Revised Version: 15.06.07
Reihe Informatik. TR-2007-001

Liveness in Interaction Systems

Moritz Martens*

June 15, 2007

Abstract

Interaction systems are a formal model for component-based systems. It has been shown that in this formalism the question whether a component is live is NP-hard. Therefore it is desirable to find sufficient criteria that can be checked in polynomial time. In this report we present and discuss two criteria for liveness. We also establish a new characterization of liveness.

1 Introduction

We consider interaction systems, a model where components are combined via connectors to form more complex systems [GS02, GS03b, GS05, BBS06, Sif04, Sif05, GS03a]. Each single component i offers ports $a_i, b_i, \ldots \in A_i$ for cooperation with other components. Each port in A_i represents an action of component i. The behavior of a component is represented by a labeled transition system. Components are glued together via connectors, where each connector connects certain ports. In the global system obtained by gluing components together deadlocks may arise where groups of components are waiting for each other cyclically and will thus no longer participate in the progress of the global system (cf. [Tan01]). If a system is deadlock-free it is always able to proceed. Then one can ask the question whether a subset of components K' is live, i.e. in every infinite sequence of transitions there are infinitely many interactions that let a component from K' participate. In [MMMC06] it has been shown that deciding liveness is NP-hard. Here we present and discuss two criteria that ensure liveness and can be tested in polynomial time. In addition we give a new characterization of liveness.

The report is organized as follows. Section 2 contains the basic definitions and Section 3 contains the definitions concerning deadlock-freedom and liveness. Sections 4 and 5 constitute the main part of the report where the criteria respectively the characterization are presented and the various proofs are given.

*E-mail: mmartens@informatik.uni-mannheim.de

2 Components, Connectors and Interaction Systems

We consider *interaction systems*, a model for component-based systems that was proposed and discussed in detail in [GS02, GS03b, GS05, BBS06, Sif04, Sif05, GS03a, GGMC$^+$07b, GGMC$^+$07a, MCMM07].

In this framework we consider a set of components K where we usually refer to a component as $i \in K$. For every component $i \in K$ a set A_i of actions or ports is specified which the component can use to cooperate with other components. This cooperation is determined by so-called connectors. A connector is a finite nonempty set of ports that contains at most one port for every component in K. Any nonempty subset of a connector constitutes an interaction of the system. Any interaction models a step of the system where the ports contained in that interaction are performed simultaneously.

Definition 1. *A* component system $CS = \left(K, \{A_i\}_{i\in K}\right)$ *is a pair where* K *is the set of* components, A_i *is the* port set *of component* i*, and any two port sets are disjoint. Ports are also referred to as* actions.

The union $A = \bigcup_{i\in K} A_i$ *of all port sets is the port set of* K*. A finite nonempty subset* c *of* A *is called a* connector *or* maximal interaction *for* CS*, if it contains at most one port of each component* $i \in K$ *that is* $|c \cap A_i| \le 1$ *for all* $i \in K$*. A* connector set *is a set* C *of connectors for* CS *that covers all ports and contains only maximal elements:*

1. $\bigcup_{c\in C} c = A$
2. $c \subseteq c' \Rightarrow c = c'$ *for all* $c, c' \in C$.

If c *is a connector,* $I(c)$ *denotes the set of all nonempty subsets of* c *and is called the* set of interactions of c*. For a set* C *of connectors*

$$I(C) = \bigcup_{c\in C} I(c)$$

is the set of interactions of C*.*

For component i *and interaction* α*, we put* $i(\alpha) = A_i \cap \alpha$*. We say that component* i participates *in* α*, if* $i(\alpha) \neq \emptyset$*.*

We give a small example to illustrate these concepts. We will extend this example throughout the report whenever we encounter new notions.

Example 1. *We consider a component system* $CS_5 = (K_5, \{A_i\}_{i\in K_5})$ *consisting of five components, where* $K_5 := \{1, 2, 3, 4, 5\}$ *and the port sets of the components are given by* $A_1 := \{a_1\}$*,* $A_2 := \{b_1, b_2\}$*,* $A_3 := \{d_1, d_2\}$*,* $A_4 := \{e_1, e_2\}$*, and* $A_5 := \{f_1, f_2\}$*. In addition we fix a connector set as follows:* $C_5 := \{\{a_1, b_1\}, \{a_1, e_1\}, \{f_1, d_1\}, \{e_2, f_2, a_1\}, \{e_2, f_2, b_2\}, \{e_2, d_2, b_2\}, \{e_2, d_2, a_1\}\}$.

For example components 1 *and* 2 *may perform their respective first actions together whereas* 4*,* 5*, and* 2 *may perform their respective second actions together.*

In the following, we always assume that $K = \{1, \ldots, n\}$ for some $n \in \mathbb{N}$ or that K is countably infinite.

An *interaction model* for a component system CS is defined by a connector set C together with an arbitrary set $Comp$ of interactions that are declared to be complete. If an interaction is declared complete it can be performed independently of the environment. In particular if we want an action of a single component to be independent of the environment it should be declared to be complete. Note that it is a design decision which interactions are chosen to be complete. This choice is not restricted in any way and only depends on the system one wishes to model[1].

Definition 2. *Let C be a connector set for the component system CS and let $Comp \subseteq I(C)$ be an arbitrary subset of interactions.*

$$IM := (C, Comp)$$

is an interaction model *for CS. The elements of $Comp$ are called* complete interactions*.*

Example (Example 1 continued)**.** *In the above example we do not want any interaction to be available independently of the connectors that have been specified. Therefore we choose $Comp_5$ to be the empty set.*

If for some reason the interactions $\{e_2\}$ and $\{f_2, b_2\}$ for example should be independent of other actions, we could set $Comp := \{\{e_2\}, \{f_2, b_2\}\}$.

The notions presented so far are only concerned with the possible structure of communication between the different components. We provide a further level of description of the components that restricts the order in which a component may perform the actions it provides. For every component $i \in K$ a labeled transition system T_i describing the behavior of that component is introduced.

Definition 3. *Let $CS = \left(K, \{A_i\}_{i\in K}\right)$ be a component system and $IM = (C, Comp)$ an interaction model for CS. Let for each component $i \in K$ a transition system $T_i = \left(Q_i, A_i, \to_i, Q_i^0\right)$ be given where $\to_i \subseteq Q_i \times A_i \times Q_i$ and $Q_i^0 \subseteq Q_i$ is a non-empty set of* initial states[2]*. We write $q_i \xrightarrow{a_i}_i q_i'$ instead of $(q_i, a_i, q_i') \in \to_i$.*

The induced interaction system *is given by*

$$Sys := (CS, IM, T)$$

where the global behavior $T = \left(Q, I(C), \to, Q^0\right)$ *is obtained from the local transition systems of the individual components in a straightforward manner:*

1. *$Q := \prod_{i\in K} Q_i$, the Cartesian product of the Q_i which we consider to be order independent. We denote states by tuples $q := (q_1, \ldots, q_j, \ldots)$ and call them* (global) states*.*

2. *$Q^0 := \prod_{i\in K} Q_i^0$, the Cartesian product of the local initial states. We call the elements of Q^0 (global)* initial states*.*

[1] A slightly more restrictive definition of *Comp* has been introduced in [GS05]. There it is required that every superset in $I(C)$ of a complete interaction should also be complete. Formally this is realized by introducing a certain notion of closure of a set of sets and requiring that the set of complete interactions is closed in this sense. For our results this requirement is not relevant.

[2] There are versions of the framework that consider systems without designated initial states and allow the system to be initialized in any global state (cf. [GGMC+07b]). This point of view is a special case of our definition where $Q_i^0 = Q_i$ for all i.

3. $\rightarrow \subseteq Q \times I(C) \times Q$*, the transition relation for Sys defined by*

$$\forall \alpha \in I(C)\, \forall q, q' \in Q : q = (q_1, \ldots, q_j, \ldots) \xrightarrow{\alpha} q' = \left(q'_1, \ldots, q'_j, \ldots\right) \Leftrightarrow$$

$$\forall i \in K : q_i \xrightarrow{i(\alpha)}_i q'_i \text{ if } i \text{ participates in } \alpha \text{ and } q'_i = q_i \text{ otherwise.}$$

A state $q_i \in Q_i$, respectively a global state $q \in Q$ is called complete *if there is some interaction $\alpha \in C \cup Comp$ and some q'_i such that $q_i \xrightarrow{\alpha}_i q'_i$, respectively some q' such that $q \xrightarrow{\alpha} q'$. Otherwise it is called* incomplete.

Note that a global state q is complete if q_i is complete for some i. But q may still be complete even if all q_i are incomplete.

Example (Example 1 continued)**.** *The behavior of component i is given in Fig. 1 for $i \in \{1, \ldots, 5\}$. For every component i we put $Q_i^0 = Q_i$. The induced tran-*

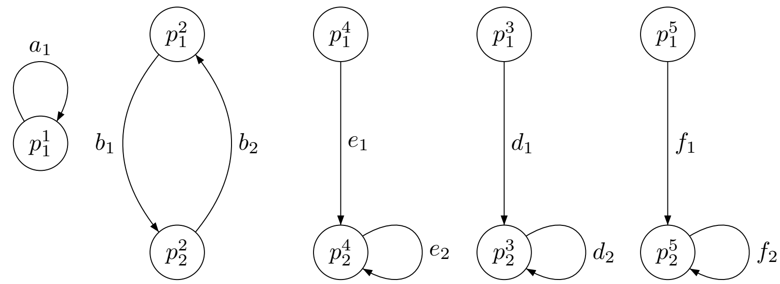

Figure 1: The local behavior of the components of Example 1

sition system is called $T(5)$. For example in the global state $\left(p_1^1, p_1^2, p_1^3, p_1^4, p_1^5\right)$ a transition labeled with $\{d_1, f_1\}$ is enabled. Our example system $Sys_5 := (CS_5, IM_5, T(5))$ is now completely specified. Note that no local state of any component is complete because there is no connector of length one and no complete interaction at all. As we will see in the next section all global states are complete.

Remark 1. *In what follows, we often mention $Sys = (CS, IM, T)$. It is understood that $CS = \left(K, \{A_i\}_{i \in K}\right)$, $IM = (C, Comp)$, $T_i = \left(Q_i, A_i, \rightarrow_i, Q_i^0\right)$ for $i \in K$, and T are given as above. Usually we will display the local transition systems graphically. If not explicitly stated otherwise the local initial states will be marked by an ingoing arrow.*

3 Properties of Interaction Systems

In this section we will define the property of liveness of a component. In order to define liveness we need the notion of deadlock-freedom first. Deadlock-freedom in interaction systems has been thoroughly studied in diverse works. We refer the reader to [GGMC+07b, MCMM07].

For a system under consideration it is desirable that a situation where all components need some other component to proceed which in turn does not offer the action needed never occurs. Such a situation would result in a setting where groups of components are waiting for each other cyclically such that no interaction will ever be performed again. This kind of event is called a global deadlock of the system and as mentioned above we want to avoid such deadlocks. The references mentioned above give sufficient criteria for deadlock-freedom of an interaction system that (in some cases) can be tested in polynomial time. It has been shown in [Min06] that deciding deadlock-freedom in component systems is NP-hard which justifies the search for such criteria.

From now on we will assume that the local transition systems have the property that every state offers at least one action. This means that a deadlock can only be caused by cyclic waiting conditions as above and not because of the nonexistence of actions in the states at hand. This is not a strong restriction as the general case can be reduced to this case by introducing idle actions.

Definition 4. *Let Sys be an interaction system.*

1. *Let $q \in Q$. q is* reachable *in Sys if there is a sequence $q^0 \overset{\alpha_0}{\to} q^1 \overset{\alpha_1}{\to} \dots \overset{\alpha_{n-1}}{\to} q$ such that $q^0 \in Q^0$ and $\alpha_i \in C \cup Comp$ for all $0 \le i \le n-1$.*

2. *Sys is called* deadlock-free *if for every reachable state q there exists $\alpha \in C \cup Comp$ and $q' \in Q$ such that $q \overset{\alpha}{\to} q'$.*

A deadlock-free system may always proceed with some maximal or complete interaction. Deadlock-freedom of a system is equivalent to the fact that every global state of the system is complete.

Example (Example 1 continued)**.** *In Sys_5 every state is reachable because the system may be initialized in any state. Sys_5 is deadlock-free. This can be seen by distinguishing several cases. If components 2 or 4 are in their first state it is always possible to perform the connector $\{a_1, b_1\}$ or $\{a_1, e_1\}$. Therefore it suffices to consider the global states in which these two components are in their second state. Then if at least one of components 3 or 5 is in the second state one of the four connectors containing three elements is possible. If this is not the case 3 and 5 are in their first state and $\{f_1, d_1\}$ is possible. Therefore every global state offers some maximal interaction and there is no deadlock.*

The definition of deadlock-freedom leads to the notion of a run which simply is an infinite thread of execution of the system.

Definition 5. *Let Sys be a deadlock-free interaction system and let $q \in Q$ be a reachable state. A* run *of Sys is an infinite sequence*

$$\sigma = q \overset{\alpha_0}{\to} q^1 \overset{\alpha_1}{\to} q^2 \dots$$

with $q^l \in Q$ and $\alpha_l \in C \cup Comp$ for all $l \in \mathbb{N}$.

Let $i \in K$ be a component and let σ be a run of Sys. If there exists l such that i participates in α_l we say that i participates in σ.

Now we can say when a set of components is live. Basically a component is live if for any point of time no matter how the system behaves the component will eventually participate in some interaction which means that the component participates infinitely many often in every run of the system. From now on we identify singleton sets with their element if it is convenient to do so.

Definition 6. *Let Sys be a deadlock-free interaction system and let $K' \subseteq K$ be a nonempty set of components. We say that K' is* live *in Sys if for every run σ of Sys there exists a natural number $n \in \mathbb{N}$ such that there is some $i \in K'$ with $i(\alpha_n) \neq \emptyset$.*

We say that K' is strongly live *in Sys if every component $i \in K'$ is live in Sys.*

For a single component the notions of liveness and strong liveness coincide. Moreover if i is live in Sys then any set of components containing i is also live whereas the converse does not hold: even if K' is live in Sys there does not need to be any $i \in K'$ that is live. Finally note that for $K' = K$ liveness follows from deadlock-freedom.

Definition 3 introduced transitions for every $\alpha \in I(C)$ even though for the notions presented above only the maximal and complete interactions are relevant. The extra information about the interactions in $I(C)$ is not needed in this work and it could be omitted. The reason we include those transitions into the global transition system is because they are needed to define a composition operator that allows to build complex systems from smaller subsystems. Definitions and results concerning this operator are presented in [GGMC+07b]. We want to stick with this general definition of an interaction system even though we do not need all the information.

Example (Example 1 continued)**.** *Component 1 is live in our example-system. We will consider the different global states and argue why for every state only finitely many steps can be performed before 1 participates in an interaction. If 2 is in its first state the only connector that does not involve 1 and might be available is $\{f_1, d_1\}$. After the execution of this connector an interaction involving 1 must be performed. If 2 is not in its first state it is possible to perform an interaction not involving 1 but this interaction (unless it is $\{f_1, d_1\}$ which can only be performed once) will force 2 to move to the first state which is the case covered by the argumentation above.*

4 Criteria for Liveness

In [MMMC06] we showed that deciding liveness in interaction systems is NP-hard. This motivates the search for sufficient criteria that can be checked in polynomial time. We present and discuss two criteria that can both be tested in polynomial time.

In this section we always assume that Sys is a deadlock-free interaction system with a finite set of components K and finite port sets A_i.

Definition 7. *Let Sys be an interaction system as above and let $j \in K$ be an arbitrary component.*

1. *Let $A'_j \subseteq A_j$ be a subset of actions of j. A'_j is* inevitable *in T_j if only finitely many transitions labeled with $a_j \in A_j \backslash A'_j$ can be performed in T_j before some action from A'_j must be performed.*

2. *Let $\Lambda \subseteq I(C)$ be an arbitrary nonempty set of interactions and let $j \in K$ be a component. We define*

$$\Lambda[j] := A_j \cap \bigcup_{\alpha \in \Lambda} \alpha$$

the set of ports of j that participate in one of the interactions of Λ.

We will now give the definitions that are needed for the graph.

Definition 8. *Let $\tilde{K} \subseteq K$ be an arbitrary subset of components. Let*

$$excl\left(\tilde{K}\right) := \left\{\alpha \in C \cup Comp | \forall i \in \tilde{K} : i\left(\alpha\right) = \emptyset\right\}$$

denote the set of maximal or complete interactions that do not allow any component from $\tilde{K}$ to participate.

We define the graph G_{live} as follows.

Definition 9. *Set*

$$G_{live} := (K, E_0)$$

where we have $(i,j) \in E_0$ if and only if $A_j \backslash excl\left(i\right)[j]$ is inevitable in T_j.

Informally an edge from i to j in the graph has the meaning that j can only proceed finitely many times before i also has to participate in a (global) step.

Definition 10. *Let Sys be an interaction system as above and let $i \in K$ be a component.*

Let $Reach^0\left(i\right) := \{j \in K | j \text{ is reachable from } i \text{ in } G_{live}\}$ denote the set of components that can be reached from i in G_{live}.

We inductively define the following subsets of K.

$$R_0\left(i\right) := Reach^0\left(i\right)$$

and

$$R_{n+1}\left(i\right) := R_n\left(i\right) \cup$$

$$\{j \in K \backslash R_n\left(i\right) | \forall \alpha \in C \cup Comp\left(j\left(\alpha\right) \neq \emptyset \Rightarrow \exists k \in R_n\left(i\right) : k\left(\alpha\right) \neq \emptyset\right)\}.$$

The first condition is given in the following statement.

Proposition 1. *Let Sys be a deadlock-free interaction system and let $k \in K$. If $K = \bigcup_{n \geq 0} R_n\left(k\right)$ then k is live in Sys.*

For the proof we need the following auxiliary lemma:

Lemma 1. *Let $\sigma = q_0 \stackrel{\alpha_0}{\to} q_1 \stackrel{\alpha_1}{\to} q_2 \ldots$ be a run. If there is a path $k_0 \to k_1 \to \ldots \to k_l$ in G_{live} and k_l participates infinitely often in σ then k_0 participates infinitely often in σ.*

Proof. We will perform an induction on the length l of the path.
Start of induction: $l = 1$. Then there is an edge $k_0 \to k_1$. As k_1 participates infinitely often in transitions of σ and as the set of actions of k_1 that need cooperation of k_0 is inevitable in T_{k_1} we conclude that k_0 participates infinitely often in transitions of σ.
Induction step: $l \to l+1$. Let $k_0 \to k_1 \to \ldots \to k_l \to k_{l+1}$ be a path of length $l+1$ and let k_{l+1} participate infinitely often in σ then by induction assumption k_1 participates infinitely often in σ and as above we conclude that k_0 participates infinitely often. □

Now we can give the proof of Proposition 1:

Proof. Let $\sigma = q_0 \stackrel{\alpha_0}{\to} q_1 \stackrel{\alpha_1}{\to} q_2 \ldots$ be a run. We have to show that σ encompasses an infinite number of transitions where k participates. As K is finite and σ infinite there must be some component $\hat{k}$ that participates in infinitely many transitions of σ.

1. $\hat{k} = k$, then we are done.
2. $\hat{k} \neq k$ then we now that $\hat{k} \in \bigcup R_i(k)$.

   case 1: if $\hat{k} \in R_0(k)$ then by the above lemma and the definition of $R_0(k)$ we conclude that k participates infinitely often in σ.

   case 2: let $\hat{k} \in R_i(k)$ for some $i > 0$. Then we show by induction on i that k participates infinitely often in σ.

   Start of induction $i = 1$: if $\hat{k} \in R_1(k)$ then for all $\alpha \in C \cup \mathit{Comp}$ with $\hat{k}(\alpha) \neq \emptyset$ $\exists j \in R_0(k)$ with $j(\alpha) \neq \emptyset$. As $\hat{k}$ participates infinitely often in σ and as there are only finitely many elements in $C \cup \mathit{Comp}$ there must be some α with $\hat{k}(\alpha) \neq \emptyset$ which occurs infinitely often in σ. By definition of $R_1(k)$ $\exists j \in R_0(k)$ with $j(\alpha) \neq \emptyset$. Hence j participates infinitely often in σ. As $j \in R_0(k)$ case 1 above implies that k participates infinitely often in σ.

   Induction step $i \to i + 1$: let $\hat{k} \in R_{i+1}(k)$. As before there is an $\alpha \in C \cup \mathit{Comp}$ with $\hat{k}(\alpha) \neq \emptyset$ and α occurs infinitely often in σ. Some $j \in R_i(k)$ participates in this α, hence j participates infinitely often in σ and by induction assumption k participates infinitely often in σ.

□

Corollary 1. *Let Sys be a deadlock-free interaction system and let $i \in K$. If $K = \mathit{Reach}^0(i)$ then i is live in Sys.*

Definition 11. *Let E_0 be defined as above and define E_{n+1} inductively as follows.*

$$E_{n+1} := \{(i,j) \,|\, A_j \backslash \mathit{excl}\,(\mathit{Reach}^n\,(i))\,[j] \text{ is inevitable in } T_j\}$$

where $\mathit{Reach}^n(i) := \{j | j \text{ is reachable from } i \text{ in } (K, \bigcup_{m=0}^{n} E_m)\}$.
Define $E := \bigcup_{m=0}^{\infty} E_m$ *and* $G := (K, E)$.

The second criterion is as follows.

Proposition 2. *Let $K' \subseteq K$ be a set of components. If all components in $K \backslash K'$ are reachable from K' in G then K' is live in Sys.*

Proof. First we prove the following two facts by induction over $l \in \mathbb{N}$.

1. $(i, j) \in E_l$ implies that j can only participate finitely many times in any run σ of Sys before i has to participate.
2. $j \in \mathit{Reach}^l(i)$ implies that j can only participate finitely many times in any run σ of Sys before i has to participate.

For $l = 0$ both statements follow from Corollary 1.

Therefore let both statements be true for l and consider $(i, j) \in E_{l+1}$ as well as a run σ. Assume that j participates infinitely many often in σ. From the definition of E_{l+1} we know that $A_j \backslash excl\left(Reach^l(i)\right)[j]$ is inevitable in T_j. Because j participates infinitely many often in σ this means that j has to perform infinitely many often some action from $A_j \backslash excl\left(Reach^l(i)\right)[j]$. But $excl\left(Reach^l(i)\right)[j]$ is the set of actions of j that occur in some connector not involving any component from $Reach^l(i)$. Therefore $A_j \backslash excl\left(Reach^l(i)\right)[j]$ is the set of actions of j that only occur in connectors also involving components from $Reach^l(i)$. Because K is finite this means that there must be some $j \in Reach^l(i)$ that participates infinitely many often in σ. From the induction hypothesis and $j \in Reach^l(i)$ we conclude that i participates in σ.

Next we consider i and j such that $j \in Reach^{l+1}(i)$. We show that the second statement is true by induction over the length of a path visiting only edges from $\bigcup\limits_{m=0}^{l+1} E_m$. If $i \to j$ is such a path of length one the claim follows from the first part of the proof. Now let $p = i \to \ldots \to k \to j$ be a path of length $s+1$ that only visits edges from $\bigcup\limits_{m=0}^{l+1} E_m$ and let σ be a run of Sys. If j participates infinitely many often we conclude that k participates infinitely many often also because $k \to j \in \bigcup\limits_{m=0}^{l+1} \to_m$ and because of the first part of the proof. k is reachable from i over a path of length s. Therefore by induction we conclude that i has to participate in σ.

The proof of the proposition is straightforward now. Let j be reachable from i in G over a path p. This path visits only finitely many edges which means that there exists $n_0 \in \mathbb{N}$ such that all edges along p lie in $\bigcup\limits_{m=0}^{n_0} E_m$. The second fact of the proof above implies that for any run σ the component j can only participate finitely many times before i also has to participate.

Then it is clear that K' is live in Sys if all components in $K \backslash K'$ are reachable from K' in G. Indeed, if $K' = K$ liveness follows from deadlock-freedom. Otherwise for any run there must be some component j that participates infinitely many often because K is finite. j is reachable from some component in K' and the above argument yields that K' participates. $\square$

It is not hard to see that every stage of the construction of the edges causes cost polynomial in $|K|$, $|C \cup Comp|$ and the sum of the sizes of the local transition systems. Once no new edges are added the construction can be stopped. Note that this is the case after at most $|K|^2$ stages because this is the maximal number of edges that G can have. This argument shows that the criterion can indeed be tested in polynomial time.

The second criterion covers a larger class of interaction systems than the first one.

Proposition 3. *Let Sys be an interaction system as above and let $i \in K$ be live in Sys.*

An interaction system that satisfies the conditions of Proposition 1 also satisfies the condition of Proposition 2 but not viceversa.

Proof. First we will show that $j \in R_n(i)$ implies $j \in Reach^n(i)$ by induction on $n \geq 0$.

For $j \in R_0(i)$ nothing has to be done because $R_0(i) = Reach^0(i)$

Let j be in $R_{n+1}(i)$. Then we know by definition of $R_n(i)$ that whenever j participates in a maximal or complete interaction α there is some $k \in R_n(i)$ that also participates in α. From the induction hypothesis we conclude that each such k is in $Reach^n(i)$. This means that whenever j participates in some $\alpha \in C \cup Comp$ some component from $Reach^n(i)$ also participates and therefore $excl(Reach^n(i))[j] = \emptyset$. Then $A_j \backslash excl(Reach^n(i))[j] = A_j$ and it is clear that this set of actions is inevitable in T_j. Therefore (i, j) is added to E_{n+1} and $j \in Reach^n(i)$.

Now it is clear that $K = \bigcup_{n \geq 0} R_n(i)$ implies that $K \backslash \{i\}$ is reachable from i in G. For the remainder of the proof we refer to the following example.

□

Example (Example 1 continued)**.** *We have already explained why the example introduced above is deadlock-free and why component* 1 *is live in this system. Now we will argue that Proposition 1 cannot be used to show this whereas Proposition 2 is sufficient.*

G_{live} *for this system is given by Figure 2 where all edges are in* E_0*.*

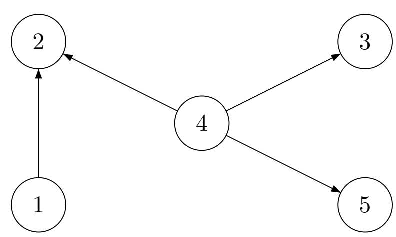

Figure 2: G_{live} for Example 1

Components 3*,* 4 *and* 5 *are not reachable from* 1 *therefore Corollary 1 cannot be applied. But it is also not possible to fall back to Proposition 1 to prove liveness of* 1*. We have* $Reach^0(1) = R_0(1) = \{1, 2\}$*. Computing* $R_1(1)$ *we additionally get* $\{4\}$ *because every connector involving* 4 *also involves some action of either* 1 *or* 2*. For all other* n *we get* $R_n(1) \cap \{3, 5\} = \emptyset$ *because the connector* $\{g_1, e_1\}$ *does not involve any component from* $Reach(1) \cup R_1(1)$*. Therefore the condition is not fulfilled.*

Proposition 2 can be used to prove liveness of 1*.* $(1, 4)$ *will be added to* E_1*. Then every component is reachable from* 1 *in* $\left(K, \bigcup_{m=0}^{1} E_m\right)$ *and therefore also in* G *and liveness of* 1 *follows from the proposition.*

5 Characterizing Liveness of a Set of Components

In this section we consider a not necessarily finite deadlock-free interaction system. We present the characterization of all subsets $K' \subseteq K$ that are live.

Definition 12. *Let Sys be a deadlock-free interaction system and let $K' \subseteq K$ be a non-empty subset of components. We define*

$$\bar{K}' := \{k \in K | \exists \alpha \in excl\,(K') : k\,(\alpha) \neq \emptyset\}\,.$$

Further we define the following labeled transition system

$$\bar{T} := \left(\bar{Q}, excl\,(K')\,, \rightarrow\right)$$

where $\bar{Q} := \prod_{k \in \bar{K}'} Q_k$ and $\rightarrow \subseteq \bar{Q} \times excl\,(K') \times \bar{Q}$ is the transition relation which is defined as follows for any two states $\bar{p}, \bar{q} \in \bar{Q}$ and any interaction $\alpha \in excl\,(K')$:

$$\bar{p} \xrightarrow{\alpha} \bar{q} \Leftrightarrow \exists p, q \in Q : \left(p \xrightarrow{\alpha} q \wedge \forall i \in \bar{K}' : (p_i = \bar{p}_i \wedge q_i = \bar{q}_i)\right)$$

Proposition 4. *Let Sys be deadlock-free and let $K' \subseteq K$.*

K' is live in Sys if and only if $\bar{T}$ neither contains any cycle visiting a state $\bar{q}$ for which there exists $q' \in \prod_{i \in K \setminus \bar{K}'} Q_i$ such that $(\bar{q}, q')$ is reachable in Sys nor any infinite path starting in such a state.

Proof. If $\bar{T}$ contains a cycle visiting a state $\bar{q}$ as above this cycle also exists in the global system. This is because all interactions used to label the transitions of $\bar{T}$ are in $excl\,(K') \subseteq C \cup Comp$.

In detail let

$$\bar{q} \xrightarrow{\alpha_0} q^1 \xrightarrow{\alpha_1} \ldots \xrightarrow{\alpha_n} \bar{q}$$

constitute the cycle in $\bar{T}$ where $\alpha_l \in excl\,(K')$. Choose a state $q' \in \prod\limits_{i \in K \setminus \bar{K}'} Q_i$ such that $(\bar{q}, q')$ is reachable in Sys. Then

$$(\bar{q}, q') \xrightarrow{\alpha_0} \left(q^1, q'\right) \xrightarrow{\alpha_1} \ldots \xrightarrow{\alpha_n} (\bar{q}, q') \xrightarrow{\alpha_0} \left(q^1, q'\right) \xrightarrow{\alpha_1} \ldots$$

yields a run in Sys which does not involve any component from K', and K' is not live. The case where $\bar{T}$ contains an infinite path starting in a state as above is treated analogously.

For the other direction we assume that in $\bar{T}$ there is neither any cycle nor any infinite path as described above. We want to show that K' is live in Sys. Assume that this is not the case. Then there must be a run

$$\sigma = q^0 \xrightarrow{\alpha_0} q^1 \xrightarrow{\alpha_1} \ldots$$

in Sys such that q^0 is reachable and no α_l involves any component from K'. This means that every α_l is in $excl\,(K')$. By deleting all local states q_i^l where $i \in K \setminus \bar{K}'$ we get an infinite sequence $\bar{\sigma}$ in $\bar{T}$. Note that all $\bar{q}^l$ on $\bar{\sigma}$ have the property described in the proposition. Either all states on $\bar{\sigma}$ are pairwise distinct or $\bar{\sigma}$ contains a cycle. In both cases we obtain a contradiction. ☐

References

[BBS06] A. Basu, M. Bozga, and J. Sifakis. Modeling Heterogeneous Real-time Systems in BIP, 2006. Invited Lecture SEFM 2006.

[GGMC+07a] G. Gössler, S. Graf, M. Majster-Cederbaum, M. Martens, and J. Sifakis. An Approach to Modelling and Verification of Component Based Systems. In *Proceedings of the 33rd International Conference on Current Trends in Theory and Practice of Computer, Science SOFSEM07*, volume 4362 of *LNCS*, 2007.

[GGMC+07b] Gregor Gössler, Susanne Graf, Mila Majster-Cederbaum, Moritz Martens, and Joseph Sifakis. Ensuring Properties of Interaction Systems. In *Program Analysis and Computation, Theory and Practice*, volume 4444 of *LNCS*, 2007.

[GS02] Gregor Gössler and Joseph Sifakis. Composition for Component-Based Modeling. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem-Paul de Roever, editors, *Formal Methods for Components and Objects*, Leiden, 2002.

[GS03a] Gregor Gössler and Joseph Sifakis. Component-Based Construction of Deadlock-Free Systems. In *proceedings of FSTTCS 2003, Mumbai, India*, volume 2914 of *LNCS*, pages 420–433, December 2003.

[GS03b] Gregor Gössler and Joseph Sifakis. Priority Systems. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem-Paul de Roever, editors, *Formal Methods for Components and Objects*, Leiden, 2003.

[GS05] Gregor Gössler and Joseph Sifakis. Composition for Component-Based Modeling. *Sci. Comput. Program.*, 55(1-3):161–183, 2005.

[MCMM07] M. Majster-Cederbaum, M. Martens, and C. Minnameier. A Polynomial-Time-Checkable Sufficient Condition for Deadlock-freeness of Component Based Systems. In *Proceedings of the 33rd International Conference on Current Trends in Theory and Practice of Computer, Science SOFSEM07*, volume 4362 of *LNCS*, 2007.

[Min06] Christoph Minnameier. Deadlock-Detection in Component--Based Systems is NP-hard. Technical report TR-2006-015, Universität Mannheim, 2006.

[MMMC06] M. Martens, C. Minnameier, and M. Majster-Cederbaum. Deciding Liveness in Component-Based Systems is NP-hard. Technical report TR-2006-017, University of Mannheim, Fakultät Mathematik und Informatik, 2006.

[Sif04] Joseph Sifakis. Modeling Real-time Systems, 2004. Keynote talk RTSS04.

[Sif05] Joseph Sifakis. A Framework for Component-based Construction, 2005. SEFM 2005: pp. 293 - 300.

[Tan01] Andrew S. Tanenbaum. *Modern Operating Systems*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2001.