{<, \leq, >, \geq\}$) are given, then for each state $s \in M$:


- If $Pr(s, \varphi) = 0$ then state s is placed in class Impossible.
- If $0 < Pr(s, \varphi) < 1$ then state s is placed in class Gobothways.
- If $Pr(s, \varphi) = 1$ then state s is placed into class Certain.

Theorem 3.2 *Compositional model checking of global time-unbounded CSL Until formula is in general not possible, even for Boucherie product form models.*

Proof. We prove theorem 3.2 by providing a counterexample. Take the example in Fig. 1 and consider the CSL global time-unbounded formula $P_{\geq b}(s_0, \varphi)$ where:

$$\varphi = (q \vee (r \wedge t)) U ((p \wedge s) \vee (r \wedge u))$$

This is a global formula since it refers to state properties of more than one component, i.e. both components. Model checking a state $s_0 = (x, y)$ of the combined process means computing the probability $Pr((x, y), \varphi)$ and comparing it to the probability bound b . This path probability can be computed on the combined model, which is shown in Fig. 2, by applying the standard CSL model checking algorithm for time-unbounded Until. The resulting satisfaction probabilities are given in Table 1 and the partitioned state space is shown in Fig. 3. As we see in Fig. 3, from now on, we draw Gobothways state in blue colour, Impossible states in red and Certain states in green.



(x, y)	$Pr((x, y), (q \vee (r \wedge t)) U ((p \wedge s) \vee (r \wedge u)))$
(1, 2)	0
(1, 3)	0
(3, 1)	0
(1, 1)	1
(3, 3)	1
(2, 1)	$\frac{\delta}{\delta + \beta}$
(2, 2)	$\frac{\nu \beta}{(\mu + \nu + \gamma)(\beta + \delta) - \gamma \beta}$
(3, 2)	$\frac{\nu(\beta + \delta)}{(\mu + \nu + \gamma)(\beta + \delta) - \gamma \beta}$

Fig. 3: Flat model with partitioned states

Table 1: Satisfaction probabilities for the combined process

However, since this is a product form model, one would wish to compute $Pr((x, y), \varphi)$ compositionally, i.e. by checking some sub-property on state x of process C_1 , checking another sub-property on state y of process C_2 , and then combining the results. It is clear that such sub-properties would also have to be Until-type formulas, since any path satisfying φ could potentially make several steps both in the C_1 - and in the C_2 -dimension of the combined model. For this purpose, let us consider the “projections” of φ on C_1 and C_2 :

$$\varphi_1 = (q \vee (r \wedge ?)) U ((p \wedge ?) \vee (r \wedge ?)) \quad (1)$$

$$\varphi_2 = (? \vee (? \wedge t)) U ((? \wedge s) \vee (? \wedge u)) \quad (2)$$

where the “?” stand for any unknown state formula that cannot be decided locally. How could one possibly evaluate φ_1 on process C_1 in isolation? One would have to make assumptions on the truth values of the three “?”, corresponding to the atomic propositions t, s, u of process C_2 . There are $2^3 = 8$ possible combinations of those truth values, all of which we should consider. But we can make our life a bit easier by having a look at process C_2 (this is cheating a bit, since we actually want to analyse process C_1 in isolation, but this cheating helps us to discard irrelevant cases). We see that in each state of C_2 exactly one of t, s, u is true, so it is enough to consider the three combinations shown in Table 2, where s_1 is the state of component C_1 :

$t \ s \ u$	φ_1	$Pr_1(s_1, \varphi_1)$			
		$s_1 :$	1	2	3
$tt \ ff \ ff$	$(q \vee r) \ U \ ff$		0	0	0
$ff \ tt \ ff$	$q \ U \ p$		1	$\frac{\delta}{\beta+\delta}$	0
$ff \ ff \ tt$	$q \ U \ r$		0	$\frac{\beta}{\beta+\delta}$	1

 Table 2: Satisfaction probabilities for φ_1 on process C_1

It is actually possible to derive the probability $Pr((2, 1), \varphi)$ from this table, in the following way: In state $(2, 1)$, atomic proposition s (from C_2) is true, so we look at the second row of Table 2. There we find the result for state 2 of C_1 : It is $\frac{\delta}{\beta+\delta}$, as we had already established in Table 1. The reason why this probability for the combined process can be obtained by only evaluating φ_1 on process C_1 is that from state $(2, 1)$, all paths that satisfy formula φ (there is actually only a single such path) move only along the C_1 -dimension of the global state space.

Since this partial result could strengthen our hope that compositional model checking is indeed possible, let's try to find the similar argument for state $(2, 2)$. In state $(2, 2)$ atomic proposition t of process C_2 is satisfied, so we have to look at the first row of Table 2. For state 2 we find the probability 0. That means that from state $(2, 2)$ there is no path satisfying φ that moves only along the C_1 -dimension of the combined model. So if there is any satisfying path originating in $(2, 2)$, such a path would have to include also moves along the C_1 -dimension. Maybe we can combine the result we just found with a complementary result for process C_2 in order to obtain the satisfaction probability for $(2, 2)$?

Therefore, let us now focus our attention on process C_2 and try to evaluate φ_2 in isolation. Again, we have to make assumptions about the four “?”, corresponding to atomic propositions p, q, r of process C_1 . Again, using the extra bit information that in process C_1 those atomic propositions are mutually exclusive, it is enough to consider three combinations as shown in Table 3 where s_2 is the state of C_2 :

$p \ q \ r$	φ_2	$Pr_2(s_2, \varphi_2)$			
		$s_2 :$	1	2	3
$tt \ ff \ ff$	$ff \ U \ s$		1	0	0
$ff \ tt \ ff$	$tt \ U \ ff$		0	0	0
$ff \ ff \ tt$	$t \ U \ u$		0	$\frac{\nu}{\nu+\mu}$	1

 Table 3: Satisfaction probabilities for φ_2 on process C_2

We are interested in state $(2, 2)$. Here the state of C_1 is 2 where q holds, so we look at the second row of Table 3. For state 2 we read the result 0. Now obviously, there is no way we can combine the two partial results we obtained (0 and 0) in order to obtain the true satisfaction probability of $(2, 2)$ which is: $\frac{\nu\beta}{(\mu+\nu+\gamma)(\beta+\delta)-\gamma\beta}$! The reason is that all paths originating in $(2, 2)$ and satisfying φ , consist of moves along both the C_1 - and C_2 -dimensions. There are actually infinitely many such paths (since the loop $(2, 2) \rightleftharpoons (3, 2)$ may be taken infinitely often) and the partial results from Tables 2 and 3 are not able to capture this behaviour. If we are interested in state $(3, 2)$ of the composed model, we would find similar reasons why $Pr((3, 2), \varphi)$ cannot be composed from two results gained on process C_1 and C_2 in isolation.

For this model it is thus not possible to compute the probability φ from probabilities of sub-properties evaluated on the constituent processes C_1 and C_2 . This is a counterexample from which we can conclude that compositional model checking of time-unbounded Until formula is in general not possible for composed models, even if they are of Boucherie product form. \square

4 Compositional model repair for product form

The model checking problem consists in verifying whether a model M satisfies a given user requirement ϕ , i.e. whether $s \models \phi$ for some state $s \in S$ of the model M (or possibly for all states of M). If the requirement is violated by the model, one can attempt to modify the model, with the goal that the modified model M' should indeed satisfy the given requirement. This process is referred to as “model repair”. Tati et al. [22,23] introduced a model repair solution using rate reduction. In this method, controllable states are defined as

those states where the modification of some well-defined subset of transition rates will improve the behaviour of the SMC without contradicting to the goal of the model repair problem. In their work, the states of M are partitioned according to Definition 3.1 into the three classes Certain, Impossible and Gobothways, and model repair is performed by reducing some specific transition rates of the model. They showed that as a result of this, all states of class Gobothways will eventually satisfy the requirement.

The algorithms presented in Tati's work all work at the level of the flat, low-level Markov chain, which has the advantage that the approach can be used for many different types of models. On the negative side, this means that even if the model has been specified in a compositional way, that structural information is not used at all during model repair. Even worse, once suitable rate modification factors have been found during model repair, it is difficult or even impossible to lift those factors from the flat Markov chain to the compositional high-level model. This is clearly a disadvantage, since users wish to know how they should modify their model specification, in order to get the requirements satisfied. We show this limitation using a simple example, and in the following subsection, we will address the problem whether and under which circumstances truly compositional model repair is possible.

Example 4.1 In Fig. 4, inspired by a simpler example in [21], on the left there are two SMCs A and B which are composed, synchronising over actions a and d . Note that this is not a Boucherie model, but a model as it would typically be specified with an SPA. The resultant flat model is drawn in Fig. 4 on the right, where the three diagonal transitions represent the synchronisations. Assume that the property that needs to be verified




Fig. 4: Processes A and B and the resultant flat model

is $\Phi = Prob_{\geq b}(((t \wedge \neg p) \vee (r \wedge s)) U p)$. It is clear that state $(2, 1)$ is of type Impossible, states $(2, 2)$, $(3, 1)$ and $(3, 2)$ are Gobothways and states $(1, 1)$ and $(1, 2)$ are Certain, as colour-coded in the figure. If property Φ is violated for state $(2, 2)$, the flat model can be repaired using Tati's model repair algorithm. In this example, the heuristics from [23] suggests to reduce the rate of transition $(2, 2) \rightarrow (2, 1)$ by some reduction factor, which makes the desired transition $(2, 2) \rightarrow (1, 2)$ more likely. But now the problem is that this transition stems from model component B under action e . By changing the rate of action e in component B , apart from transition $(2, 2) \rightarrow (2, 1)$, transition rates $(1, 2) \rightarrow (1, 1)$ and $(3, 2) \rightarrow (3, 1)$ will change as well. While reducing the rate $(1, 2) \rightarrow (1, 1)$ is only an unnecessary side-effect, reducing $(3, 2) \rightarrow (3, 1)$ is a very negative side-effect, since this will make the undesired transition from $(3, 2)$ to $(2, 1)$ more likely, thereby lowering the satisfaction probability of state $(3, 2)$. Therefore, this example shows that the model repair solution obtained for the flat model cannot, in general, be lifted to the compositional high-level model.

4.1 Impossibility of compositional model repair even for product form

Having seen that even for a simple compositional model, compositional model checking and compositional model repair are not possible, we now address the interesting question whether models of Boucherie product form can be repaired in a compositional fashion. Unfortunately it turns out that this is not the case.

Theorem 4.2 *Compositional model repair (by transition rate modification) for global time-unbounded CSL Until formulas is in general not possible, even for Boucherie product form models.*

Proof. We prove Theorem 4.2 by counterexample. We consider the same Boucherie-type model as in Fig. 1, now with the following requirement:

$$\phi = P_{\geq b}(\varphi) = P_{\geq b}(q U ((p \wedge s) \vee (r \wedge t)))$$

Fig. 5 shows the state space where the colours now indicate the partitioning for this path formula φ . Model checking this requirement on the composed model yields the results in Table 4.




Fig. 5: Flat model with partitioned state space

(x, y)	$Pr((x, y), q U ((p \wedge s) \vee (r \wedge t)))$
(1, 2)	0
(1, 3)	0
(3, 1)	0
(3, 3)	0
(1, 1)	1
(3, 2)	1
(2, 1)	$\frac{\delta}{\beta+\delta}$
(2, 2)	$\frac{\beta}{\beta+\delta}$

Table 4: Satisfaction probabilities for the combined process

The only interesting states are (2, 1) and (2, 2), since the other states trivially satisfy or violate the requirement. Suppose that both $\frac{\delta}{\beta+\delta}$ and $\frac{\beta}{\beta+\delta}$ are less than the requested probability bound b , i.e. those two states violate the requirement. We therefore wish to repair the model by adjusting its rates. So we could increase the probability $Pr((2, 1), \varphi) = \frac{\delta}{\beta+\delta}$, by either increasing δ or decreasing β (or both at the same time). On the other hand, in order to repair state (2, 2) and increase the probability $Pr((2, 2), \varphi) = \frac{\beta}{\beta+\delta}$, we would have to decrease δ or increase β , which means that the repair strategies for these two states are exactly against each other! Looking at the flat model, we could solve this dilemma by distinguishing between the two δ -transitions (and between the two β -transitions), and adapt those rates individually. However, in the composed model, both δ -transitions (and both β -transitions) stem from the same transition in process C_1 , so we can only adjust them together in the same direction. So we see that compositional repair is not possible for this example. \square

4.2 Model Repair For Boucherie Framework

In the previous section, we have shown that, in general, compositional model repair is not possible, even for models of Boucherie product form. In this section, we provide a positive result, that under some conditions, compositional model repair of Boucherie-type models for specific CSL property types is possible.

Proposition 4.3 (a) *Given a Boucherie-type product form model with components C_1, \dots, C_K (and some passive resources which are redundant within the state vector) whose combined reachable state space is denoted as S . Consider an time-unbounded CSL Until requirement of the form*

$$\Phi = P_{\geq b}(\varphi) = P_{\geq b}((p_1 \wedge \dots \wedge p_K) U (q_1 \wedge \dots \wedge q_K))$$

where p_k and q_k are state formulas which refer only to component C_k , such that the following conditions hold:

- (i) $\forall k \in \{1, \dots, K\} : q_k \rightarrow p_k$
- (ii) For every reachable state $s = (s_1, \dots, s_K) \in S$ it is required that if $s_k \models q_k$ then C_k in state s_k does not possess a shared resource.
- (iii) For C_k , let the path formula φ_k be defined as $\varphi_k = p_k U q_k$. Each C_k , with respect to φ_k , should have states of type Gobothways (unless $q_k = tt$).

We define the following three types of repair measures (relating to φ_k):

- $Incr_k(m, n)$: in C_k , increase the rate from Gobothways-state m to Certain-state n by a factor of $i_k(m, n) \geq 1$.
- $Decr_k(m, n)$: in C_k , decrease the rate from Gobothways-state m to Impossible-state n by a factor of $0 < d_k(m, n) \leq 1$.
- $Hold_k(m)$: in C_k , for $m \in S_k$ such that $m \models q_k$, decrease the rate from m to any other state in S_k by a factor of $0 < h_k(m) \leq 1$.

Let $V = \{s \in S \mid 0 < Pr(s, \varphi) < b\}$, i.e. the subset of states which non-trivially violate Φ . Then one can perform compositional model repair for the states of V (i.e. one can increase the probabilities by which these states satisfy the Until formula φ to the desired probability bound b) by a combination of the repair measures $Incr_k(m, n)$, $Decr_k(m, n)$ and $Hold_k(m)$ ($k = 1, \dots, K$ and $m, n \in S_k$) with appropriately chosen factors $i_k(m, n)$, $d_k(m, n)$ and $h_k(m)$.

(b) For a Boucherie product form model that violates the non-blocking condition (ii) or a model that is not of product form, compositional model repair as in (a) is in general not possible.

Note 1 The cases $p_k = tt$ and $q_k = tt$ are covered by the Proposition 4.3. In particular, if $q_k = tt$, component C_k remains unchanged during model repair. Furthermore, the often used Eventually-operator ($ttU \dots$) is covered by setting all p_k to tt , which shows that the conditions imposed by the Proposition are not too restrictive for practical application.

Note 2 The Proposition (part (a)) provides an existence result, but it does not specify which combination of repair factors' values will be successful.

Note 3 If the requirement at hand is of the form $P_{<b}(\varphi)$, then during model repair one may wish to decrease the satisfaction probabilities (instead of increasing them). For this case, one can formulate a dual statement.

Proof. Part (a): In a Boucherie-type PF model, a component is either free to move as if it was fully independent, or it is completely blocked (i.e. cannot move at all). So the embedded Markov chain of component C_k (the probabilistic decision to which state the component should move next) remains the same, whether the component is in isolation or in the context of the other components and resources. For a state $s \in V$, since $Pr(s, \varphi) > 0$, there exist satisfying paths, but their combined probability mass is too low. The rate adjustments as suggested in the Proposition do not alter the set of satisfying paths (i.e. no new satisfying path is added and no satisfying path is removed), but they increase the probability mass of the set of satisfying paths. In particular, when adjusting the rates in the components by applying $Incr_k(m, n)$ and/or $Decr_k(m, n)$, the likelihood of C_k moving via p_k -states towards q_k -states is increased, which means that for the composed model the probability of the combination of those paths, which is simply an interleaving of the components' paths, is increased. (One can prove this last statement formally by an inductive argument, assuming that several repair measures of type $Incr$ and / or $Decr$ are applied one after the other, and showing that each one of them increases the probability mass of the desired set of paths.)

The $Hold_k(m)$ repair measure has the effect that a component C_k , once it has reached a q_k -state (where because of condition (i) p_k also still holds), will stay in this state longer. Thereby the probability is increased that, while C_k remains in this q_k -state, the other components $C_j \neq C_k$ will reach their q_j -states before component C_k moves on.

Condition (i), i.e. the fact that $q_k \rightarrow p_k$, means that it is not possible for a component C_k to move to one of its q_k -states too early (and thereby not satisfying the $(p_1 \wedge \dots \wedge p_K)$ -condition any more), since even if the other components $j \neq k$ have not yet reached their q_j -states, p_k is still satisfied in any q_k -state.

Condition (ii) is important, since it guarantees that if C_k is in state s_k that satisfies q_k then none of the other processes is blocked by C_k , i.e. the other processes $j \neq k$ can still move along the j -dimension of the state space towards their q_j -states.

A further important point in the proof of the Proposition concerns the special form of the path formula, which is



$$\varphi = ((p_1 \wedge \dots \wedge p_K) U (q_1 \wedge \dots \wedge q_K))$$

Let us, without loss of generality, assume that all states of C_k which satisfy p_k have adjacent indices and all states which satisfy q_k also have adjacent indices (this can always be achieved by a suitable state reordering). Then the special form of φ implies that the set of states $P = \{s \in S \mid s \models p_1 \wedge \dots \wedge p_K\}$ will form a K -dimensional rectangle in S (possibly with some holes, in case there are unreachable states due to resource conflicts). Likewise, the set of states $Q = \{s \in S \mid s \models q_1 \wedge \dots \wedge q_K\}$ is also a K -dimensional rectangle, which is completely contained in the first rectangle due to condition (i). An instance of this situation is depicted in Fig. 6 for the two-dimensional case. In the figure, only two possible repair measures as suggested by the Proposition are indicated. Model repair by the three repair measures specified in the Proposition then means that the probability for moving from the the set $P \setminus Q$ to Q , without visiting any other states in between, can be made sufficiently high.

Part (b): Example 4.1 is a non-Boucherie model where we saw that compositional model repair is not possible, which is a counterexample that proves the second statement of part (b). The first statement of part (b) (concerning the non-blocking condition (ii)) can also be shown by counterexample, but we do not provide such an example in this paper. \square

Example 4.4 a) Now we provide an example of how to use Proposition 4.3 for compositional model repair. In Fig. 7, C_1 and C_2 are two CTMCs forming a Boucherie product form model. Their states are labelled by atomic propositions from the set $AP = \{p_1, p_2, q_1, q_2\}$. All transition rates are assumed to have the value 1. There is one resource in this system and its use is highlighted in Fig. 7 by the bold circles (state 5 in C_1 and state 3 in C_2). The CSL property which needs to be checked is:

$$P_{\geq b}(\varphi) = P_{\geq b}((p_1 \wedge p_2) U (q_1 \wedge q_2))$$


 Fig. 6: The $(q_1 \wedge q_2)$ -rectangle is a subset of the $(p_1 \wedge p_2)$ -rectangle

 Fig. 7: Components C_1 and C_2



Thus, φ is a global time-unbounded Until formula of the type as in Proposition 4.3. Assume that the state of interest is $s_0 = (1, 1)$ and that the probability $Pr((1, 1), \varphi)$ is less than the desired probability bound b , so model repair is needed.

The flat model of this system is shown in Fig. 8. Condition (i) of Proposition 4.3 indicates that if there is a state which satisfies q_k , it must satisfy p_k as well and it is clear that this condition is satisfied in this example. In Fig. 8, the small rectangle which is the set of all $(q_1 \wedge q_2)$ -states is a subset of the big rectangle which is the set of all $(p_1 \wedge p_2)$ -states. This issue is a direct consequence of condition (i) in Proposition 4.3. We can extract φ_1 and φ_2 out of φ :

$$\varphi_1 = p_1 U q_1; \quad \varphi_2 = p_2 U q_2$$

According to φ_1 and considering C_1 in isolation (Fig. 9), states 1 and 4 are Gobothways states, state 5 is Impossible and states 2 and 3 are Certain. Similarly, considering φ_2 and C_2 in isolation, we see that in C_2 , state 3 is Impossible, state 2 is Certain and state 1 is Gobothways. Based on Proposition 4.3, we can obtain repair measures in order to use them in model repair. Table 5 shows the three types of repair measures in C_1 and C_2 as suggested in the Proposition, and Fig. 9 shows the partitioned states in C_1 and C_2 with only important rates for model repair.

According to Proposition 4.3, there always exists a solution for model repair using a combination of the repair measures specified in Table 5. By applying a combination of $Incr_i$, $Decr_i$ and $Hold_i$, we can increase the probability $Pr((1, 1), \varphi)$ to any desired level. In this example, it is not possible to increase the probability more than a specific limit by using only a single one of the repair measures. For instance, by forming and solving a set of equations, we can show that by increasing only rate b_1 and holding all


 Fig. 8: Flat combined model of C_1 and C_2

 Fig. 9: Partitioned state sets in C_1 and C_2

$Incr_1(1, 2)$	a_1
$Incr_1(4, 3)$	a_4
$Incr_2(1, 2)$	b_1
$Decr_1(4, 5)$	a_5
$Decr_2(1, 3)$	b_4
$Hold_1(2)$	a_3
$Hold_1(3)$	a_7
$Hold_2(2)$	b_2

Table 5: Repair measures in Example 4.4

other rates equal to one, it is only possible to reach $Pr((1, 1), \varphi) = 0.75$:

$$Pr((1, 1), \varphi) = \frac{b_1(3b_1^2 + 28b_1 + 66)}{4b_1^3 + 44b_1^2 + 146b_1 + 126}$$

However, in order to achieve a higher satisfaction probability, we can use combinations of repair measures. One possible solution is to increase rates a_1 and b_1 simultaneously. Fig. 10 shows the satisfaction probability depending on transition rates a_1 and b_1 (all other rates are fixed at the value 1). The black plane marks the original satisfaction probability (when also both a_1 and b_1 are equal 1), which is $97/320$. We have also verified that the satisfaction probability can be increased even further than shown in the plot, i.e. any given probability bound $b < 1$ can be topped.

Note 4 Proposition 4.3 states that we can always find a successful combination of the specified three types of repair measures. However, this does not exclude the existence of other rate adaptations which also might have a positive effect on the satisfaction probabilities. For instance, in Example 4.4 it is possible to increase the probability of φ by decreasing transition rate a_2 in C_1 (which is a transition within the set of p_1 -states).

5 Conclusion

In this paper, we have studied compositional model checking and model repair for CSL global time-unbounded Until formulas for Boucherie product form models, where we consider model repair by rate adaptation. Using




Fig. 10: Satisfaction probability $Pr((1, 1), (p_1 \wedge p_2)U(q_1 \wedge q_2))$ depending on transition rates a_1 and b_1

counterexamples, it is proved that, in general, compositional model checking and model repair is not possible for such CSL formulas, even for product form models. We have also proved a proposition which – in the Boucherie framework – makes compositional model repair possible for a specific, but still quite general type of CSL Until formula. Although not every time-unbounded Until formula can be written in this form, it is very useful since such formulas often appear in practice. In particular, the widely used Eventually operator is an instance of the case covered by the proposition.

As future work, we intend to generalise Proposition 4.3, since we think that condition (ii) could be omitted if the $Hold_k(m)$ repair measure was refined to $Hold_k(m, n)$ (which means that the reduction factor could be different, depending on the target state of the particular transition to be reduced). This generalisation, however, would make the proof more difficult. Furthermore, since the proposition provides only an existential result, we intend to develop from it a constructive algorithm which identifies successful combinations of repair measures and determines concrete rate adaptation factors.

References

- [1] C. Baier, B. Haverkort, H. Hermanns, and J. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(6):524–541, June 2003.
- [2] P. Ballarini and A. Horvath. Compositional Model Checking of product-form CTMCs. *Electronic Notes in Theoretical Computer Science*, 250(1):21 – 37, 2009. Proceedings of the Seventh International Workshop on Automated Verification of Critical Systems (AVoCS 2007).
- [3] R. J. Boucherie. A characterization of independence for competing Markov chains with applications to stochastic Petri nets. *IEEE Transactions on Software Engineering*, 20(7):536–544, July 1994.
- [4] P. Buchholz. Structured analysis approaches for large Markov chains. *Applied Numerical Mathematics*, 31(4):375 – 404, 1999.
- [5] P. Buchholz and J. Kriege. Approximate aggregation of Markovian models using alternating least squares. *Perform. Eval.*, 73:73–90, 2014.
- [6] S. Donatelli and G. Franceschinis. The PSR methodology: Integrating hardware and software models. In *Application and Theory of Petri Nets 1996*, pages 133–152. Springer Berlin Heidelberg, 1996.
- [7] S. Donatelli and M. Sereno. On the product form solution for Stochastic Petri Nets. In K. Jensen, editor, *Application and Theory of Petri Nets 1992*, pages 154–172, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg.
- [8] J. Fourneau. Collaboration of discrete-time Markov chains: Tensor and product form. *Performance Evaluation*, 67:779796, 09 2010.
- [9] J. Fourneau, B. Plateau, and W. Stewart. An algebraic condition for product form in stochastic automata networks without synchronizations. *Perform. Eval.*, 65:854–868, 11 2008.
- [10] A. Gouberman, M. Siegle, and B. Tati. Markov chains with perturbed rates to absorption: Theory and application to model repair. *Performance Evaluation*, 130:32 – 50, 2019.
- [11] H. Hermanns, U. Herzog, V. Mertsiotakis, and M. Rettetbach. Exploiting stochastic process algebra achievements for generalized stochastic Petri nets. *Proceedings of the Seventh International Workshop on Petri Nets and Performance Models*, pages 183–192, 1997.

- [12] H. Hermanns, M. Kwiatkowska, G. Norman, D. Parker, and M. Siegle. On the use of MTBDDs for performability analysis and verification of stochastic systems. *The Journal of Logic and Algebraic Programming*, 56(1):23 – 67, 2003. Probabilistic Techniques for the Design and Analysis of Systems.
- [13] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, New York, NY, USA, 1996.
- [14] J. Hillston and N. Thomas. Product Form Solution for a Class of PEPA Models. *Perform. Eval.*, 35(3-4):171–192, May 1999.
- [15] J. R. Jackson. Jobshop-Like Queueing Systems. *Manage. Sci.*, 50(12 Supplement):1796–1802, December 2004.
- [16] J.-P. Katoen, D. Klink, and M. R. Neuhäuser. Compositional Abstraction for Stochastic Systems. In Ouaknine J. and F. Vaandrager, editors, *Formal Modeling and Analysis of Timed Systems*, pages 195–211. Springer Berlin Heidelberg, 2009.
- [17] M. Kuntz, M. Siegle, and E. Werner. Symbolic Performance and Dependability Evaluation with the Tool CASPA. In M. Nunez, Z. Maamar, and F.L. Pelayo, editors, *Applying Formal Methods: Testing, Performance and M/E Commerce: FORTE 2004 Workshops, European Performance Engineering Workshop*, pages 293–307. Springer, LNCS 3236, 2004.
- [18] D. Milios and S. Gilmore. Compositional Approximate Markov Chain Aggregation for PEPA Models. In M. Tribastone and S. Gilmore, editors, *Computer Performance Engineering*, pages 96–110, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [19] W. H. Sanders and J. F. Meyer. Stochastic Activity Networks: Formal Definitions and Concepts. In E. Brinksma, H. Hermanns, and J. Katoen, editors, *Lectures on Formal Methods and Performance Analysis: First EEF/Euro Summer School on Trends in Computer Science, Revised Lectures*, pages 315–343. Springer Berlin Heidelberg, 2001.
- [20] M. Sereno. Towards a Product Form Solution for Stochastic Process Algebras. *The Computer Journal*, 38(7):622–632, 01 1995.
- [21] B. Tati. *Quantitative Model Repair of Stochastic Systems*. PhD thesis, Bundeswehr University Munich, Department of Computer Science, 2018.
- [22] B. Tati and M. Siegle. Parameter and Controller Synthesis for Markov Chains with Actions and State Labels. In É. André and G. Frehse, editors, *2nd International Workshop on Synthesis of Complex Parameters (SynCoP’15)*, volume 44 of *OpenAccess Series in Informatics (OASICS)*, pages 63–76, Dagstuhl, Germany, 2015.
- [23] B. Tati and M. Siegle. Rate Reduction for State-labelled Markov Chains with Upper Time-bounded CSL Requirements. In *Proceedings Casting Workshop on Games for the Synthesis of Complex Systems and 3rd International Workshop on Synthesis of Complex Parameters, Casting/SynCoP 2016, Eindhoven, The Netherlands, April 2-3, 2016.*, pages 77–89, 2016.
- [24] N. Thomas and P. G. Harrison. Semi-Product-Form Solution for PEPA Models with Functional Rates. In A. Dudin and K. De Turck, editors, *Analytical and Stochastic Modeling Techniques and Applications*, pages 416–430, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [25] M. Wan, G. Ciardo, and A. S. Miner. Approximate Steady-state Analysis of Large Markov Models Based on the Structure of their Decision Diagram Encoding. *Perform. Eval.*, 68(5):463–486, May 2011.