Restricting Generalised State Machines

Frank Heitmann and Michael Köhler-Bußmeier

University of Hamburg, Department for Informatics
Vogt-Kölln-Straße 30, D-22527 Hamburg
{heitmann,koehler}@informatik.uni-hamburg.de

Abstract. Generalised state machines (GSMs) are a restriction of elementary object nets (Eos) a Petri net formalism which again uses Petri nets as tokens. Using GSMs or Eos many applications can be modelled which involve the mobility of e.g. active objects or agents.
To understand GSMs better we investigate the complexity of the reachability problem in certain restrictions here. We show that the problem is solvable in polynomial time for a strongly deterministic GSM where the system net and all object nets are marked graphs and that dropping the restrictions only slightly already makes the problem NP-hard.

1 Introduction

Elementary object systems (Eos for short), are Petri nets in which tokens may be Petri nets again. Originally proposed by Valk [14, 15] for a two levelled structure, the formalism was later generalised in [5, 6] for arbitrary nesting structures.[1]

Even if restricted to a depth of two, Eos are already Turing complete and thus many problems like reachability are undecidable [4] for them.

In most cases, however, the modelling power of elementary net systems is not needed. Generalised state machines (GSM) are a subclass in which the duplication or destruction of object nets is not allowed and is thus nicely suited to model physical entities.

To understand GSMs and their complexity better we focus on deterministic and strongly deterministic GSMs in this paper and investigate the complexity of the reachability problem if we restrict the allowed structure of the object nets or the system net.

After presenting elementary object systems and generalised state machines in Section 2, we show in Section 3 that the reachability problem for strongly deterministic GSMs is solvable in polynomial time, if all object nets and the system net are marked graphs, i.e. if $|^\bullet p| = |p^\bullet| = 1$ holds for all places. In Section 4 we show that by dropping the restrictions slightly the reachability problem becomes NP-hard. In the Outlook we discuss the boundary between these cases and the associated open problems.

In the following we assume basic knowledge of Petri nets, see e.g. [13].

[1] Many related approaches like recursive nets [2], nested nets [11], adaptive workflow nets [12], Mobile Systems [10], and many others are known. See [8] for a detailed discussion.

CONCURRENCY, SPECIFICATION AND PROGRAMMING
M. Szczuka et al. (eds.): Proceedings of the international workshop CS&P 2011
September 28-30, Pułtusk, Poland, pp. 194-205

2 Fundamentals

An elementary object system (EOS) is composed of a system net, which is a p/t net $\widehat{N} = (\widehat{P}, \widehat{T}, \mathbf{pre}, \mathbf{post})$, and a set of object nets $\mathcal{N} = \{N_1, \ldots, N_n\}$, which are p/t nets given as $N_i = (P_{N_i}, T_{N_i}, \mathbf{pre}_{N_i}, \mathbf{post}_{N_i})$. We assume $\widehat{N} \notin \mathcal{N}$ and the existence of the object net $N_\bullet \in \mathcal{N}$ which has no places or transitions and is used to model black tokens. Moreover we assume that all sets of nodes (places and transitions) are pairwise disjoint and set $P_\mathcal{N} = \cup_{N\in\mathcal{N}} P_N$ and $T_\mathcal{N} = \cup_{N\in\mathcal{N}} T_N$. The system net places are typed by the mapping $d : \widehat{P} \to \mathcal{N}$ with the meaning, that if $d(\widehat{p}) = N$, then the place $\widehat{p}$ of the system net may contain only net-tokens of the object net type N. The transitions in an EOS are labelled with synchronisation channels by the synchronisation labelling l. For this we assume a fixed set of channels C. In addition we allow the label τ which is used to describe that no synchronisation is desired (i.e. autonomous firing). The synchronisation labelling is then a tuple $l = (\widehat{l}, (l_N)_{N\in\mathcal{N}})$ where $\widehat{l} : \widehat{T} \to (\mathcal{N} \to (C \cup \{\tau\}))$ and $l_N : T_N \to (C \cup \{\tau\})$ for all $N \in \mathcal{N}$. All these functions are total. The intended meaning is as follows: $l_N(t) = \tau$ means that the transition t of the object net N fires (object-)autonomously. $l_N(t) = c \neq \tau$ means that t synchronises via the channel c with the system net. $\widehat{l}(\widehat{t})(N) = \tau$ means that the system net transition $\widehat{t}$ fires independent (or autonomous) from the object net N. $\widehat{l}(\widehat{t})(N) = c \neq \tau$ means that $\widehat{t}$ synchronises via the channel c with the object net N. In case of a synchronous event the system net and the object net transitions have to be labelled with the same channel. A system net transition $\widehat{t}$ fires system-autonomously, if $\widehat{l}(\widehat{t})(N) = \tau$ for all $N \in \mathcal{N}$.

A *marking* of an EOS is a *nested* multiset, denoted $\mu = \sum_{k=1}^{n} \widehat{p}_k[M_k]$, where $\widehat{p}_k$ is a place in the system net and M_k is the marking of the net-token of type $d(\widehat{p}_k)$. The set of all markings is denoted $\mathcal{M}$. We define the partial order $\leq$ on nested multisets by setting $\mu_1 \leq \mu_2$ iff $\exists \mu : \mu_2 = \mu_1 + \mu$.

$\Pi^1(\mu)$ denotes the projection of the nested marking μ to the system net level and $\Pi^2_N(\mu)$ denotes the projection to the marking belonging to the object net N, i.e. $\Pi^1(\sum_{k=1}^{n} \widehat{p}_k[M_k]) = \sum_{k=1}^{n} \widehat{p}_k$ and $\Pi^2_N(\sum_{k=1}^{n} \widehat{p}_k[M_k]) = \sum_{k=1}^{n} \mathbf{1}_N(\widehat{p}_k) \cdot M_k$, where $\mathbf{1}_N : \widehat{P} \to \{0, 1\}$ with $\mathbf{1}_N(\widehat{p}) = 1$ iff $d(\widehat{p}) = N$.

Definition 1 (EOS). *An elementary object system (*EOS*) is a tuple $OS = (\widehat{N}, \mathcal{N}, d, l)$ such that:*

1. *$\widehat{N}$ is a p/t net, called the* system net.
2. *$\mathcal{N}$ is a finite set of disjoint p/t nets, called* object nets.
3. *$d : \widehat{P} \to \mathcal{N}$ is the typing of the system net places.*
4. *$l = (\widehat{l}, (l_N)_{N\in\mathcal{N}})$ is the labelling.*

An EOS *with initial marking is a tuple $OS = (\widehat{N}, \mathcal{N}, d, l, \mu_0)$ where $\mu_0 \in \mathcal{M}$ is the initial marking.*

The synchronisation labelling generates the set of system events Θ, which consists of the disjoint sets of synchronous events Θ_l, object-autonomous events

Θ_o, and system-autonomous events Θ_s. An event is a pair, denoted $\widehat{t}[\vartheta]$ in the following, where $\widehat{t}$ is a transition of the system net or $\widehat{\epsilon}$ if object-autonomous firing is desired and ϑ maps each object net to one of its transitions or to ϵ if no firing is desired in this object net, that is $\vartheta : \mathcal{N} \to T_\mathcal{N} \cup \{\epsilon\}$ where $\vartheta(N) \neq \epsilon$ implies $\vartheta(N) \in T_N$ for all $N \in \mathcal{N}$. Is $\vartheta(N) = \epsilon$ for all N the system net transition fires autonomously. We also use the shortcut ϑ_ϵ for this function. The labelling functions are extended to $l_N(\epsilon) = \tau$ and $\widehat{l}(\widehat{\epsilon})(N) = \tau$ for all $N \in \mathcal{N}$.

We now distinguish three cases: For a synchronous event $\widehat{t}[\vartheta] \in \Theta_l$, the system net transition $\widehat{t} \neq \widehat{\epsilon}$, fires synchronously with all the object net transitions $\vartheta(N), N \in \mathcal{N}$. Thus at least one $N \in \mathcal{N}$ must exist with $\widehat{l}(\widehat{t})(N) \neq \tau$ and $\vartheta(N) \neq \epsilon$. We demand $\vartheta(N) \neq \epsilon \Leftrightarrow \widehat{l}(\widehat{t})(N) \neq \tau$ and that the labels have to match, i.e. $\widehat{l}(\widehat{t})(N) = l_N(\vartheta(N))$ for all $N \in \mathcal{N}$. Note that for object nets which do not participate in the event (either because they are not in the preset of the system net transition or because no object net transitions fires synchronously) $\widehat{l}(\widehat{t})(N) = \tau$ holds, which forces $\vartheta(N) = \epsilon$ and thus $l_N(\vartheta(N)) = l_N(\epsilon) = \tau = \widehat{l}(\widehat{t})(N)$.

In the case of a system-autonomous event $\widehat{t}[\vartheta] \in \Theta_s$, $\widehat{t} \neq \widehat{\epsilon}$ fires autonomously. Therefore we demand that $\widehat{l}(\widehat{t})(N) = \tau$ for all $N \in \mathcal{N}$ and $\vartheta = \vartheta_\epsilon$, that is $\vartheta(N) = \epsilon$ for all $N \in \mathcal{N}$.[2]

In the third case of a object-autonomous event $\widehat{\epsilon}[\vartheta] \in \Theta_o$, $\vartheta(N) \neq \epsilon$ for exactly one object net N. Moreover the transition $\vartheta(N)$ must not use a channel, that is $l_N(\vartheta(N)) = \tau$ has to hold.[3]

If we write $\widehat{t}[\vartheta] \in \Theta$ in the following, this includes the possibility that the event is an system- or object-autonomous event, i.e. $\vartheta = \vartheta_\epsilon$ or $\widehat{t} = \widehat{\epsilon}$ is possible. Moreover, since the sets of transitions are all disjoint, we usually write $\widehat{t}[\vartheta(N_1), \vartheta(N_2), \ldots]$ and also skip the object nets which are mapped to ϵ, that is, we simply list the object net's transitions with which a system net transition synchronises.

Example 1. Figure 1 shows an Eos with the system net $\widehat{N}$ and the object nets $\mathcal{N} = \{N_1, N_2\}$. The system has four net-tokens: two on place p_1 and one on p_2 and p_3 each. The net-tokens on p_1 and p_2 share the same net structure, but have independent markings.

Formally we have the system net $\widehat{N} = (\widehat{P}, \widehat{T}, \mathbf{pre}, \mathbf{post})$ with the places and transitions given by $\widehat{P} = \{p_1, \ldots, p_6\}$ and $\widehat{T} = \{t\}$, the object net $N_1 = (P_1, T_1, \mathbf{pre}_1, \mathbf{post}_1)$ with $P_1 = \{a_1, b_1\}$ and $T_1 = \{t_1\}$ and the the object net $N_2 = (P_2, T_2, \mathbf{pre}_2, \mathbf{post}_2)$ with $P_2 = \{a_2, b_2, c_2\}$ and $T_2 = \{t_2\}$. The typing is given by $d(p_1) = d(p_2) = d(p_4) = N_1$ and $d(p_3) = d(p_5) = d(p_6) = N_2$.

[2] Note that this implies $\vartheta(N) \neq \epsilon \Leftrightarrow \widehat{l}(\widehat{t})(N) \neq \tau$, the equivalence we had to demand in the case above. Moreover $l_N(\vartheta(N)) = \widehat{l}(\widehat{t})(N)$ follows, too.

[3] Note that the labels match again for all N, i.e. $\widehat{l}(\widehat{t})(N) = \widehat{l}(\widehat{\epsilon})(N) = \tau = l_N(\vartheta(N))$ for all $N \in \mathcal{N}$, but the equivalence $\vartheta(N) \neq \epsilon \Leftrightarrow \widehat{l}(\widehat{t})(N) \neq \tau$ does not hold for exactly one N, namely for the N for which $\vartheta(N) \neq \epsilon$ holds. $\vartheta(N) \in T_N$ is the transition intended to fire object-autonomously.

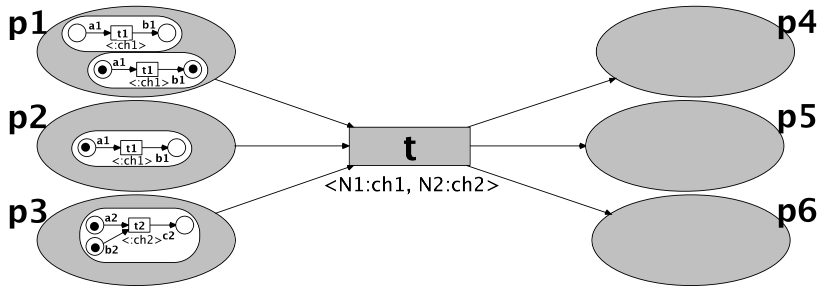

Fig. 1. An Elementary Object Net System

We have two channels ch_1 and ch_2. The labelling function $\widehat{l}$ of the system net is defined by $\widehat{l}(t)(N_1) = ch_1$ and $\widehat{l}(t)(N_2) = ch_2$. The labelling l_{N_1} of the first object net is defined by setting $l_{N_1}(t_1) = ch_1$. Similarly, l_{N_2} is defined by $l_{N_2}(t_2) = ch_2$.

There is only one (synchronous) event: $\Theta = \Theta_l = \{t[N_1 \mapsto t_1, N_2 \mapsto t_2]\}$. The event is also written shortly as $t[t_1, t_2]$.

The initial marking μ has two net-tokens on p_1, one on p_2, and one on p_3:

$$\mu = p_1[a_1 + b_1] + p_1[\mathbf{0}] + p_2[a_1] + p_3[a_2 + b_2]$$

Note that for Figure 1 the structure is the same for the three net-tokens on p_1 and p_2 but the net-tokens' markings are different.

To explain firing we distinguish two cases: Firing a system-autonomous or synchronous event $\widehat{t}[\vartheta] \in \Theta_l \cup \Theta_s$ removes net-tokens together with their individual internal markings. The new net-tokens are placed according to the system net transition and the new internal markings are determined by the internal markings just removed and ϑ. Thus a nested multiset $\lambda \in \mathcal{M}$ that is part of the current marking μ, i.e. $\lambda \leq \mu$, is replaced by a nested multiset ρ.

The enabling condition is expressed by the *enabling predicate* ϕ_{OS} (or just ϕ whenever OS is clear from the context):

$$\begin{aligned} \phi(\widehat{t}[\vartheta], \lambda, \rho) \iff & \Pi^1(\lambda) = \mathbf{pre}(\widehat{t}) \wedge \Pi^1(\rho) = \mathbf{post}(\widehat{t}) \wedge \\ & \forall N \in \mathcal{N} : \Pi^2_N(\lambda) \geq \mathbf{pre}_N(\vartheta(N)) \wedge \\ & \forall N \in \mathcal{N} : \Pi^2_N(\rho) = \Pi^2_N(\lambda) - \mathbf{pre}_N(\vartheta(N)) + \mathbf{post}_N(\vartheta(N)), \end{aligned} \tag{1}$$

where $\mathbf{pre}_N(\epsilon) = \mathbf{post}_N(\epsilon) = \mathbf{0}$ for all $N \in \mathcal{N}$.

For an object-autonomous event $\widehat{\epsilon}[\vartheta] \in \Theta_o$ let N be the object net for which $\vartheta(N) \neq \epsilon$ holds. Now $\phi(\widehat{\epsilon}[\vartheta], \lambda, \rho)$ holds iff. $\Pi^1(\lambda) = \Pi^1(\rho) = \widehat{p}$ for a $\widehat{p} \in \widehat{P}$ with $d(\widehat{p}) = N$ and $\Pi^2_N(\lambda) \geq \mathbf{pre}_N(\vartheta(N))$ and $\Pi^2_N(\rho) = \Pi^2_N(\lambda) - \mathbf{pre}_N(\vartheta(N)) + \mathbf{post}_N(\vartheta(N))$. In case of an object-autonomous event λ and ρ are thus essentially markings of an object net, but 'preceded' by a system net place typed with this object net.

Definition 2 (Firing Rule). *Let OS be an* Eos *and $\mu, \mu' \in \mathcal{M}$ markings. The event $\widehat{t}[\vartheta] \in \Theta$ is* enabled in μ for the mode $(\lambda, \rho) \in \mathcal{M}^2$ *iff $\lambda \leq \mu \wedge \phi(\widehat{t}[\vartheta], \lambda, \rho)$ holds.*

An event $\widehat{t}[\vartheta]$ that is enabled in μ for the mode (λ, ρ) can fire: $\mu \xrightarrow[OS]{\widehat{t}[\vartheta](\lambda,\rho)} \mu'$. The resulting successor marking is defined as $\mu' = \mu - \lambda + \rho$.

If the mode is not relevant we write $\mu \xrightarrow[OS]{\widehat{t}[\vartheta]} \mu'$.

We say that $\widehat{t}[\vartheta]$ is enabled in μ *or simply* active *if a mode (λ, ρ) exists such that $\widehat{t}[\vartheta]$ is enabled in μ for (λ, ρ).*

Example 2. Consider the Eos of Figure 1 again. The current marking μ of the Eos enables $t[N_1 \mapsto t_1, N_2 \mapsto t_2]$ in the mode (λ, ρ), where

$$
\begin{aligned}
\mu &= p_1[\mathbf{0}] + p_1[a_1 + b_1] + p_2[a_1] + p_3[a_2 + b_2] = p_1[\mathbf{0}] + \lambda \\
\lambda &= p_1[a_1 + b_1] + p_2[a_1] + p_3[a_2 + b_2] \\
\rho &= p_4[a_1 + b_1 + b_1] + p_5[\mathbf{0}] + p_6[c_2]
\end{aligned}
$$

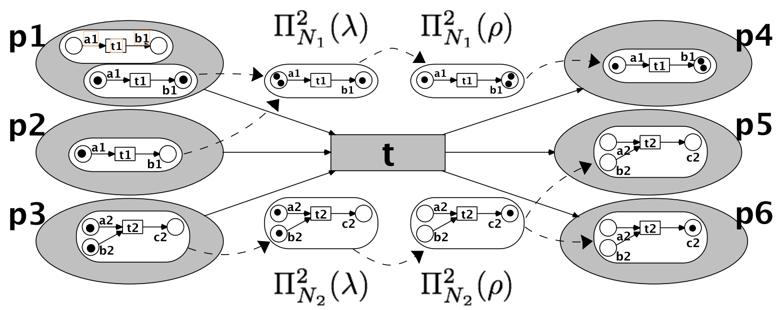

Fig. 2. The EOS of Figure 1 illustrating the projections $\Pi^2_N(\lambda)$ and $\Pi^2_N(\rho)$

The net-tokens' markings are added by the projections Π^2_N resulting in the markings $\Pi^2_N(\lambda)$. The sub-synchronisation generate $\Pi^2_N(\rho)$. (The results are shown above and below the transition t in Figure 2.) After the synchronisation we obtain the successor marking μ' with net-tokens on p_4, p_5, and p_6 as shown in Figure 2:

$$
\begin{aligned}
\mu' &= (\mu - \lambda) + \rho = p_1[\mathbf{0}] + \rho \\
&= p_1[\mathbf{0}] + p_4[a_1 + b_1 + b_1] + p_5[\mathbf{0}] + p_6[c_2]
\end{aligned}
$$

The state space of an Eos is of infinite size in general and many problems like reachability are undecidable for them. In [9] we therefore introduced four different notions of safeness for Eos. Here we concentrate on a safeness definition that guarantees finiteness of the state space.[4]

[4] In [9] this definition is called safe(3) and positive results for safe(3) Eos immediately carry over to safe(4) Eos.

Definition 3 (Safeness). *An* EOS *OS is safe iff for all reachable markings there is at most one token on each system net place and each net-token is safe:*

$$\forall \mu \in RS(OS) : \forall \widehat{p} \in \widehat{P} : \Pi^1(\mu)(\widehat{p}) \leq 1 \wedge \\ \forall N \in \mathcal{N} : \forall p \in P_N : \forall \widehat{p}[M] \leq \mu : M(p) \leq 1$$

2.1 Generalised State Machines

A *generalised state machine* (GSM) is an EOS such that every system net transition has either exactly one place in its preset and one in its postset typed with the same object net or there are no such places. Additionally the initial marking has at most one net-token of each type.

Definition 4. *Let* $G = (\widehat{N}, \mathcal{N}, d, l, \mu_0)$ *be an* EOS. *G is a* generalised state machine *(GSM) iff for all* $N \in \mathcal{N} \setminus \{N_\bullet\}$

1. $\forall \widehat{t} \in \widehat{T} : |\{\widehat{p} \in {}^\bullet\widehat{t} \mid d(\widehat{p}) = N\}| = |\{\widehat{p} \in \widehat{t}^\bullet \mid d(\widehat{p}) = N\}| \leq 1$
2. $\sum_{\widehat{p} \in \widehat{P}, d(\widehat{p}) = N} \Pi^1(\mu_0)(\widehat{p}) \leq 1$

holds.

Note that, if the second item holds at the beginning, it will hold for all reachable markings, due to the first item.

Generalised state machines where first introduced in [7] (there called *ordinary object-net systems*) where it was proven that reference and value semantics are equivalent for them.

For each GSM a p/t net, called reference net, can be easily constructed (see [7]).

The reference net $\mathrm{RN}(G)$ of a GSM G is a p/t net and is obtained by taking as set of places the disjoint union of all places of G and as set of transitions the events of G. Given a marking μ of G the projections $(\Pi^1(\mu), (\Pi^2_N(\mu))_{N \in \mathcal{N}})$ can be identified with the multiset:

$$\mathrm{RN}(\mu) := \Pi^1(\mu) + \sum_{N \in \mathcal{N}} \Pi^2_N(\mu),$$

since the places of all nets in $\mathcal{N}$ are disjoint. These are the markings in the reference net.

Definition 5. *Let* $G = (\widehat{N}, \mathcal{N}, d, l, \mu_0)$ *be a GSM.*[5] *The* reference net*, denoted by* $\mathrm{RN}(G)$*, is defined as the p/t net:*

$$\mathrm{RN}(G) = \left(\left(\widehat{P} \cup \bigcup\nolimits_{N \in \mathcal{N}} P_N\right), \Theta, \mathbf{pre}^{\mathrm{RN}}, \mathbf{post}^{\mathrm{RN}}, \mathrm{RN}(\mu_0)\right)$$

[5] The definition of a reference net for an EOS is analogously, but Theorem 1 in general only holds in one direction then. We focus on GSMs here.

where $\mathbf{pre}^{\mathrm{RN}}$ *and* $\mathbf{post}^{\mathrm{RN}}$ *are defined for an event* $\widehat{t}[\vartheta]$ *by:*

$$\mathbf{pre}^{\mathrm{RN}}(\widehat{t}[\vartheta]) = \mathbf{pre}(\widehat{t}) + \sum\nolimits_{N \in \mathcal{N}} \mathbf{pre}_N(\vartheta(N))$$
$$\mathbf{post}^{\mathrm{RN}}(\widehat{t}[\vartheta]) = \mathbf{post}(\widehat{t}) + \sum\nolimits_{N \in \mathcal{N}} \mathbf{post}_N(\vartheta(N)),$$

with $\mathbf{pre}(\widehat{\epsilon}) = \mathbf{post}(\widehat{\epsilon}) = \mathbf{0}$ *and* $\mathbf{pre}_N(\epsilon) = \mathbf{post}_N(\epsilon) = \mathbf{0}$ *for all* $N \in \mathcal{N}$.

The term *reference net* stems from the fact that $\mathrm{RN}(G)$ behaves as if each object net (in G) would have been accessed via pointers and not like a value. Since each object-net exists in a GSM at most once, the difference between references and values does not truly exist (cf. [7]). We still use the term reference net, since the above definition can also be used for EOS.

We repeat two easy to prove theorems (cf. [4] and [7]) which allow to carry over results for p/t nets to generalised state machines:

Theorem 1. *Let G be a generalised state machine. An event $\widehat{t}[\vartheta]$ is activated in G for (λ, ρ) iff it is in $\mathrm{RN}(G)$:*

$$\mu \xrightarrow[G]{\widehat{t}[\vartheta](\lambda,\rho)} \mu' \quad \iff \quad \mathrm{RN}(\mu) \xrightarrow[\mathrm{RN}(G)]{\widehat{t}[\vartheta]} \mathrm{RN}(\mu')$$

Theorem 2. *The reachability problem is decidable for generalised state machines.*

In the definition of the reference net (Definition 5) the set of events is present. Unfortunately, given a GSM $G = (\widehat{N}, \mathcal{N}, d, l, \mu_0)$ the number of events may become huge. To see this let T_i be the set of transitions of the object net N_i. Let $\widehat{l}(\widehat{t})(N_i) = c_i$ for all i and one system net transition $\widehat{t}$, where the c_i are channels. Let $l_{N_i}(t) = c_i$ for all $t \in T_i$ and all i. Now $\widehat{t}$ may fire synchronously with each transition in N_1, each in N_2 and so on. Each of these possibilities results in a different event, so we already have at least $|T_1| \cdot |T_2| \cdot \ldots \cdot |T_n|$ events, a number exponential in the number of object nets and thus in the size of the GSM. Note that this is possible for each system net transition resulting in an even larger number of events.

Lemma 1. *Let $|T| := \max\{|T_N| \mid N \in \mathcal{N}\}$ then the space needed to store Θ is in $O(|\widehat{T}| \cdot |T|^{|\mathcal{N}|} \cdot e)$, where e is the space needed to store one event $\widehat{t}[\vartheta]$, that is $e = O(|\mathcal{N}|)$.*

Given a GSM it might thus be very expensive to construct its reference net. Note that this is due to the nondeterminism introduced above by the labelling. All transitions of one object net are labelled with the same channel, so one of the transitions is chosen nondeterministically to fire synchronously with $\widehat{t}$. To prevent this, we introduced deterministic GSMs in [3]:

Definition 6. *A GSM G is called* deterministic *if for each $N \in \mathcal{N}$ the value of $l_N(t)$ differs for all $t \in T_N$ with $l_N(t) \neq \tau$. G is* strongly deterministic *if in addition for all $\widehat{t}$ and N with $\widehat{l}(\widehat{t})(N) \neq \tau$ the value of $\widehat{l}(\widehat{t})(N)$ differs.*

Thus, in a deterministic GSM each channel is used at most once in each object net. In a strongly deterministic GSM each channel is also used at most once in the system net.

On the one hand GSMs can be used to model mobility and communication of processes or agents - even if only to a smaller degree compared to Eos or general object nets. On the other hand a GSM G can be 'flatten' to the reference net $\mathrm{Rn}(G)$, which is a p/t net and which thus makes it possible to use analyses techniques for p/t nets. Since GSMs are thus as powerful as p/t nets, it is interesting to investigate if certain restrictions known for p/t nets can be transferred to GSMs and if they retain their complexity.

3 GSMs and Marked Graphs

In the following we restrict the structure of the object nets and the system net and investigate the complexity of the reachability problem. At first we severely restrict both the object nets and the system net to marked graphs, i.e. $|{}^\bullet p| = |p^\bullet| = 1$ holds for all $p \in \widehat{P} \cup P_{\mathcal{N}}$. Given a strongly deterministic GSM G, the reference net $\mathrm{Rn}(G)$ is then also a marked graph and thus the reachability problem is in P. The proof is very similar to a proof in [3], where it was proven that for a strongly deterministic and conflict-free GSM G the reference net $\mathrm{Rn}(G)$ is a conflict-free p/t net. For a strongly deterministic GSM conflict-freedom is just conflict-freedom of the system net and all object nets.

Theorem 3. *Let $G = (\widehat{N}, \mathcal{N}, d, l, \mu_0)$ be a strongly deterministic GSM with $|{}^\bullet p| = |p^\bullet| = 1$ for all $p \in \widehat{P} \cup P_{\mathcal{N}}$, i.e. the system net and all object nets interpreted as p/t nets are marked graphs. Then the p/t net $\mathrm{Rn}(G)$ is a marked graph.*

Proof. Let $p \in P(\mathrm{Rn}(G))$. Our goal is to show $|\mathbf{pre}^{\mathrm{Rn}}(p)| = |\mathbf{post}^{\mathrm{Rn}}(p)| = 1$.[6] We distinguish two cases: $p \in \widehat{P}$ or $p \in P_N$ for a $N \in \mathcal{N}$. Let $p \in \widehat{P}$. Since $|p^\bullet| = 1$ holds, p is connected with a system net transition and thus associated with an event, thus $|\mathbf{post}^{\mathrm{Rn}}(p)| \geq 1$. Assume $|\mathbf{post}^{\mathrm{Rn}}(p)| > 1$ holds. Then two events $\widehat{t}[\vartheta], \widehat{t'}[\vartheta'] \in \mathbf{post}^{\mathrm{Rn}}(p)$ with $\widehat{t}[\vartheta] \neq \widehat{t'}[\vartheta']$ exist and $p \in \mathbf{pre}(\widehat{t}) \wedge p \in \mathbf{pre}(\widehat{t'})$ follows from Definition 5 and $p \in \widehat{P}$. This implies $\widehat{t} \neq \widehat{\epsilon} \neq \widehat{t'}$, since $\mathbf{pre}(\widehat{\epsilon}) = \mathbf{0}$.

The events are thus not object-autonomous. With this and from $\widehat{t}[\vartheta] \neq \widehat{t'}[\vartheta']$, $\widehat{t} \neq \widehat{t'}$ follows. Assume otherwise. Then a net $N \in \mathcal{N}$ would have to exist with $\vartheta(N) \neq \vartheta'(N)$. But with $\vartheta(N) = t_1 \neq t_2 = \vartheta'(N)$ synchronisation would require that $l_N(t_1) = \widehat{l}(\widehat{t})(N) = l_N(t_2)$, which is not possible in a deterministic GSM, if $t_1, t_2 \in T_N$.[7] If w.l.o.g. $t_1 \in T_N$ and $t_2 = \epsilon$, then $l_N(t_2) = l_N(\epsilon) = \tau = l_N(t_1)$ would mean that t_1 fires object-autonomously in contrast to $\widehat{t'} \neq \widehat{\epsilon}$. But with $\widehat{t} \neq \widehat{t'}$ and $p \in \mathbf{pre}(\widehat{t}) \wedge p \in \mathbf{pre}(\widehat{t'})$ we have a contradiction to $|p^\bullet| = 1$!

[6] In Definition 5 $\mathbf{pre}^{\mathrm{Rn}}$ and $\mathbf{post}^{\mathrm{Rn}}$ took events as argument. Here we mean the pre- and postsets of a place p in the reference net. If we mean the pre- and postsets in the original system or object nets we write ${}^\bullet p$, resp. $p^\bullet$.

[7] Note that for $t \in T_N$, $\vartheta(N) = t$ implies $l_N(t) \neq \tau$.

The case $|\mathbf{pre}^{\mathrm{RN}}(p)| > 1$ analogously leads to a contradiction and thus $|\mathbf{pre}^{\mathrm{RN}}(p)| = |\mathbf{post}^{\mathrm{RN}}(p)| = 1$ for $p \in \widehat{P}$.

Now let $p \in P_N$ for a $N \in \mathcal{N}$. Like above $|\mathbf{post}^{\mathrm{RN}}(p)|$ is at least 1. Assume $|\mathbf{post}^{\mathrm{RN}}(p)| > 1$. Again two events $\widehat{t}[\vartheta], \widehat{t}'[\vartheta'] \in \mathbf{post}^{\mathrm{RN}}(p)$ with $\widehat{t}[\vartheta] \neq \widehat{t}'[\vartheta']$ exist. This time $p \in \mathbf{pre}_N(\vartheta(N)) \wedge p \in \mathbf{pre}_N(\vartheta'(N))$ follows from Definition 5 and $p \in P_N$.[8] This implies $\vartheta(N) \neq \epsilon \neq \vartheta'(N)$ (because $\mathbf{pre}_N(\epsilon) = \mathbf{0}$). The events are thus not system-autonomous.

The case $\vartheta(N) \neq \vartheta'(N)$ immediately implies a contradiction to $|p^\bullet| = 1$.

In the case of $\vartheta(N) = \vartheta'(N)$, neither $\widehat{t} = \widehat{t}' = \widehat{\epsilon}$, nor $\widehat{t} = \widehat{t}' \neq \widehat{\epsilon}$ is possible. In the first case both events would be equal (remember that $\vartheta(N)$ and $\vartheta'(N)$ are both not ϵ, thus N is the only object net not mapped to ϵ by ϑ, resp. ϑ', which implies that the events are equal). The second case is not possible as was already shown above. Thus $\widehat{t} \neq \widehat{t}'$ holds and only two possible cases are left: Either both events are synchronous events or one is a synchronous the other an object-autonomous event. In the first case $\widehat{l}(\widehat{t})(N) = \widehat{l}(\widehat{t}')(N)$ follows from $\widehat{l}(\widehat{t})(N) = l_N(\vartheta(N))$, $\widehat{l}(\widehat{t}')(N) = l_N(\vartheta'(N))$ and $\vartheta(N) = \vartheta(N')$, but this is not possible in a *strongly deterministic* GSM. In the second case let w.l.o.g. $\widehat{t}[\vartheta] \in \Theta_l$ and $\widehat{t}'[\vartheta'] = \widehat{\epsilon}[\vartheta'] \in \Theta_o$. But since $\vartheta(N)$ and $\vartheta'(N)$ are both not ϵ (but transitions of N), this would mean that $\vartheta(N)$ fires synchronously with $\widehat{t}$ and $\vartheta'(N)$ fires object-autonomously (i.e. is not labelled with a channel). Since $\vartheta(N) = \vartheta'(N)$, this is not possible. Thus like above (in the case $p \in \widehat{P}$) where $\widehat{t} = \widehat{t}'$ was not possible, here (in the case $p \in P_N$) $\vartheta(N) = \vartheta'(N)$ is not possible.[9]

The case $|\mathbf{pre}^{\mathrm{RN}}(p)| > 1$ is again analogously and from this $|\mathbf{pre}^{\mathrm{RN}}(p)| = |\mathbf{post}^{\mathrm{RN}}(p)| = 1$ follows for all $p \in P(\mathrm{RN}(G))$. $\mathrm{RN}(G)$ is therefore a marked graph.

Since in a (strongly) deterministic GSM the number of events is bounded by a polynomial[10] the reference net can be constructed in polynomial time. Since reachability can be decided in polynomial time for a marked graph [1], the following corollary follows:

Corollary 1. *The reachability problem for strongly deterministic GSMs with* $|{}^\bullet p| = |p^\bullet| = 1$ *for all* $p \in \widehat{P} \cup P_{\mathcal{N}}$ *is solvable in polynomial time.*

4 GSMs and NP-completeness

By dropping the condition that the GSM has to be *strongly* deterministic and only requiring a deterministic GSM the reference net must not necessarily be a

[8] Note that the object nets can not be different, since $p \in P_N$ holds.

[9] Note, that $\vartheta(N) = \vartheta'(N)$ is thus not possible and that in this part of the proof the marked graph characteristic was not used. For this reason this part is the same as in Theorem 4.8 in [3].

[10] There are at most $|\widehat{T}|$ system-autonomous and $\sum_{N \in \mathcal{N}} |T_N|$ object-autonomous events. The sum of these is a upper bound for the number of events, because each synchronous event can be thought of here as a combination of a system- and several object-autonomous events.

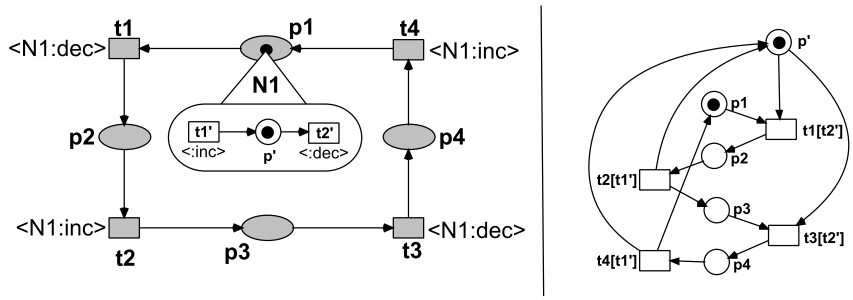

Fig. 3. A GSM G (left) with its reference net $\textsc{Rn}(G)$.

marked graph. Figure 3 shows an example. Note that the system net on the left is a deterministic GSM with $|^\bullet p| = |p^\bullet| = 1$ for all $p \in \widehat{P} \cup P_\mathcal{N}$. The reference net on the right is not a marked graph.

If we also change the condition for the system net from $|^\bullet p| = |p^\bullet| = 1$ to $|^\bullet t| = |t^\bullet| = 1$, we end up with a formalism in which the reachability problem is NP-hard.

Theorem 4. *Let $G = (\widehat{N}, \mathcal{N}, d, l, \mu_0)$ be a deterministic GSM with $|^\bullet p| = |p^\bullet| = 1$ for all $p \in P_\mathcal{N}$, and $|^\bullet t| = |t^\bullet| = 1$ for all $t \in \widehat{T}$. Moreover there is only one object net ($|\mathcal{N}| = 1$) and $|\Pi^1(\mu_0)| = 1$ holds. Then the reachability problem is NP-hard.*

Proof. We only sketch the proof here. The reduction is from Longest Path, where a directed graph $G = (V, E)$, two nodes v_1, v_2 and a number k is given and the question asked is, if a path exists from v_1 to v_2 with a length of at least k.

We have one object net consisting of a place p and one transition t_1 in p's preset and one transition t_2 in p's postset. t_1 is labelled with channel *inc*, t_2 is labelled with τ (object-autonomous firing). The idea is that the object net travels along the path and that whenever a edge is used, t_1 fires. Thus the number of tokens on p are the number of edges used.

For this we create for each node $v \in V$ a place $\widehat{p}_v$ in the system net and for each edge $e = (v, v') \in E$ a transition $\widehat{t}_e$ and two arcs $(\widehat{p}_v, \widehat{t}_e)$, $(\widehat{t}_e, \widehat{p}_{v'})$. The transitions $\widehat{t}_e$ for $e \in E$ are all labelled with the channel *inc*.

The object net is placed on the (start) node $\widehat{p}_{v_1}$.

If a path from v_1 to v_2 exists with length $m \geq k$. The object net can 'travel' along this path from $\widehat{p}_{v_1}$ to $\widehat{p}_{v_2}$ and we end up in the marking $\widehat{p}_{v_2}[m \cdot p]$. The other direction is proven similarly.

To be able to formulate this as a reachability problem (where we ask for an exact number of tokens on the place p), we need to be able to remove tokens from place p. For this the object-autonomous event $\widehat{\epsilon}[t_2]$ may remove tokens from p.

Thus given an instance (G, v_1, v_2, k) from Longest Path we construct the GSM outlined above and asked if the marking $\widehat{p}_{v_2}[k \cdot p]$ is reachable.

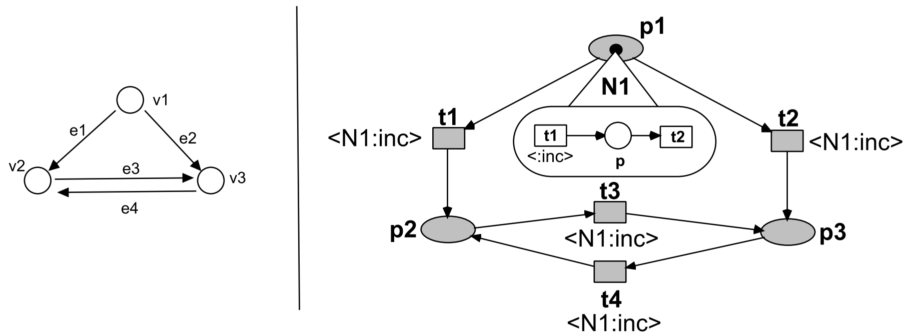

Fig. 4. Example of the construction in the proof of Theorem 4

Figure 4 shows an example. This construction is clearly possible in polynomial time and thus the problem is NP-hard.

Note that the GSM sketched above only uses one object net type and no black tokens on the system net level. Net systems with this property are called *unary*. Thus even in the case of unary GSMs dropping the requirement to be strongly deterministic to just being deterministic and changing the condition for the system net to $|{}^\bullet t| = |t^\bullet| = 1$ and $|\Pi^1(\mu_0)| = 1$, already results in a formalism where the reachability problem is NP-hard. Note that all object nets are marked graphs and that for them the reachability problem is in P. Also note that a for a p/t net with the requirements inflicted upon the system net, the reachability problem would also be in P.

5 Conclusion and Outlook

The focus of the paper was the definition and analysis of restrictions of generalised state machines. We showed that for strongly deterministic GSMs where each net is a marked graph, the reachability problem can be decided in polynomial time.

If we look at deterministic GSMs instead of strongly deterministic ones and require $|{}^\bullet t| = |t^\bullet| = 1$ for all $t \in \widehat{T}$ and $|\Pi^1(\mu_0)| = 1$ (there is only one object net), the reachability problem is already NP-hard. It is not known if this problem is also *in* NP, but we suspect so.

The complexity of the problem 'in between', that is, the reachability problem for deterministic GSMs where each object net and the system net are marked graphs, is currently unknown. One problem is, that GSMs which are only deterministic might have a huge number of different events and it is not yet clear, if it is necessary to construct them all.

References

1. Javier Esparza and Mogens Nielsen. Decidability issues for petri nets - a survey. *J. Inform. Process. Cybernet*, 30(3):143–160, 1994.
2. Serge Haddad and Denis Poitrenaud. Theoretical aspects of recursive Petri nets. In S. Donatelli and J. Kleijn, editors, *Application and Theory of Petri Nets*, volume 1639 of *Lecture Notes in Computer Science*, pages 228–247. Springer-Verlag, 1999.
3. Frank Heitmann and Michael Köhler-Bußmeier. On defining conflict-freedom for object nets. To appear in Proceedings of the International Workshop on Logic, Agents, and Mobility (LAM 2011).
4. Michael Köhler. The reachability problem for object nets. *Fundamenta Informaticae*, 79(3-4):401 – 413, 2007.
5. Michael Köhler and Heiko Rölke. Concurrency for mobile object-net systems. *Fundamenta Informaticae*, 54(2-3), 2003.
6. Michael Köhler and Heiko Rölke. Properties of Object Petri Nets. In J. Cortadella and W. Reisig, editors, *Application and Theory of Petri Nets*, volume 3099 of *Lecture Notes in Computer Science*, pages 278–297. Springer-Verlag, 2004.
7. Michael Köhler and Heiko Rölke. Reference and value semantics are equivalent for ordinary object petri nets. In G. Ciardo and P. Darondeau, editors, *Application and Theory of Petri Nets*, volume 3536 of *Lecture Notes in Computer Science*, pages 309–328. Springer-Verlag, 2005.
8. Michael Köhler-Bußmeier. Decidability results for elementary object systems. Report of the department of informatics, Universität Hamburg, Fachbereich Informatik, 2011.
9. Michael Köhler-Bußmeier and Frank Heitmann. Safeness for object nets. *Fundamenta Informaticae*, 101(1-2):29–43, 2010.
10. Charles Lakos. A Petri net view of mobility. In *Formal Techniques for Networked and Distributed Systems (FORTE 2005)*, volume 3731 of *Lecture Notes in Computer Science*, pages 174–188. Springer-Verlag, 2005.
11. Irina A. Lomazova. Nested Petri nets – a formalism for specification of multi-agent distributed systems. *Fundamenta Informaticae*, 43(1-4):195–214, 2000.
12. Irina A. Lomazova, Kees M. van Hee, Olivia Oanea, Alexander Serebrenik, Natalia Sidorova, and Marc Voorhoeve. Nested nets for adaptive systems. In *Application and Theory of Petri Nets and Other Models of Concurrency*, Lecture Notes in Computer Science, pages 241–260. Springer-Verlag, 2006.
13. Wolfgang Reisig and Grzegorz Rozenberg, editors. *Lectures on Petri Nets I: Basic Models*, volume 1491 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.
14. Rüdiger Valk. Modelling concurrency by task/flow EN systems. In *3rd Workshop on Concurrency and Compositionality*, number 191 in GMD-Studien, St. Augustin, Bonn, 1991. Gesellschaft für Mathematik und Datenverarbeitung.
15. Rüdiger Valk. Object Petri nets: Using the nets-within-nets paradigm. In Jörg Desel, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Advanced Course on Petri Nets 2003*, volume 3098 of *Lecture Notes in Computer Science*, pages 819–848. Springer-Verlag, 2003.