Deriving Complexity Results for Interaction Systems from 1-Safe Petri Nets

Mila Majster-Cederbaum and Christoph Minnameier*

Institut für Informatik
Universität Mannheim, Germany
cmm@informatik.uni-mannheim.de

Abstract. Interaction systems are a formal model for component-based systems, where components are combined via connectors to form more complex systems. We compare interaction systems (*IS*) to the well-studied model of 1-safe Petri nets (*1SN*) by giving a translation map_1: *1SN* $\rightarrow$ *IS* and a translation map_2: *IS* $\rightarrow$ *1SN*, so that a 1-safe Petri net (an interaction system) and its according interaction system (1-safe Petri net) defined by the respective mapping are isomorphic up to some label relation R. So in some sense both models share the same expressiveness. Also, the encoding map_1 is polynomial and can be used to reduce the problems of reachability, deadlock and liveness in *1SN* to the problems of reachability, deadlock and liveness in *IS*, yielding PSPACE-hardness for these questions.

1 Introduction

In [GS03], Gössler and Sifakis presented *interaction systems*, a model for component-based concurrent systems. As typical for component-based systems, interaction systems display two different layers of description: On the one hand the components, which are used to describe the communicating units, together with the ports over which they communicate. On the other hand the *glue code*, i.e. the information about the way components may communicate with each other. I/O-Automata [LT89] and interface automata [dAH01] can be considered as subclasses of interaction systems, for the latter feature a more general notion of communication. E.g. interaction systems allow different degrees of parallelism, i.e. different interactions may involve different numbers of participants.

Interaction systems seem to be an appropriate model for a variety of different types of systems. More details about interaction systems and their properties can be found in [Sif04, Sif05, GGM+07b, GGM+07a, MMM07b, MMM07a]. A framework for component-based modelling using interaction systems has been implemented in the BIP-project [BBS06, GQ07, BS07] and applied to [BMP+07]. Furthermore, interaction systems have been used to model biochemical reactions [MSW07] and they serve as a commom semantic framework for the SPEEDS-project [BCSM07].

* Corresponding author.

V. Geffert et al. (Eds.): SOFSEM 2008, LNCS 4910, pp. 352–363, 2008.

The aim of this paper is to answer some relevant questions concerning the inherent complexity of properties of interaction systems. A first result, concerning these matters is given in [Min07], where it is shown that the problems of local and global deadlock are NP-hard. Here, we obtain stronger results by establishing a relation between interaction systems and the well-studied model of 1-safe Petri nets.

In [CEP93] important results about PSPACE-completeness of behavioral questions in 1-safe Petri nets have been established, which is the starting point of our investigation. In particular, we consider the traditional Petri net token-game semantics for Petri nets, which does not allow the concurrent performance of multiple transitions (even if their presets and postsets are disjoint). In other words, we restrict ourselves to the intrinsic concurrency of the Petri net model, i.e. the fact that a transition may already involve multiple places.

This decision is natural for our purpose of comparing the model of 1-safe nets to the model of interaction systems, because in the latter's semantics, we also may concurrently perform multiple actions within an interaction but only one interaction at a time.

The main part of this work consists of giving mappings from one model to the other and isomorphism relations for the resulting pairs of nets and systems.

The mappings and isomorphism relations are then used to derive PSPACE-hardness results for some important behavioral questions for interaction systems, namely reachability, global deadlock and liveness.

The paper is organized as follows. Section 2 contains the basic definitions. Sections 3 and 4 give the respective translations between 1-safe Petri nets and interaction systems. Section 5 contains a conclusion and a discussion of related work.

2 Definitions

2.1 1-Safe Petri Nets

A ***Petri net*** [CEP93] is a fourtuple $N = (P, T, F, M_0)$ such that:

- P and T are finite disjoint sets. Their elements are called ***places*** and ***transitions***, respectively.
- $F \subseteq (P \times T) \cup (T \times P)$. F is called the ***flow relation***.
- $M_0 : P \to \mathbb{N}$ is called the ***initial marking*** of N. In general, a mapping $M : P \to \mathbb{N}$ is called a ***marking*** of N. By $\mathcal{M}$ we denote the set of all markings of a net.

For places as well as transitions we define the notion of preset and postset:
For $p \in P$, ***preset(p)*** $:= \{t \in T \mid (t,p) \in F\}$, ***postset(p)*** $:= \{t \in T \mid (p,t) \in F\}$.
For $t \in T$, ***preset(t)*** $:= \{p \in P \mid (p,t) \in F\}$, ***postset(t)*** $:= \{p \in P \mid (t,p) \in F\}$.

For technical reasons we only consider nets in which every node has a nonempty preset or a nonempty postset. We let $+$ denote the union of multisets.

Let $N = (P, T, F, M_0)$ be a Petri net. A transition $t \in T$ is ***enabled*** under a marking M if $M(p) > 0$ for every place p in the preset of t. Given a transition t, we define a relation $\xrightarrow{t}_N$ as follows: $M \xrightarrow{t}_N M'$ if t is enabled under M and $M'(p) = M(p) + F(t, p) - F(p, t)$, where $F(x, y)$ is 1 if $(x, y) \in F$ and 0 otherwise. We say that the transition t is performed at M. We define the global transition system (or global behavior) T_N of N by $T_N = (\mathcal{M}, T, \rightarrow_N, M_0)$.

For $M, M' \in \mathcal{M}$, we write $\boldsymbol{M \rightarrow^*_N M'}$ if there are ($k \in \mathbb{N}$ and) markings $M^1, \ldots, M^k \in \mathcal{M}$ and transitions $t_1, \ldots, t_{k+1} \in T$ that build a transition sequence $M \xrightarrow{t_1}_N M^1 \xrightarrow{t_2}_N \ldots \xrightarrow{t_k}_N M^k \xrightarrow{t_{k+1}}_N M'$ in T_N.

A marking M of a net N is called 1-safe, if for every place p of the net $M(p) \leq 1$. We identify a 1-safe marking with the set of places such that $M(p) = 1$. A net N is called 1-safe if all its reachable markings are 1-safe. 1-safe nets are a well studied computation model. The following questions are known to be PSPACE-complete [CEP93].

The ***reachability*** problem for 1-safe nets consists of deciding, given a 1-safe net $N = (P, T, F, M_0)$ and a marking M of N, whether $M_0 \rightarrow^*_N M$.

The ***liveness*** problem for 1-safe nets consists of deciding, given a 1-safe net $N = (P, T, F, M_0)$ if every transition can always occur again. More precisely, if for every reachable marking M and every transition t, there is $M' \in \mathcal{M}$ with $M \rightarrow^*_N M'$ and M' enables t.

The ***deadlock*** problem for 1-safe nets consists of deciding, given a 1-safe net $N = (P, T, F, M_0)$, if every reachable marking enables some transition. If this is the case we call the net deadlock-free.

Example 1

The 1-safe net N_1 is given (by its graphical representation) in Figure 1. N_1 is deadlock free and even live and the set of reachable markings is $\{\{p_1, p_2, p_3\}, \{p_3, p_4, p_5\}, \{p_1, p_6\}\}$.

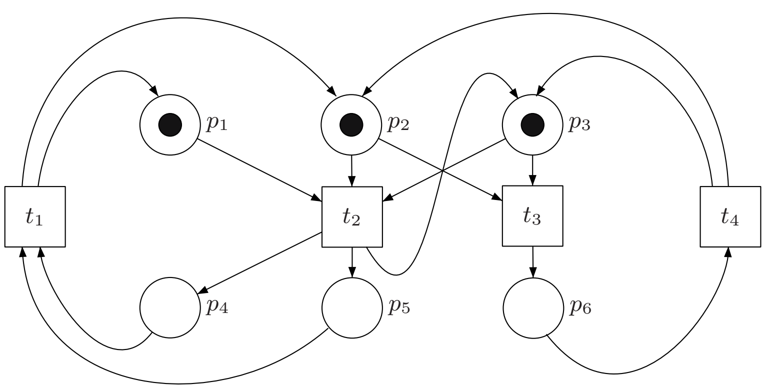

Fig. 1. A 1-safe net N_1

2.2 Interaction Systems

We review here interaction systems, a model for component-based systems that was proposed and discussed in detail in [GS03, Sif05, GS05, BBS06, GGM+07b, GGM+07a, MMM07a]. An ***interaction system*** is a tuple $Sys = (K, \{A_i\}_{i \in K}, C, Comp, \{T_i\}_{i \in K})$, where K is the set of ***components***. W.l.o.g. we assume $K = \{1, \ldots, n\}$. Each component $i \in K$ offers a finite set A_i of ***ports*** (also called ***actions***) for cooperation with other components. The ***port sets*** A_i are pairwise disjoint. Cooperation is described by connectors and complete interactions. A ***connector*** is a finite set of actions $c \subseteq \bigcup_{i \in K} A_i$, subject to the constraint that for each component i at most one action $a_i \in A_i$ is in c. A connector $c = \{a_{i_1}, \ldots, a_{i_k}\}$ with $a_{i_j} \in A_{i_j}$ describes that the components $i_1, \ldots, i_j$ cooperate via these ports.

A ***connector set*** C is a finite set of connectors, s.t. every action of every component occurs in at least one connector of C and no connector contains any other connector. Sometimes not all components involved in a connector are ready to perform their respective action. Still, we might want to allow those that are ready to go on. For this we may designate certain subsets of connectors as ***complete interactions***. Let $Comp$ be a designated set of complete interactions. $Comp$ has to be upwards-closed w.r.t. C, i.e.: $\forall \alpha \in Comp\ \forall c \in C\ ((\alpha \subset \alpha' \subseteq c) \Rightarrow \alpha' \in Comp)$.

We call $Int := C \cup Comp$ the set of ***interactions***[1]. (The distinction between connectors and complete interactions is irrelevant for our encodings).

The local behavior of each component i is described by a transition system $T_i = (Q_i, A_i, \rightarrow_i, q_i^0)$, 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 an interaction $\alpha \in Int$ and a component $i \in K$ we denote by $i(\alpha) := A_i \cap \alpha$ the ***participation*** of i in α. For ease of notation, we identify a singleton set with its element.

For $q_i \in Q_i$ we define the set of ***enabled actions*** $ea(q_i) := \{a_i \in A_i \mid \exists q_i' \in Q_i, \text{ s.t. } q_i \xrightarrow{a_i}_i q_i'\}$. 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*** $T_{Sys} = (Q, Int, \rightarrow_{Sys}, q^0)$ of Sys (henceforth also referred to as global transition system) is obtained from the behaviors of the individual components, given by the transition systems T_i, and the interactions Int 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 $\rightarrow_{Sys} \subseteq Q \times Int \times Q$, defined by
  $\forall \alpha \in Int\ \forall q, q' \in Q \quad q = (q_1, \ldots, q_n) \xrightarrow{\alpha}_{Sys} q' = (q_1', \ldots, q_n') \quad \text{iff}$
  $\forall i \in K \quad (q_i \xrightarrow{i(\alpha)}_i q_i' \text{ if } i(\alpha) \neq \emptyset \text{ and } q_i' = q_i \text{ otherwise}).$
- $q^0 = (q_1^0, \ldots, q_n^0)$ is the starting state for Sys.

[1] In the original nomenclature of [GS03], subsets of connectors in general are called interactions. This more general notion of interaction is however only needed for the purpose of composing interaction systems out of smaller interaction systems.

Less formally, a transition labeled by α may take place in the global transition system when each component i participating in α is ready to perform $i(\alpha)$.

Example 2
Let $Sys_1 = (\{1,2,3\}, \{A_i\}_{1\leq i\leq 3}, C, Comp, \{T_i\}_{1\leq i\leq 3})$, where $A_1 = \{a_1, b_1\}$, $A_2 = \{a_2, b_2\}, A_3 = \{a_3, b_3, d_3\}$, $C = \{\{a_1, a_2, a_3\}, \{b_1, b_2, b_3\}, \{d_3\}\}$, $Comp = \{\{b_1, b_2\}\}$ and the local transition systems T_i are given in Figure 2.

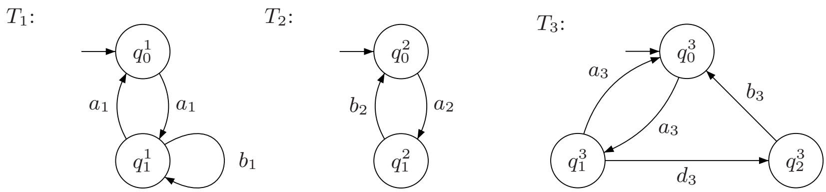

Fig. 2. The T_i's for Sys_1

For the following definitions let Sys be an interaction system:

Let $\rightarrow^*_{Sys}$ denote the reflexive and transitive closure of $\rightarrow_{Sys}$.

Given a state $q \in Q$ we denote by ***reachability of q*** the question, whether q is reachable in T_{Sys}, i.e. whether $q^0 \rightarrow^*_{Sys} q$.

The question whether Sys contains a **global deadlock** (henceforth simply referred to as a deadlock) is the question whether there is a reachable global state q such that $q \not\rightarrow$.

We say **a component** $i \in K$ ***is live***[2] in Sys, if for any reachable global state there is some $q' \in Q$ with $q \rightarrow^*_{Sys} q'$ such that there exist $\alpha \in Int$ and $q'' \in Q$ with $q' \xrightarrow{\alpha}_{Sys} q''$, where i participates in α.

If a component $i \in K$ is live in Sys then at each reachable global state a clever scheduler can continue in such a way that eventually an interaction may be performed in which i participates.

2.3 Isomorphism up to a Label Relation *R*

We define a notion of isomorphism, namely isomorphism up to a label relation R, which we use to establish a relation between transition systems that use different label sets L_1 and L_2. R then defines which labels in L_1 we want to correspond to which labels in L_2.

Let $T_i = (Q_i, L_i, \rightarrow_i, q_i^0)$, $i \in \{1,2\}$ be two labeled transition systems. Given a *label relation* $R \subseteq (L_1 \times L_2)$, that relates labels of L_1 to labels of L_2, we say that T_1 and T_2 are isomorphic up to R iff there exists a bijective function $f : Q_1 \rightarrow Q_2$, such that $f(q_1^0) = q_2^0$ and $\forall q_1 \in Q_1, q_2 \in Q_2$ the following two propositions hold:

[2] Note that this notion of liveness does not coincide with the one defined in [MMM07a].

1) $q_1 \xrightarrow{l_1}_1 q_1' \quad \Rightarrow \quad \exists l_2 \in L_2$, s.t. $(l_1, l_2) \in R \ \wedge \ f(q_1) \xrightarrow{l_2}_2 f(q_1')$.
2) $q_2 \xrightarrow{l_2}_2 q_2' \quad \Rightarrow \quad \exists l_1 \in L_1$, s.t. $(l_1, l_2) \in R \ \wedge \ f^{-1}(q_2) \xrightarrow{l_1}_1 f^{-1}(q_2')$.

We say an interaction system and a 1-safe net are isomorphic up to a label relation R iff this holds for their respective global transition systems.

3 Translating 1-Safe Nets to Interaction Systems

Let $N = (P, T, F, M_0)$ be a 1-safe net. We give a translation map_1 from 1-safe Petri nets to interaction systems as follows. We introduce a component $\hat{p}$ for each place $p \in P$. The transition system $T_{\hat{p}}$ has only two states, one state $s^1_{\hat{p}}$ to reflect the fact that p contains a token, one state $s^0_{\hat{p}}$ to reflect that it doesn't. The transitions t adjacent to p define the transition relation of $T_{\hat{p}}$, where we distinguish three cases:

a) $t \in (preset(p) \setminus postset(p))$. When such a transition is performed in N, this means that p is empty before the performance of t and contains a token afterwards. Thus, we introduce an edge from $s^0_{\hat{p}}$ to $s^1_{\hat{p}}$ labeled by $a_{(t,p)}$.
b) $t \in (postset(p) \setminus preset(p))$. Inverse to a), i.e. we introduce an edge from $s^1_{\hat{p}}$ to $s^0_{\hat{p}}$ labeled by $a_{(p,t)}$.
c) $t \in (preset(p) \cap postset(p))$. This means there has to be a token in p to perform t and there will still be one there afterwards. In this case, we introduce a loop at $s^1_{\hat{p}}$ labeled by $a_{(t,p,t)}$.

For an example of a place with pre- and postset resp. its corresponding component, see Figure 3 (a) resp. (b). (Note that only edges adjacent to p are depicted.)

Now we define a connector $c(t)$ for each transition t. For the places adjacent to t again we distinguish three cases:

a) $p \in (preset(t) \setminus postset(t))$. This means that in order to perform t, there has to be a token in p, and there will be no token in p after performing t. Thus we include the action $a_{(p,t)}$ in $c(t)$ which already occurs in the component $\hat{p}$ in such a way that this fact is perfectly reflected.
b) $p \in (postset(t) \setminus preset(t))$. Inverse to a), i.e. we include the action $a_{(t,p)}$ in $c(t)$.
c) $p \in (preset(t) \cap postset(t))$. This means that in order to perform t, there has to be a token in p, and there still be a token in p after performing t. Thus we include the action $a_{(t,p,t)}$ in $c(t)$ which already occurs in the component $\hat{p}$ in the corresponding way.

For an example of a transition with pre- and postset resp. its corresponding connector, see Figure 4 (a) resp. (b). (Note that only edges adjacent to t are depicted.)

Formal definition of map_1:

$map_1(N) = \{K, \{A_i\}_{i\in K}, C, Comp, \{T_i\}_{i\in K}\}$, where

$$\boldsymbol{K} := \{\hat{p} \mid p \in P\}$$

$$\begin{aligned}
\text{For } \hat{p} \in K : A^{in}_{\hat{p}} &:= \{a_{(t,p)} \mid \exists t \in T, \text{ s.t. } p \in (preset(t) \setminus postset(t))\},\\
A^{out}_{\hat{p}} &:= \{a_{(p,t)} \mid \exists t \in T, \text{ s.t. } p \in (postset(t) \setminus pretset(t))\},\\
A^{inout}_{\hat{p}} &:= \{a_{(t,p,t)} \mid \exists t \in T, \text{ s.t. } p \in (preset(t) \cap postset(t))\}, \text{ and}\\
\boldsymbol{A_{\hat{p}}} &:= A^{in}_{\hat{p}} \cup A^{out}_{\hat{p}} \cup A^{inout}_{\hat{p}}.
\end{aligned}$$

$$\begin{aligned}
\boldsymbol{T_{\hat{p}}} &:= (\{s^0_{\hat{p}}, s^1_{\hat{p}}\}, A_{\hat{p}}, \rightarrow_{\hat{p}}, q^{\hat{p}}_s), \text{ where } A_{\hat{p}} \text{ has already been given,}\\
\rightarrow_{\hat{p}} &:= \{(s^0_{\hat{p}}, a_{(t,p)}, s^1_{\hat{p}}) \mid a_{(t,p)} \in A^{in}_{\hat{p}}\}\\
&\cup \{(s^1_{\hat{p}}, a_{(p,t)}, s^0_{\hat{p}}) \mid a_{(p,t)} \in A^{out}_{\hat{p}}\}\\
&\cup \{(s^1_{\hat{p}}, a_{(t,p,t)}, s^1_{\hat{p}}) \mid a_{(t,p,t)} \in A^{inout}_{\hat{p}}\}\\
q^0_{\hat{p}} &:= s^0_{\hat{p}} \text{ if } M_0(p) = 0 \text{ and } q^0_{\hat{p}} := s^1_{\hat{p}} \text{ if } M_0(p) = 1.
\end{aligned}$$

In order to define a connector for a transition we now relate the actions in $\bigcup_{i\in K} A_i$ to the transitions in the way described above.

$$\begin{aligned}
\text{For } t \in T\text{: } A^{in}_t &:= \{a_{(p,t)} \mid p \in (preset(t) \setminus postset(t))\},\\
A^{out}_t &:= \{a_{(t,p)} \mid p \in (postset(t) \setminus pretset(t))\},\\
A^{inout}_t &:= \{a_{(t,p,t)} \mid p \in (preset(t) \cap postset(t))\}\\
c(t) &:= A^{in}_t \cup A^{out}_t \cup A^{inout}_t\\
\boldsymbol{C} &:= \{c(t) \mid t \in T\}\\
\boldsymbol{Comp} &:= \emptyset
\end{aligned}$$

It remains to prove that C is indeed a connector set.

We observe that $\{A^{sup}_t \mid t \in T, sup \in \{in, out, inout\}\}$ is a disjoint decomposition of $\bigcup_{i\in K} A_i$. This is due to the fact that the A^{sup}_t's are defined following the definition of the $A^{sup}_{\hat{p}}$'s. C is just a coarser decomposition obtained from the one above by merging some of the disjoint subsets. So each action occurs exactly once in a connector, i.e. it occurs in at least one connector and no connector can be a subset of another connector.

Also, as $Comp = \emptyset$ we have upwards-closedness of $Comp$ w.r.t. C.

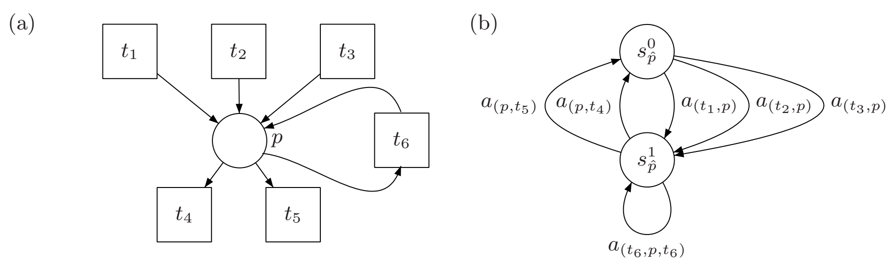

Fig. 3. A place with ingoing and outgoing transitions and its corresponding component

(a)

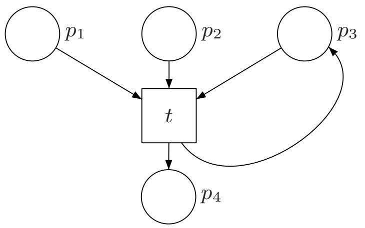

(b)

$$c(t) = \{a_{(p_1,t)}, a_{(p_2,t)}, a_{(t,p_3,t)}, a_{(t,p_4)}\}$$

Fig. 4. A transition with its pre- and postset and its corresponding connector

Example 1 continued
Let $N_1 = (P, T, F, M_0)$ be the 1-safe net from Example 1. The corresponding interaction system is $map_1(N_1) = \{\{1, \ldots, 6\}, \{A_i\}_{1 \le i \le 6}, C, \emptyset, \{T_i\}_{1 \le i \le 6}\}$, where
$C = \{\{\ a_{(p_4,t_1)}, a_{(p_5,t_1)}, a_{(t_1,p_1)}, a_{(t_1,p_2)}\}, \{a_{(p_1,t_2)}, a_{(p_2,t_2)}, a_{(t_2,p_4)}, a_{(t_2,p_5)},$
$a_{(p_3,t_2,p_3)}\}, \{a_{(p_2,t_3)}, a_{(p_3,t_3)}, a_{(t_3,p_6)}\}, \{a_{(p_6,t_4)}, a_{(t_4,p_3)}, a_{(t_4,p_2)}\}\}$

and the T_i's (and implicitely the A_i's) are given in Figure 5.

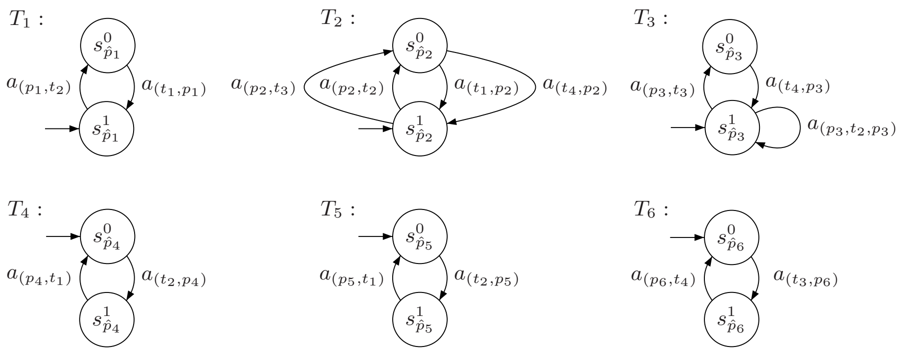

Fig. 5. The T_i's for $map_1(N_1)$

Let q be a global state of $map_1(N)$. Then $q(\hat{p})$ denotes the projection of q to $\hat{p}$.

Theorem 1. *Let N be a 1-safe net and $Sys = map_1(N)$. With $R := \{(c, t) \in (Int \times T) \mid c = c(t)\}$ and with the bijection $f : Q \to \mathcal{M}$, defined by $f(q) = \{p \in P \mid q(\hat{p}) = s^1_{\hat{p}}\}$ we have defined an isomorphism up to R for Sys and N.*

Let *Sys* be an interaction system. We consider questions for typical properties and prove them PSPACE-hard using Theorem 1 (and, of course, building on the evident fact that map_1 can be determined in polynomial time):

Corollary 1. *The question, whether some state q can be reached in Sys is PSPACE-hard.*

We know that the reachability question for 1-safe nets, i.e. the question whether some marking M is reachable in N is PSPACE-hard [CEP93].

By Theorem 1 we know that this question can be answered by answering instead the reachability question for M's corresponding global state $f^{-1}(M)$ in *Sys*.

Corollary 2. *The question, whether Sys is free of global deadlock is PSPACE-hard.*

We know that the question of deadlock for 1-safe nets, i.e. the question whether there is a reachable marking M in N where no transition is enabled is PSPACE-hard [CEP93]. By Theorem 1 we may conclude that this is the case iff there is a reachable global state in $map_1(N)$, where no interaction is enabled.

Corollary 3. *The question, whether a component $i \in K$ is live is PSPACE-hard.*

We know that the question of liveness for 1-safe nets, i.e. the question whether every transition can always occur again is PSPACE-hard [CEP93].

As liveness in 1-safe nets concerns transitions, which are translated to interactions, and, in contrast, liveness in interaction systems concerns components, we introduce a place p_t for each transition t, such that the place's corresponding component $\hat{p}_t$ will be live iff t can always occur again. This can be done by employing a (polynomial) preencoding map_{pre} on N before applying map_1.

More formally, let $map_{pre}(N) = (P \cup \{p_t \mid t \in T\}, T, F \cup \{(p_t, t), (t, p_t) \mid t \in T\}, M_0 \cup \{p_t \mid t \in T\})$.

Now N is live iff every $\hat{p}_t$ $(t \in T)$ is live in $map_1(map_{pre}(N))$.

4 Translating Interaction Systems to 1-Safe Nets

In this section, we present the encoding map_2 from interaction systems to 1-safe nets. Our interest in such a translation is mainly of theoretic nature, i.e. we want to gain more understanding of the properties of these two models. Still, as interaction models are a relatively young model, for which so far not many tools have been developed, there is some practical benefit: One could translate a system into a net and apply Petri net tools in order to investigate some behavioral questions of the system.

Let $Sys = (K, \{A_i\}_{i \in K}, C, Comp, \{T_i\}_{i \in K})$ be an interaction system. We introduce a place $\hat{q}_i$ for each local state $q_i \in Q_i$ of a component $i \in K$. A global state of *Sys* is a tuple of the present local states of the components, so for every reachable state in N, there will always be exactly one place $\hat{q}_i$ for each $i \in K$ that contains a token. This reflects that q_i is the present state of component i.

It remains to translate the glue code given by the interactions *Int* to the notion of transition. An action a_i in A_i may occur multiple times in the local transition system T_i of component i. Thus the performance of an interaction α may cause differnet state changes in Sys.

As a consequence we are going to map an interaction α not to a single transition but to a set of transitions $T(\alpha)$. Each transition in $T(\alpha)$ represents one of these possible global state changes and will shift the tokens in N according to the local state changes that are caused for the components that participate in α.

More formally, we define the mapping map_2 from interaction systems to 1-safe nets as follows:
$map_2(Sys) = (P, T, F, M_0)$, where
$P = \bigcup_{i \in K} \{\hat{q}_i \mid q_i \in Q_i\}$.

For $\alpha = \{a_{i_1}, a_{i_2}, \ldots, a_{i_k}\} \in Int$, we introduce a set of transitions $T(\alpha) := \{\{(q_{i_1}, a_{i_1}, q'_{i_1}), \ldots, (q_{i_k}, a_{i_k}, q'_{i_k})\} \mid \forall 1 \leq j \leq k (q_{i_j}, a_{i_j}, q'_{i_j}) \in \rightarrow_{i_j}\}$.

Then we define $T = \bigcup_{\alpha \in Int} T(\alpha)$.

For each α and each transition $t = \{(q_{i_1}, a_{i_1}, q'_{i_1}), \ldots, (q_{i_k}, a_{i_k}, q'_{i_k})\}$ in $T(\alpha)$ we introduce arcs as follows:
$F(t) = \{(\hat{q}_{i_1}, t), \ldots, (\hat{q}_{i_k}, t)\} \cup \{(t, \hat{q}'_{i_1}), \ldots, (t, \hat{q}'_{i_k})\}$
$F(\alpha) = \bigcup_{t \in T(\alpha)} F(t)$.
$F = \bigcup_{\alpha \in Int} F(\alpha)$.
$M_0 = \{\hat{q}_i \in P \mid q_i = q_i^0\}$.

This means that in the initital marking exactly those places that correspond to the local starting states of the components contain a token.

Remark: Let T_i be the local labeled transition system of component i and let $a_i \in A_i$ be an action of i. We denote the number of ***occurences*** of a_i in T_i by $occ(a_i)$. Note that for one interaction $\alpha = \{a_{i_1}, \ldots, a_{i_k}\}$ there are $(occ(a_{i_1}) \cdot \ldots \cdot occ(a_{i_k}))$ instances of α. This means we might have exponentially (in n) many instances for a single interaction α, which will result in an exponential blowup in our mapping from interaction systems to 1-safe nets. (See, e.g. Example 2, where we would gain $occ(a_1) \cdot occ(a_2) \cdot occ(a_3) = 2 \cdot 1 \cdot 2 = 4$ transitions of the interaction $\{a_1, a_2, a_3\}$ in $T(\{a_1, a_2, a_3\})$.)

Theorem 2. *Let Sys be an interaction system and $N = map_2(Sys)$. With $R := \{(\alpha, t) \in (Int \times T) \mid t \in T(\alpha)\}$ and with the bijection $f : Q \rightarrow \mathcal{M}$, defined by $f(q_1, \ldots, q_n) = \{\hat{q}_1, \ldots, \hat{q}_n\}$ we have defined an isomorphism up to R for Sys and N.*

Remark: One application of our translation of interaction systems to Petri nets is to answer behavioral questions for an interaction system *Sys* by translating it to a 1-safe net and answering the (corresponding) question there. Also the translation preserves component identity, i.e. a component i is represented in $map_2(Sys)$ exactly by the places $\{\hat{q}_i \mid q_i \in Q_i\}$.

5 Conclusion and Related Work

Interaction systems are a model for component-based systems. The increasing relevance of interaction systems demands a profound theoretical basis for this model. In this paper we study complexity results for interaction systems. We do so by establishing a relation between the model of interaction systems and the well-studied model of 1-safe Petri nets for which complexity results have been investigated in [CEP93]. We show that anything described by a 1-safe net can easily be described by an interaction system without a blowup in notation. Similarly, interaction systems can be translated into 1-safe nets. However, it seems unavoidable to have a (worst case) exponential blowup for this translation.

The results with the greatest impact are that the problems of deadlock-freeness and reachability are PSPACE-hard for interaction systems. These are the first PSPACE-hardness results concerning interaction systems and they partially outrun the complexity results given in [Min07]. The established results provide an essential basis for future work: Given these "master"-reductions we may extend the PSPACE-hardness results (by polynomial reductions) to almost all behavioral questions for interaction systems.

Furthermore these results suggest that there is no polynomial time algorithm for solving the questions of deadlock, reachability or liveness in interaction systems and thus provide further motivation for approaches to establish desired properties: e.g. finding sufficient conditions for deadlock-freeness and other properties of interaction systems such as the ones given in [MMM07a] and [GGM+07a] that can be tested in polynomial time or methods making use of compositionality. The model of interaction systems is particularly suited for applying these approaches because they exploit local information about components, whose identities are preserved when composing the interaction system. In contrast to this Petri nets lack compositionality and the identity of a component is lost when a composite system is modeled by a Petri net.

References

[BBS06] Basu, A., Bozga, M., Sifakis, J.: Modeling Heterogeneous Real-time Components in BIP. In: SEFM 2006. Proceedings of the Fourth IEEE International Conference on Software Engineering and Formal Methods, pp. 3–12. IEEE Computer Society, Washington, DC, USA (2006)

[BCSM07] Bozga, M., Constant, O., Skipper, M., Ma, Q.: Speeds Meta-Model Syntax and Static Semantics (2007)

[BMP+07] Basu, A., Mounier, L., Poulhis, M., Pulou, J., Sifakis, J.: Using BIP for Modeling and Verification of Networked Systems - A Case-Study on Tinyos-Based Networks. Technical Report, Verimag, Centre Équation (2007)

[BS07] Bliudze, S., Sifakis, J.: The algebra of connectors: structuring interaction in bip. In: EMSOFT 2007. Proceedings of the 7th ACM & IEEE international conference on Embedded software, pp. 11–20. ACM, New York (2007)

[CEP93] Cheng, A., Esparza, J., Palsberg, J.: Complexity Results for 1-safe Nets. In: Shyamasundar, R.K. (ed.) FSTTCS 1993. LNCS, vol. 761, pp. 326–337. Springer, Heidelberg (1993)

[dAH01] de Alfaro, L., Henzinger, T.: Interface automata. In: Proceedings of FSE 2001, ACM Press, New York (2001)

[GGM+07a] Goessler, G., Graf, S., Majster-Cederbaum, M., Martens, M., Sifakis, J.: An Approach to Modelling and Verification of Component Based Systems. In: van Leeuwen, J., Italiano, G.F., van der Hoek, W., Meinel, C., Sack, H., Plášil, F. (eds.) SOFSEM 2007. LNCS, vol. 4362, Springer, Heidelberg (2007)

[GGM+07b] Goessler, G., Graf, S., Majster -Cederbaum, M., Martens, M., Sifakis, J.: Ensuring Properties of Interaction Systems by Construction. In: Program Analysis and Compilation, Theory and Practice. LNCS, vol. 4444, Springer, Heidelberg (2007)

[GQ07] Graf, S., Quinton, S.: Contracts for BIP: Hierarchical Interaction Models for Compositional Verification. In: FORTE 2007. LNCS, vol. 4574, pp. 1–18. Springer, Heidelberg (2007)

[GS03] Goessler, G., Sifakis, J.: Component-based Construction of Deadlock-free Systems. In: Pandya, P.K., Radhakrishnan, J. (eds.) FST TCS 2003. LNCS, vol. 2914, pp. 420–433. Springer, Heidelberg (2003)

[GS05] Goessler, G., Sifakis, J.: Composition for Component-based Modeling. Sci. Comput. Program. 55(1-3), 161–183 (2005)

[LT89] Lynch, N.A., Tuttle, M.R.: An Introduction to Input/Output Automata. In: CWI-Quarterly, pp. 219–246 (1989)

[Min07] Minnameier, C.: Local and Global Deadlock-Detection in Component-based Systems are NP-hard. In: Information Processing Letters 3630 (2007)

[MMM07a] Majster-Cederbaum, M., Martens, M., Minnameier, C.: A Polynomial-time Checkable Sufficient Condition for Deadlock-Freedom of Component-based Systems. In: van Leeuwen, J., Italiano, G.F., van der Hoek, W., Meinel, C., Sack, H., Plášil, F. (eds.) SOFSEM 2007. LNCS, vol. 4362, Springer, Heidelberg (2007)

[MMM07b] Majster-Cederbaum, M., Martens, M., Minnameier, C.: Liveness in Interaction Systems. In: Proceedings of FACS 2007, ENTCS (2007)

[MSW07] Majster-Cederbaum, M., Semmelrock, N., Wolf, V.: Interaction Models for Biochemical Reactions. In: BioComp (2007)

[Sif04] Sifakis, J.: Modeling Real-time Systems. In: Keynote talk RTSS 2004 (2004)

[Sif05] Sifakis, J.: A Framework for Component-based Construction (extended abstract). In: SEFM, pp. 293–300 (2005)