Reihe Informatik. TR-2006-017
Revised Version: January 12th, 2007

Deciding Liveness in Component-Based Systems is NP-hard

Mila Majster-Cederbaum, Moritz Martens*, Christoph Minnameier

Institut für Informatik
Universität Mannheim, Germany

Abstract. Interaction systems are a formal model for component-based systems. Combining components via connectors to form more complex systems may give rise to deadlock situations. In a system that has been shown to be deadlock-free one can ask if a set of components is live. We present here a polynomial time reduction from 3-SAT to the question whether a set of components is live in a deadlock-free system.

1 Introduction

We consider a setting where components are combined via connectors to form more complex systems [GS03,Sif04,Sif05,GGMCS06,BBS06]. 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 can be represented via a labeled transition system with starting state, where in each state there is at least one action available. Components are glued together via connectors, where each connector connects certain ports. In the global system obtained by gluing components (local) deadlocks may arise where a group of components is engaged in a cyclic waiting and will thus no longer participate in the progress of the global system (cf. [Tan01]). In a deadlock-free system at least one interaction is enabled in every reachable state. Then one can ask the question whether a subset of components K' is live, i.e. in every infinite sequence of connectors there are infinitely many interactions that let a component from K' participate. We show that deciding whether a set of components is live in a deadlock-free system is NP-hard by encoding the classic 3-SAT problem in interaction systems in such a way that a formula is not satisfiable if and only if a certain component is live in the system corresponding to the formula. The system will be constructed such that the component in question can participate in every connector except one. Then it will be shown that a state where this connector can be perfomed repeatedly is reachable if and only if the formula is satisfiable.

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

The paper is organized as follows. Section 2 contains the basic definitions. Section 3 gives the polynomial time reduction from 3-SAT to interaction systems and the proofs. There will be a short example as well. Section 4 contains the conclusion and a discussion of related work.

2 Components, Connectors and Interaction Systems

We consider *interaction systems*, a model for component-based systems that was proposed and discussed in detail in [GS03,Sif05,GS05,BBS06]. An *interaction system* is a tuple $Sys = (K, \{A_i\}_{i \in K}, C, \{T_i\}_{i \in K})$[1], where K is the set of *components*. Without loss of generality we assume $K = \{1, \ldots, n\}$. Each component $i \in K$ offers a finite set of *ports* A_i for cooperation with other components. The port sets A_i are pairwise disjoint. Cooperation is described by *connectors*. A *connector* is a set of actions $c \subseteq \bigcup_{i \in K} A_i$, where for each component i at most one action $a_i \in A_i$ is in c. A connector set C is a set of connectors, such that every action of every component occurs in at least one connector of C and no connector contains any other connector.

The local behavior of each component i is described by $T_i = (Q_i, A_i, \rightarrow_i, q_i^0)$, where Q_i is the finite set of local states, $\rightarrow_i \subseteq Q_i \times A_i \times Q_i$ the local transition relation and $q_i^0 \in Q_i$ is the local starting state. Given a connector $c \in C$ and a component $i \in K$ we denote by $i(c) := A_i \cap c$ the participation of i in c.

For $q_i \in Q_i$ we define the set of *enabled actions* $ea(q_i) := \{a_i \in A_i \mid \exists q_i' \in Q_i,$ s.t. $q_i \xrightarrow{a_i} q_i'\}$. We assume that the T_i's are non-terminating, i.e. $\forall i \in K \; \forall q_i \in Q_i \; ea(q_i) \neq \emptyset$.

The global behavior $T_{Sys} = (Q, C, \rightarrow, q^0)$ of Sys (henceforth called *global transition system*) is obtained from the behaviors of the individual components, given by the transition systems T_i, and the connectors C in a straightforward manner:

- $Q = \prod_{i \in K} Q_i$, the Cartesian product of the Q_i, which we consider to be order independent. We denote states by tuples $(q_1, \ldots, q_n)$ and call them global states.
- the relation $\rightarrow \subseteq Q \times C \times Q$, defined by
  $\forall c \in C \; \forall q, q' \in Q \quad q = (q_1, \ldots, q_n) \xrightarrow{c} q' = (q_1', \ldots, q_n')$ iff
  $\forall i \in K \quad (q_i \xrightarrow{i(c)}_i q_i'$ if $i(c) \neq \emptyset$ and $q_i' = q_i$ otherwise).
- $q^0 = (q_1^0, \ldots, q_n^0)$ is the starting state for *Sys*.

[1] The model in [GS03] is more general. It introduces a notion of complete interaction, which is a subset of a connector and distinguishes between connectors and complete interactions. We are able to show NP-hardness for liveness-check of a set of components in interaction systems without the use of complete interactions, so we omit them for ease of notation. Note that this yields a stronger not weaker result than using complete interactions. Readers who are familiar with interaction systems may simply assume $Comp = \emptyset$ for $Sys(F)$ in Section 3.

In the global transition system a transition labeled with c may take place when each component participating in c is ready to perform $i(c)$.
For an example of an interaction system see Example 1 at the end of section 3.

Let $q = (q_1, \ldots, q_n) \in Q$ be a global state. We say that some non-empty set $D = \{j_1, j_2, \ldots, j_{|D|}\} \subseteq K$ of components is in *deadlock* in q if $\forall i \in D\ \forall c \in C$, such that $c \cap ea(q_i) \neq \emptyset\ \exists j \in D$, where $j(c) \not\subseteq ea(q_j)$. Then we say that i waits for j. A system has a local *deadlock* in some global state q if there is $D \subseteq K$, that is in deadlock in q. If $D = K$, the system is globally deadlocked. Hence a global deadlock is a special case of a local deadlock. A system is *deadlock-free*, if there is no reachable state q and $D \subseteq K$, such that D is in deadlock in q. A system is *globally deadlock-free*, if there is no reachable state q such that K is in deadlock in q.

In a globally deadlock-free system it is always possible to proceed from a reachable state. We define a *run* of a system to be an infinite sequence

$$\rho := q^0 \stackrel{\alpha_0}{\to} q^1 \stackrel{\alpha_1}{\to} \ldots$$

such that α_n is in C for all n. In a globally deadlock-free system we define a subset $K' \subset K$ to be *live* if in every run there are infinitely many interactions α_n such that there is some $k \in K'$ with $k(\alpha_n) \neq \emptyset$. If $K' = \{k\}$ only contains one element we simply speak of liveness of k.

We denote by IS the set of all interaction systems and by DLFIS the set of interaction systems that are globally deadlock-free:

$$DLFIS := \{Sys \in IS \mid Sys \text{ is globally deadlock-free}\}$$

We consider the well-studied NP-complete 3-SAT problem [GJ79] where the formula is a conjunction of clauses k_i, each of which is a disjunction of 3 literals, (i.e. possibly negated variables) and reduce the question whether a formula is not satisfiable to the question of deciding whether a set of components is live in a globally deadlock-free system.

3 The Reduction

Let $F = k_1 \wedge \ldots \wedge k_n$ with $k_i = (l_{(i,1)} \vee l_{(i,2)} \vee l_{(i,3)})$ be a propositional formula in 3-KNF, where $l_{(i,1)}, l_{(i,2)}, l_{(i,3)}$ are positive literals (i.e. variables) or negative literals (i.e. negated variables). We agree upon the following notations. For a literal l let $var(l)$ denote the variable occuring in that literal. For F let $var(F) := \left\{var\left(l_{(i,j)}\right) | 1 \leq i \leq n, 1 \leq j \leq 3\right\}$ denote the set of variables occuring in F.

In the following, we construct an interaction system $Sys(F) \in DLFIS$ with component-set $K \cup \{\kappa_l\}$, such that $(F \notin 3-SAT) \Leftrightarrow (\kappa_l$ is live in $Sys(F))$. Besides the component κ_l we represent each clause k_i by a component $(i,0)$ and each literal $l_{(i,j)}$ by a component (i,j). We denote by $\tilde{K}$ the set $K \backslash \{\kappa_l\}$.

$$Sys(F) := (K, \{A_{(i,j)}\}_{(i,j)\in\tilde{K}} \cup \{A_{\kappa_l}\}, C, \{T_{i,j}\}_{(i,j)\in\tilde{K}} \cup \{T_{\kappa_l}\}),$$

where:

$$\begin{aligned} K &:= \{(i,j) \mid 1 \leq i \leq n, 0 \leq j \leq 3\} \cup \{\kappa_l\} \\ A_{(i,0)} &:= \{true_i, SAT_i\} \text{ for } 1 \leq i \leq n \\ A_{(i,j)} &:= \left\{set1_{(i,j)}, set0_{(i,j)}, true_{(i,j)}, a_{(i,j)}\right\} \text{ for } 1 \leq i \leq n, j \neq 0 \\ A_{\kappa_l} &:= \{a_{\kappa_l}\} \end{aligned}$$

We define the following connectors. First we have:

$$sat := \{SAT_1, \ldots, SAT_n\}$$

Next there are two sorts of connectors representing the assignment of 1 or 0 to a variable $x \in var(F)$:

$$set1_x := \left\{a_{\kappa_l}, set1_{(i_1,j_1)}, \ldots, set1_{(i_m,j_m)}\right\}$$

where $x = var\left(l_{(i_1,j_1)}\right) = \ldots = var\left(l_{(i_m,j_m)}\right)$ and there is no other literal l with $x = var(l)$. Analogously we define

$$set0_x := \left\{a_{\kappa_l}, set0_{(i_1,j_1)}, \ldots, set0_{(i_m,j_m)}\right\}$$

where $x = var\left(l_{(i_1,j_1)}\right) = \ldots = var\left(l_{(i_m,j_m)}\right)$ and there is no other literal l with $x = var(l)$. Another group of connectors is defined by:

$$t_{(i,j)} := \left\{a_{\kappa_l}, true_{(i,j)}, true_i\right\}$$

Finally:

$$c_a := \{a_{\kappa_l}\} \cup \left\{a_{(i,j)} | 1 \leq i \leq n, j \neq 0\right\}$$

Then we set

$$C := \{sat\} \cup \bigcup_{x \in var(F)} \left(\{set1_x\} \cup \{set0_x\}\right) \cup \bigcup_{1 \leq i \leq n, j \neq 0} \left\{t_{(i,j)}\right\} \cup \{c_a\}$$

The local transition system for κ_l is given in Figure 1. The local transition system $T_{(i,0)}$ for $1 \leq i \leq n$ is given in Figure 2 (a). The local transition system $T_{(i,j)}$ for $1 \leq i \leq n$, $j \neq 0$ and $l_{(i,j)}$ is a positive (resp. negative) literal is given in Figure 2 (b) (resp. (c)).

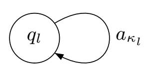

Fig. 1. The transition system for κ_l

We call components $(i,0)$ clause-components and components (i,j) where $j \neq 0$ literal-components. For a component (i,j) we call the state $q^t_{(i,j)}$ its *true*-state and if it exists $q^f_{(i,j)}$ its *false*-state. For a global state q we assume that the components are ordered as follows

$$q^0 := (q^0_{(1,0)}, q^0_{(1,1)}, q^0_{(1,2)}, q^0_{(1,3)}, q^0_{(2,0)}, \ldots, q^0_{(n,3)}, q_l).$$

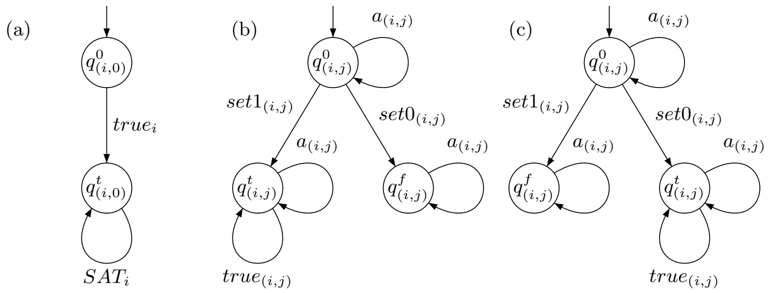

Fig. 2. The transition systems for clause-components (a) and literal-components (b) and (c)

Remark 1: Note that there is no blow-up in notation when we go from F to $Sys(F)$. The four transition systems we introduce for each clause are of constant size as is the transition system for κ_l. The $set1_x$- and $set0_x$-connectors have an overall size which is linear in the number of literals in F and their number is linear in the number of different variables in F. The *sat*-connector is linear in the number of clauses in F. Other than that there are $3n$ connectors of constant length and one of length $3n+1$.

Lemma 1 *In $Sys\,(F)$ a state where sat can be performed is reachable from the initial state if and only if F is satisfiable.*

Proof. We show both implications.
⇒: Let

$$\rho_s := q^0 \stackrel{\alpha_0}{\rightarrow} \ldots \stackrel{\alpha_{s-1}}{\rightarrow} q^s$$

be an initial fragment of a run such that *sat* is enabled in q^s. This means that all clause-components have to be in their *true*-state in q^s which is only possible if for every i there exists exactly one j_i such that on ρ_s the connector $t_{(i,j_i)}$ was performed. This means that for every i the component (i, i_j) must have moved to its *true*-state somewhere along ρ_s before the execution of $t_{(i,j_i)}$. We define an assignment $\sigma_\rho : var\,(F) \rightarrow \{1, 0\}$ as follows.

$$\sigma_\rho\,(x) := \begin{cases} 1 \text{ if } \nexists i \text{ such that } x = var\left(l_{(i,j_i)}\right) \\ 1 \text{ if } \exists i \text{ such that } x = l_{(i,j_i)} \\ 0 \text{ if } \exists i \text{ such that } \neg x = l_{(i,j_i)} \end{cases}$$

σ_ρ is well-defined. It is clear that the first case and one of the two other cases exclude each other. Assume there is a variable x such that the condition of the two last cases is satisfied. Then there exist i_1 and i_2 such that

$x = var\left(l_{\left(i_1,j_{i_1}\right)}\right) = var\left(l_{\left(i_2,j_{i_2}\right)}\right)$ where $l_{\left(i_1,j_{i_1}\right)}$ is positive and $l_{\left(i_2,j_{i_2}\right)}$ is negative. The components (i_1, j_{i_1}) and (i_2, j_{i_2}) occur in the connectors $set1_x$ and $set0_x$ together. Both components have moved to their *true*-state along ρ_s which means that (i_1, j_{i_1}) must have performed $set1_{\left(i_1,j_{i_1}\right)}$ and (i_2, j_{i_2}) must have performed $set0_{\left(i_2,j_{i_2}\right)}$. Therefore both $set1_x$ and $set0_x$ must have been performed somewhere along ρ_s which is not possible.

Now it remains so show that $\sigma_\rho(F) = 1$. Let k_i be a clause of F. If (i, j_i) represents a positive literal $\sigma_\rho\left(var\left(l_{(i,j_i)}\right)\right) = 1$ and therefore $\sigma_\rho(k_i) = 1$. Otherwise $\sigma_\rho\left(var\left(l_{(i,j_i)}\right)\right) = 0$ and again $\sigma_\rho(k_i) = 1$. We conclude $\sigma_\rho(F) = 1$ and F is satisfiable.
$\Leftarrow$: If F is satisfiable let σ be an assignment such that $\sigma(F) = 1$. We construct an initial fragment ρ_s of a run that ends in a state where SAT is enabled.

Every variable in $var(F)$ is assigned a value by σ. In the order of appearance of variables in F we perform the following connectors. If $\sigma(x) = 1$ we perform $set1_x$. Otherwise we perform $set0_x$. This is possible because no literal-component (i, j) can leave its initial state without the other literal-components $\left(\tilde{i}, \tilde{j}\right)$ with $var\left(l_{\left(\tilde{i},\tilde{j}\right)}\right) = var\left(l_{(i,j)}\right)$. At this point of the run every literal-component is either in its *false*-state or in its *true*-state. All other components are in their respective initial state.

Because $\sigma(F) = 1$ we know that for every clause k_i there is a literal-component (i, j) such that $\sigma\left(l_{(i,j)}\right) = 1$. The corresponding literal-component must be in its *true*-state because of the definition of the transition systems for these components. This means that for all $1 \leq i \leq n$ we can perform the connector $t_{(i,j)}$.

Thus we reach a state q where all clause-components are in their *true*-state. *sat* can be performed in q and we are done.

Proposition 1 *$Sys(F)$ is deadlock-free and F is not satisfiable if and only if κ_l is live in $Sys(F)$.*

Proof. It is clear that $Sys(F)$ is deadlock-free because c_a is enabled in every global state.
$\Leftarrow$: Let κ_l be live in $Sys(F)$. We want to show that F is not satisfiable. Assume that there is an assignment σ such that $\sigma(F) = 1$. From Lemma 1 we conclude that there is an initial fragment of a run

$$\rho_s := q^0 \stackrel{\alpha_0}{\rightarrow} \ldots \stackrel{\alpha_{s-1}}{\rightarrow} q^s$$

such that *sat* can be performed in q^s. κ_l only participates finitely many often in ρ defined by

$$\rho_s := q^0 \stackrel{\alpha_0}{\rightarrow} \ldots \stackrel{\alpha_{s-1}}{\rightarrow} q^s \stackrel{sat}{\rightarrow} q^s \stackrel{sat}{\rightarrow} \ldots$$

This is a contradiction.
$\Rightarrow$: Let F not be satisfiable. We want to show that κ_l is live in $Sys(F)$. Assume this is not the case. Then there is a run

$$\rho := q^0 \stackrel{\alpha_0}{\rightarrow} q^1 \stackrel{\alpha_1}{\rightarrow} \ldots$$

such that from some point on κ_l does not participate any more. Note that the only connector that does not let κ_l participate is *sat*. This means that there is some n_0 such that $\alpha_n = sat$ for all $n \geq n_0$. Therefore a state where *sat* can be performed is reachable in $Sys\,(F)$. Lemma 1 implies that F is satisfiable. This is a contradiction.

Example 1: Let $F = (x_1 \vee \overline{x_2} \vee x_3) \wedge (\overline{x_1} \vee x_2 \vee \overline{x_3}) \wedge (\overline{x_1} \vee \overline{x_2} \vee \overline{x_3})$. F is satisfiable, namely $\sigma(F) = 1$ for $\sigma(x_1) = 1, \sigma(x_2) = 1, \sigma(x_3) = 0$.

Consider $Sys(F) = (K, \{A_i\}_{i \in K}, C, \{T_i\}_{i \in K})$. The set of components is given by $K = \{(1,0), (1,1), (1,2), (1,3), (2,0), \ldots, (3,3), \kappa_l\}$ and the port sets A_i for $i \in K$, the connector-set C as well as the local transition systems $\{T_i\}_{i \in K}$ are defined as above. One possible fragment of a run constructed from σ according to Lemma 1 is as follows.

$$\rho := q^0 \overset{set1_{x_1}}{\rightarrow} q^1 \overset{set1_{x_2}}{\rightarrow} q^2 \overset{set0_{x_3}}{\rightarrow} q^4 \overset{t_{(1,1)}}{\rightarrow} q^5 \overset{t_{(2,2)}}{\rightarrow} q^5 \overset{t_{(3,3)}}{\rightarrow} q^6$$

In the first step component $(1,1)$ moves to its *true*-state and $(2,1)$ and $(3,1)$ move to their respective *false*-state. Analoguously the other six literal-components change their state according to $set1_{x_2}$ and $set0_{x_3}$. In steps four to six components $(1,0)$, $(2,0)$, and $(3,0)$ move to their *true*-state together with $(1,1)$, $(2,2)$, respectively $(3,3)$. Note that in the fifth step $t_{(2,3)}$ could also have been performed because $q^5_{(2,3)} = q^t_{(2,3)}$ as well. *sat* is enabled in q^6.

4 Conclusion and Related Work

This work is closely based on [Min06] where it was shown that the detection of deadlocks in interaction systems is NP-hard. We followed the reduction presented there making the necessary adjustments. This led to the result that the question whether a set of components is live is NP-hard for interaction systems even without the use of complete interactions. This yields a motivation for sufficient criteria for liveness that work in polynomial-time such as the polynomial-time checkable sufficient condition for liveness of a subset of components in interaction systems presented in [GGMC⁺07]. In [GSG⁺06] one can find sufficient criteria for other properties of component-based systems as well as considerations about ensuring properties by construction using composability.

References

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

[GGMC⁺07] 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.

[GGMCS06] Gregor Gößler, Susanne Graf, Mila Majster-Cederbaum, and Joseph Sifakis. Dynamic Construction of Deadlock-free Interaction Systems. Technical report, INRIA, 2006. in preparation.

[GJ79] M. R. Gary and D. S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness*. New York: W.H. Freeman, 1979.

[GS03] 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. downloadable through http://www-verimag.imag.fr/ sifakis/.

[GS05] Gregor Gößler and Joseph Sifakis. Composition for component-based modeling. *Sci. Comput. Program.*, 55(1-3):161–183, 2005.

[GSG+06] Gregor Gößler, Joseph Sifakis, Susanne Graf, Mila Majster-Cederbaum, and Moritz Martens. Ensuring Properties of Interaction Systems, 2006. accepted for publication.

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