Complexity Studies for Safe and Fan-Bounded Elementary Hornets

Michael Köhler-Bußmeier and Frank Heitmann

michael.koehler-bussmeier@haw-hamburg.de
heitmann@informatik.uni-hamburg.de

Abstract. HORNETS are Petri nets that have nets as tokens. There are an algebraic extension of elementary object nets (EOS) with the possibility to mofify the structure of the net-tokens. In previous contributions we investigated elementary HORNETS as well as their subclass of safe elementary HORNETS. We showed that the reachability problem for safe elementary HORNETS requires at least exponential space. We have also showed that exponential space is sufficient. This shows that safe elementary HORNETS are much more complicated than safe elementary object nets (safe EOS), where reachability is known to be PSPACE-complete. In this contribution we study structural restrictions of elementary HORNETS that have a better complexity: fan-bounded HORNETS. It turns out that reachability is again in PSPACE for this class of HORNETS.

Key words: Hornets, nets-within-nets, object nets, reachability, safeness

1 Hornets: Higher-Order Object Nets

In this paper we study self-modifying systems in the formalisms of HORNETS. HORNETS are a generalisation of object nets [1, 2], which follow the *nets-within-nets* paradigm as proposed by Valk [3].

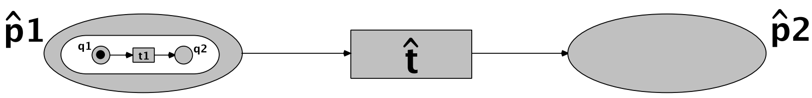

Fig. 1. An Elementary Object Net System (EOS)

With object nets we study Petri nets where the tokens are nets again, i.e. we have a nested marking. Events are also nested. We have three different kinds of events – as illustrated by the example given in Figure 1:

1. System-autonomous: The system net transition $\widehat{t}$ fires autonomously, which moves the net-token from $\widehat{p}_1$ to $\widehat{p}_2$ without changing its marking.
2. Object-autonomous: The object net fires transition t_1 "moving" the black token from q_1 to q_2. The object net remains at its location $\widehat{p}_1$.

3. Synchronisation: Whenever we add matching synchronisation inscriptions at the system net transition $\widehat{t}$ and the object net transition t_1, then both must fire synchronously: The object net is moved to $\widehat{p}_2$ and the black token moves from q_1 to q_2 inside. Whenever synchronisation is specified, autonomous actions are forbidden.

For HORNETS we extend object-nets with algebraic concepts that allow to modify the structure of the net-tokens as a result of a firing transition. This is a generalisation of the approach of algebraic nets [4], where algebraic data types replace the anonymous black tokens.

The use of algebraic operations in HORNETS relates them to *algebraic higher-order (AHO) systems* [5], which are restricted to two-levelled systems but have a greater flexibility for the operations on net-tokens, since each net transformation is allowed. There is also a relationship to Nested Nets [6], which are used for adaptive systems.

It is not hard to prove that the general HORNET formalism is Turing-complete. In [7] we have proven that there are several possibilities to simulate counter programs: One could use the nesting to encode counters. Another possibility is to encode counters in the algebraic structure of the net operators.

In our general research we like to study the *complexity* that arises due the algebraic structure. Therefore, we restrict HORNETS to guarantee that the system has a finite state space: First, we allow at most one token on each place, which results in the class of *safe* HORNETS. However this restriction does not guarantee finite state spaces, since we have the nesting depth as a second source of undecidability [2]. Second, we restrict the universe of object nets to finite sets. Finally, we restrict the nesting depth and introduce the class of *elementary* HORNETS, which have a two-levelled nesting structure. This is done in analogy to the class of *elementary object net systems (*EOS*)* [1], which are the two-level specialisation of general object nets [1, 2].

If we rule out these sources of complexity the main origin of complexity is the use of algebraic transformations, which are still allowed for safe, elementary HORNETS. As a result we obtain the class of safe, elementary HORNETS – in analogy to the class of *safe* EOS [8]. We have shown in [8–10] that most problems for *safe* EOS are PSPACE-complete. More precisely: All problems that are expressible in LTL or CTL, which includes reachability and liveness, are PSPACE-complete. This means that with respect to these problems *safe* EOS are no more complex than p/t nets. In a previous publication [11] we have shown that *safe, elementary* HORNETS are beyond PSPACE. We have shown a lower bound, i.e. that "the reachability problem requires exponential space" for safe, elementary HORNETS – similarly to well known result of for *bounded* p/t nets [12]. In [13] we give an algorithm thats needs at most exponential space, which shows that lower and upper bound coincide.

In this paper we would like to study restrictions of Elementary HORNETS to obtain net classes where the reachability requires less than exponential space. From [11] we know that the main source of complexity for EHORNETS is mainly due to the huge number of different of net-tokens, which is double-exponential for safe EHORNETS. A closer look reveals that the number of net-token's marking is rather small – "only" single-exponential, while the number of different object nets is double-exponential. We conclude that restricting the net-tokens' marking beyond safeness would not improve complexity. Instead, we have to impose *structural* restrictions on the object-nets. Petri net theory offers several well known candidates for structural restrictions, like state machines, free-choice nets etc. Here, we restrict object-nets to state-machines.

The paper has the following structure: Section 2 defines Elementary HORNETS. Since the reachability problem is known to be undecidable even for EOS, we restrict elementary HORNETS to safe ones, which have finite state spaces. State Machines have at most one place in the pre- and in the post-set. In Section 3 we generalise this notion in the way that the number of all places in pre- and postset of an object-net is below a given bound. So, we obtain the maximal synchronisation degree of the objects nets (i.e. the maximal pre- and postset size) as a fresh complexity parameter. Section 4 shows that the reachability problem is PSPACE-complete.

2 Definition of Elementary Hornets (EHORNETS)

A multiset $\mathbf{m}$ on the set D is a mapping $\mathbf{m} : D \to \mathbb{N}$. Multisets can also be represented as a formal sum in the form $\mathbf{m} = \sum_{i=1}^{n} x_i$, where $x_i \in D$.

Multiset addition is defined component-wise: $(\mathbf{m}_1 + \mathbf{m}_2)(d) := \mathbf{m}_1(d) + \mathbf{m}_2(d)$. The empty multiset $\mathbf{0}$ is defined as $\mathbf{0}(d) = 0$ for all $d \in D$. Multiset-difference $\mathbf{m}_1 - \mathbf{m}_2$ is defined by $(\mathbf{m}_1 - \mathbf{m}_2)(d) := \max(\mathbf{m}_1(d) - \mathbf{m}_2(d), 0)$.

The cardinality of a multiset is $|\mathbf{m}| := \sum_{d \in D} \mathbf{m}(d)$. A multiset $\mathbf{m}$ is finite if $|\mathbf{m}| < \infty$. The set of all finite multisets over the set D is denoted $MS(D)$.

Multiset notations are used for sets as well. The meaning will be apparent from its use.

Any mapping $f : D \to D'$ extends to a multiset-homomorphism $f^\sharp : MS(D) \to MS(D')$ by $f^\sharp\left(\sum_{i=1}^{n} x_i\right) = \sum_{i=1}^{n} f(x_i)$.

A *p/t net* N is a tuple $N = (P, T, \mathbf{pre}, \mathbf{post})$, such that P is a set of places, T is a set of transitions, with $P \cap T = \emptyset$, and $\mathbf{pre}, \mathbf{post} : T \to MS(P)$ are the pre- and post-condition functions. A marking of N is a multiset of places: $\mathbf{m} \in MS(P)$. We denote the enabling of t in marking $\mathbf{m}$ by $\mathbf{m} \xrightarrow{t}$. Firing of t is denoted by $\mathbf{m} \xrightarrow{t} \mathbf{m}'$.

Net-Algebras We define the algebraic structure of object nets. For a general introduction of algebraic specifications cf. [14].

Let K be a set of net-types (kinds). A (many-sorted) *specification* (Σ, X, E) consists of a signature Σ, a family of variables $X = (X_k)_{k \in K}$, and a family of axioms $E = (E_k)_{k \in K}$.

A signature is a disjoint family $\Sigma = (\Sigma_{k_1 \cdots k_n, k})_{k_1, \cdots, k_n, k \in K}$ of operators. The set of terms of type k over a signature Σ and variables X is denoted $\mathbb{T}^k_\Sigma(X)$.

We use (many-sorted) predicate logic, where the terms are generated by a signature Σ and formulae are defined by a family of predicates $\Psi = (\Psi_n)_{n \in \mathbb{N}}$. The set of formulae is denoted PL_Γ, where $\Gamma = (\Sigma, X, E, \Psi)$ is the *logic structure*.

Let Σ be a signature over K. A *net-algebra* assigns to each type $k \in K$ a set $\mathcal{U}_k$ of object nets – the net universe. Each object $N \in \mathcal{U}_k, k \in K$ net is a p/t net $N = (P_N, T_N, \mathbf{pre}_N, \mathbf{post}_N)$. We identify $\mathcal{U}$ with $\bigcup_{k \in K} \mathcal{U}_k$ in the following. We assume the family $\mathcal{U} = (\mathcal{U}_k)_{k \in K}$ to be disjoint.

The nodes of the object nets in $\mathcal{U}_k$ are not disjoint, since the firing rule allows to transfer tokens between net tokens within the same set $\mathcal{U}_k$. Such a transfer is possible, if we assume that all nets $N \in \mathcal{U}_k$ have the same set of places P_k. P_k is the place universe for all object nets of kind k.

In general, P_k is not finite. Since we like each object net to be finite in some sense, we require that the transitions T_N of each $N \in \mathcal{U}_k$ use only a finite subset of P_k, i.e. $\forall N \in \mathcal{U} : |{}^\bullet T_N \cup T_N{}^\bullet| < \infty$.

The family of object nets $\mathcal{U}$ is the universe of the algebra. A net-algebra $(\mathcal{U}, \mathcal{I})$ assigns to each constant $\sigma \in \Sigma_{\lambda,k}$ an object net $\sigma^\mathcal{I} \in \mathcal{U}_k$ and to each operator $\sigma \in \Sigma_{k_1\cdots k_n,k}$ with $n > 0$ a mapping $\sigma^\mathcal{I} : (\mathcal{U}_{k_1} \times \cdots \times \mathcal{U}_{k_n}) \to \mathcal{U}_k$.

A net-algebra is called *finite* if P_k is a finite set for each $k \in K$.

Since all nets $N \in \mathcal{U}_k$ have the same set of places P_k, which is finite for EHORNETS, there is an upper bound for the cardinality of $\mathcal{U}_k$.

Proposition 1 (Lemma 2.1 in [11]). *For each $k \in K$ the cardinality of each net universe $\mathcal{U}_k$ is bound as follows: $|\mathcal{U}_k| \leq 2^{\left(2^{4|P_k|}\right)}$.*

A variable assignment $\alpha = (\alpha_k : X_k \to \mathcal{U}_k)_{k \in K}$ maps each variable onto an element of the algebra. For a variable assignment α the evaluation of a term $t \in \mathbb{T}^k_\Sigma(X)$ is uniquely defined and will be denoted as $\alpha(t)$.

A net-algebra, such that all axioms of (Σ, X, E) are valid, is called *net-theory*.

Nested Markings A marking of an EHORNET assigns to each system net place one or many net-tokens. The places of the system net are typed by the function $k : \widehat{P} \to K$, meaning that a place $\widehat{p}$ contains net-tokens of kind $k(\widehat{p})$. Since the net-tokens are instances of object nets, a *marking* is a *nested* multiset of the form:

$$\mu = \sum_{i=1}^{n} \widehat{p}_i[N_i, M_i] \quad \text{where} \quad \widehat{p}_i \in \widehat{P}, N_i \in \mathcal{U}_{k(\widehat{p}_i)}, M_i \in MS(P_{N_i}), n \in \mathbb{N}$$

Each addend $\widehat{p}_i[N_i, M_i]$ denotes a net-token on the place $\widehat{p}_i$ that has the structure of the object net N_i and the marking $M_i \in MS(P_{N_i})$. The set of all nested multisets is denoted as $\mathcal{M}_H$. We define the partial order $\sqsubseteq$ on nested multisets by setting $\mu_1 \sqsubseteq \mu_2$ iff $\exists \mu : \mu_2 = \mu_1 + \mu$.

The projection $\Pi^{1,H}_N(\mu)$ is the multiset of all system-net places that contain the object-net N:[1]

$$\Pi^{1,H}_N\left(\sum\nolimits_{i=1}^{n} \widehat{p}_i[N_i, M_i]\right) := \sum\nolimits_{i=1}^{n} \mathbf{1}_N(N_i) \cdot \widehat{p}_i \tag{1}$$

where the indicator function $\mathbf{1}_N$ is defined as: $\mathbf{1}_N(N_i) = 1$ iff $N_i = N$.

Analogously, the projection $\Pi^{2,H}_N(\mu)$ is the multiset of all net-tokens' markings (that belong to the object-net N):

$$\Pi^{2,H}_N\left(\sum\nolimits_{i=1}^{n} \widehat{p}_i[N_i, M_i]\right) := \sum\nolimits_{i=1}^{n} \mathbf{1}_k(N_i) \cdot M_i \tag{2}$$

The projection $\Pi^{2,H}_k(\mu)$ is the sum of all net-tokens' markings belonging to the same type $k \in K$:

$$\Pi^{2,H}_k(\mu) := \sum\nolimits_{N \in \mathcal{U}_k} \Pi^{2,H}_N(\mu) \tag{3}$$

[1] The superscript H indicates that the function is used for HORNETS.

Synchronisation The transitions in an HORNET are labelled with synchronisation inscriptions. We assume a fixed set of channels $C = (C_k)_{k\in K}$.

– The function family $\widehat{l}_\alpha = (\widehat{l}^k_\alpha)_{k\in K}$ defines the synchronisation constraints. Each transition of the system net is labelled with a multiset $\widehat{l}^k(\widehat{t}) = (e_1, c_1) + \cdots + (e_n, c_n)$, where the expression $e_i \in \mathbb{T}^k_\Sigma(X)$ describes the called object net and $c_i \in C_k$ is a channel. The intention is that $\widehat{t}$ fires synchronously with a multiset of object net transitions with the same multiset of labels. Each variable assignment α generates the function $\widehat{l}^k_\alpha(\widehat{t})$ defined as:

$$\widehat{l}^k_\alpha(\widehat{t})(N) := \sum\nolimits_{\substack{1\le i\le n\\ \alpha(e_i)=N}} c_i \quad \text{for} \quad \widehat{l}^k(\widehat{t}) = \sum\nolimits_{1\le i\le n}(e_i, c_i) \tag{4}$$

Each function $\widehat{l}^k_\alpha(\widehat{t})$ assigns to each object net N a multiset of channels.
– For each $N \in \mathcal{U}_k$ the function l_N assigns to each transition $t \in T_N$ either a channel $c \in C_k$ or $\bot_k$, whenever t fires without synchronisation, i.e. autonomously.

System Net Assume we have a fixed logic $\Gamma = (\Sigma, X, E, \Psi)$ and a net-theory $(\mathcal{U}, \mathcal{I})$. An *elementary higher-order object net* (EHORNET) is composed of a system net $\widehat{N}$ and the set of object nets $\mathcal{U}$. W.l.o.g. we assume $\widehat{N} \notin \mathcal{U}$. To guarantee finite algebras for EHORNETS, we require that the net-theory $(\mathcal{U}, \mathcal{I})$ is finite, i.e. each place universe P_k is finite.

The system net is a net $\widehat{N} = (\widehat{P}, \widehat{T}, \mathbf{pre}, \mathbf{post}, \widehat{G})$, where each arc is labelled with a multiset of terms: $\mathbf{pre}, \mathbf{post} : \widehat{T} \to (\widehat{P} \to MS(\mathbb{T}_\Sigma(X)))$. Each transition is labelled by a guard predicate $\widehat{G} : \widehat{T} \to PL_\Gamma$. The places of the system net are typed by the function $k : \widehat{P} \to K$. As a typing constraint we have that each arc inscription has to be a multiset of terms that are all of the kind that is assigned to the arc's place:

$$\mathbf{pre}(\widehat{t})(\widehat{p}), \quad \mathbf{post}(\widehat{t})(\widehat{p}) \quad \in \quad MS(\mathbb{T}^{k(\widehat{p})}_\Sigma(X)) \tag{5}$$

For each variable binding α we obtain the evaluated functions $\mathbf{pre}_\alpha, \mathbf{post}_\alpha : \widehat{T} \to (\widehat{P} \to MS(\mathcal{U}))$ in the obvious way.

Definition 1 (Elementary Hornet, EHORNET). *Assume a fixed many-sorted predicate logic $\Gamma = (\Sigma, X, E, \Psi)$.*

An elementary HORNET *is a tuple $EH = (\widehat{N}, \mathcal{U}, \mathcal{I}, k, l, \mu_0)$ such that:*

1. *$\widehat{N}$ is an algebraic net, called the* system net.
2. *$(\mathcal{U}, \mathcal{I})$ is a finite net-theory for the logic Γ.*
3. *$k : \widehat{P} \to K$ is the typing of the system net places.*
4. *$l = (\widehat{l}, l_N)_{N\in\mathcal{U}}$ is the labelling.*
5. *$\mu_0 \in \mathcal{M}_H$ is the initial marking.*

Events The synchronisation labelling generates the set of system events Θ. We have three kinds of events:

1. Synchronised firing: There is at least one object net that has to be synchronised, i.e. there is a N such that $\widehat{l}(\widehat{t})(N)$ is not empty.
   Such an event is a pair $\theta = \widehat{t}^{\alpha}[\vartheta]$, where $\widehat{t}$ is a system net transition, α is a variable binding, and ϑ is a function that maps each object net to a multiset of its transitions, i.e. $\vartheta(N) \in MS(T_N)$. It is required that $\widehat{t}$ and $\vartheta(N)$ have matching multisets of labels, i.e. $\widehat{l}(\widehat{t})(N) = l_N^{\sharp}(\vartheta(N))$ for all $N \in \mathcal{U}$. (Remember that $l_N^{\sharp}$ denotes the multiset extension of l_N.)
   The intended meaning is that $\widehat{t}$ fires synchronously with all the object net transitions $\vartheta(N), N \in \mathcal{U}$.
2. System-autonomous firing: The transition $\widehat{t}$ of the system net fires autonomously, whenever $\widehat{l}(\widehat{t})$ is the empty multiset $\mathbf{0}$.
   We consider system-autonomous firing as a special case of synchronised firing generated by the function ϑ_{id}, defined as $\vartheta_{id}(N) = \mathbf{0}$ for all $N \in \mathcal{U}$.
3. Object-autonomous firing: An object net transition t in N fires autonomously, whenever $l_N(t) = \bot_k$.
   Object-autonomous events are denoted as $id_{\widehat{p},N}[\vartheta_t]$, where $\vartheta_t(N') = \{t\}$ if $N = N'$ and $\mathbf{0}$ otherwise. The meaning is that in object net N fires t autonomously within the place $\widehat{p}$.
   For the sake of uniformity we define for an arbitrary binding α:

$$\mathbf{pre}_{\alpha}(id_{\widehat{p},N})(\widehat{p}')(N') = \mathbf{post}_{\alpha}(id_{\widehat{p},N})(\widehat{p}')(N') = \begin{cases} 1 & \text{if } \widehat{p}' = \widehat{p} \wedge N' = N \\ 0 & \text{otherwise.} \end{cases}$$

The set of all *events* generated by the labelling l is $\Theta_l := \Theta_1 \cup \Theta_2$, where Θ_1 contains synchronous events (including system-autonomous events as a special case) and Θ_2 contains the object-autonomous events:

$$\begin{aligned} \Theta_1 &:= \left\{ \widehat{\tau}^{\alpha}[\vartheta] \quad \mid \forall N \in \mathcal{U} : \widehat{l}_{\alpha}(\widehat{t})(N) = l_N^{\sharp}(\vartheta(N)) \right\} \\ \Theta_2 &:= \left\{ id_{\widehat{p},N}[\vartheta_t] \mid \widehat{p} \in \widehat{P}, N \in \mathcal{U}_{k(\widehat{p})}, t \in T_N \right\} \end{aligned} \tag{6}$$

Firing Rule A system event $\theta = \widehat{\tau}^{\alpha}[\vartheta]$ removes net-tokens together with their individual internal markings. Firing the event replaces a nested multiset $\lambda \in \mathcal{M}_H$ that is part of the current marking μ, i.e. $\lambda \sqsubseteq \mu$, by the nested multiset ρ. The enabling condition is expressed by the *enabling predicate* ϕ_{EH} (or just ϕ whenever EH is clear from the context):

$$\begin{aligned} \phi_{EH}(\widehat{\tau}^{\alpha}[\vartheta], \lambda, \rho) \leftrightarrow \forall k \in K : & \\ & \forall \widehat{p} \in k^{-1}(k) : \forall N \in \mathcal{U}_k : \Pi_N^{1,H}(\lambda)(\widehat{p}) = \mathbf{pre}_{\alpha}(\widehat{\tau})(\widehat{p})(N) \wedge \\ & \forall \widehat{p} \in k^{-1}(k) : \forall N \in \mathcal{U}_k : \Pi_N^{1,H}(\rho)(\widehat{p}) = \mathbf{post}_{\alpha}(\widehat{\tau})(\widehat{p})(N) \wedge \\ & \Pi_k^{2,H}(\lambda) \geq \textstyle\sum_{N \in \mathcal{U}_k} \mathbf{pre}_N^{\sharp}(\vartheta(N)) \wedge \\ & \Pi_k^{2,H}(\rho) = \Pi_k^{2,H}(\lambda) + \textstyle\sum_{N \in \mathcal{U}_k} \mathbf{post}_N^{\sharp}(\vartheta(N)) - \mathbf{pre}_N^{\sharp}(\vartheta(N)) \end{aligned} \tag{7}$$

The predicate ϕ_{EH} has the following meaning: Conjunct (1) states that the removed sub-marking λ contains on $\widehat{p}$ the right number of net-tokens, that are removed by $\widehat{\tau}$.

Conjunct (2) states that generated sub-marking ρ contains on $\widehat{p}$ the right number of net-tokens, that are generated by $\widehat{\tau}$. Conjunct (3) states that the sub-marking λ enables all synchronised transitions $\vartheta(N)$ in the object N. Conjunct (4) states that the marking of each object net N is changed according to the firing of the synchronised transitions $\vartheta(N)$.

Note, that conjunct (1) and (2) assures that only net-tokens relevant for the firing are included in λ and ρ. Conditions (3) and (4) allow for additional tokens in the net-tokens.

For system-autonomous events $\widehat{t}^{\alpha}[\vartheta_{id}]$ the enabling predicate ϕ_{EH} can be simplified further: Conjunct (3) is always true since $\mathbf{pre}_N(\vartheta_{id}(N)) = \mathbf{0}$. Conjunct (4) simplifies to $\Pi_k^{2,H}(\rho) = \Pi_k^{2,H}(\lambda)$, which means that no token of the object nets get lost when a system-autonomous events fires.

Analogously, for an object-autonomous event $\widehat{\tau}[\vartheta_t]$ we have an idle-transition $\widehat{\tau} = id_{\widehat{p},N}$ and $\vartheta = \vartheta_t$ for some t. Conjunct (1) and (2) simplify to $\Pi_{N'}^{1,H}(\lambda) = \widehat{p} = \Pi_{N'}^{1,H}(\rho)$ for $N' = N$ and to $\Pi_{N'}^{1,H}(\lambda) = \mathbf{0} = \Pi_{N'}^{1,H}(\rho)$ otherwise. This means that $\lambda = \widehat{p}[M]$, M enables t, and $\rho = \widehat{p}[M - \mathbf{pre}_N(\widehat{t}) + \mathbf{post}_N(\widehat{t})]$.

Definition 2 (Firing Rule). *Let EH be an* EHORNET *and $\mu, \mu' \in \mathcal{M}_H$ markings.*

- *The event $\widehat{\tau}^{\alpha}[\vartheta]$ is enabled in μ for the mode $(\lambda, \rho) \in \mathcal{M}_H^2$ iff $\lambda \sqsubseteq \mu \wedge \phi_{EH}(\widehat{\tau}[\vartheta], \lambda, \rho)$ holds and the guard $\widehat{G}(\widehat{t})$ holds, i.e. $E \models_{\mathcal{I}}^{\alpha} \widehat{G}(\widehat{\tau})$.*
- *An event $\widehat{\tau}^{\alpha}[\vartheta]$ that is enabled in μ can fire – denoted $\mu \xrightarrow[EH]{\widehat{\tau}^{\alpha}[\vartheta](\lambda,\rho)} \mu'$.*
- *The resulting successor marking is defined as $\mu' = \mu - \lambda + \rho$.*

Note, that the firing rule has no a-priori decision how to distribute the marking on the generated net-tokens. Therefore we need the mode (λ, ρ) to formulate the firing of $\widehat{\tau}^{\alpha}[\vartheta]$ in a functional way.

3 Fan-Bounded Safe, Elementary HORNETS

We know from [11, Lemma 3.1] that a safe EHORNET has a finite reachability set. More precisely: There are at most $\left(1 + U(m) \cdot 2^m\right)^{|\widehat{P}|}$ different markings, where m is the maximum of all $|P_k|$ and $U(m)$ is the number of object nets. In the general case we have $U(m) = 2^{(2^{4m})}$, which dominates the bound. It is double-exponential, while the number of different marking wihin each net-token is 2^m, i.e. "only" single-exponential. The huge number of object nets is the source of the exponential space requirement for the reachability problem.

Therefore, if one wants to require less than exponential space one has to restrict the structure of possible object nets in $\mathcal{U}_k$. The huge number of object-nets in $\mathcal{U}_k$ arises since we allow object nets with any number of places in the preset or postset, i.e. unbounded joins or forks. Therefore it seems promising to restrict the synchronisation degree.

In the following we want to restrict the number of object-nets in $\mathcal{U}_k$. We forbid unbounded joins or forks. From a practical point of view this can be considered as a rather unlikely. From a theoretical point of view we can take this into account with an parametrised complexity analysis. The parameter considered here is the maximal number of places in the pre- or postset, i.e. the maximal synchronisation at object-net level.

Definition 3. *An elementary* HORNET $EH = (\widehat{N}, \mathcal{U}, \mathcal{I}, k, l, \mu_0)$ *is called* β-fan-bounded *whenever all transitions of all object-nets have at most* β *places in the pre- and in the postset:*

$$\forall k \in K : \forall N \in \mathcal{U}_k : \forall t \in T_N : |{}^\bullet t| \leq \beta \wedge |t^\bullet| \leq \beta$$

The *fan-bound* of EH is defined as:

$$\beta(EH) := \max\{|{}^\bullet t|, |t^\bullet| \;:\; k \in K : N \in \mathcal{U}_k : t \in T_N\}$$

Note, that an elementary HORNET is always fan-bounded, since $P_N \subseteq P_k$ and P_k is always finite in the elementary case: $\beta(EH) \leq |P_k| < \infty$.

Proposition 2. *For a safe,* β-fan-bounded EHORNET *the cardinality of each net universe* $\mathcal{U}_k$ *is bounded for each* $k \in K$ *as follows:* $|\mathcal{U}_k| \leq 2^{O\left(n^{(4\beta)}\right)}$ *where* $n := |P_k|$.

Proof. For a safe, *β-fan-bounded* EHORNET the number of possible objects is calculated as follows: Each possible transition t chooses a subset of P_k for the preset ${}^\bullet t$ and another subset for the postset $t^\bullet$ with the constraint that these subsets have a cardinality of at most β. The number of these subsets is:

$$\left|\bigcup_{i=0}^{\beta} \binom{P_k}{i}\right| = \sum_{i=0}^{\beta} \binom{|P_k|}{i} = \binom{|P_k|}{0} + \binom{|P_k|}{1} + \cdots + \binom{|P_k|}{\beta}$$

(Here $\binom{A}{i}$ denote the set of all subsets of A that have cardinality i.)

We identify t with the pair $({}^\bullet t, t^\bullet)$. The number of different transitions is:

$$\begin{aligned} |T_k| &= \left(\tbinom{|P_k|}{0} + \tbinom{|P_k|}{1} + \cdots + \tbinom{|P_k|}{\beta}\right)^2 \\ &\leq \left(1 + n + \tfrac{n\cdot(n-1)}{2!} \cdots + \tfrac{(n)_\beta}{\beta!}\right)^2 \\ &\leq \left(const \cdot n^\beta\right)^2 \\ &= const \cdot n^{(2\beta)} \end{aligned}$$

So, the number of different transitions is in $O\left(n^{2\beta}\right)$.[2]

The set of labelled transitions is $LT_k := T_k \times (C_k \cup \{\bot_k\})$ and we have $|LT_k| = |T_k \times (C_k \cup \{\bot_k\})|$ different labelled transitions. We cannot use more channels than we have transitions in the object net, i.e. we could use at most $|T_k|$ different channels from $C_k \cup \{\bot_k\}$. Thus, we have:

$$|LT_k| = |T_k| \cdot (|C_k| + 1) \leq |T_k| \cdot |T_k|$$

[2] Note, that while the bound we have given for the general case in Lemma 2.1 in [11] is strict (i.e. there are Hornets that exactly have this number of object-nets) the calculation given here gives us only an upper bound.

From $|T_k| \leq const \cdot n^{(2\beta)}$ we obtain:

$$|LT_k| \leq \left(const \cdot n^{2\beta}\right)^2 = const \cdot n^{(4\beta)}$$

Thus the set of labelled transitions is in $O\left(n^{(4\beta)}\right)$, i.e. a polynomial in the number of places $n = |P_k|$ where the degree of the polynomial is given by the fan-parameter β.

Since each object net N in $\mathcal{U}_k$ is characterised by its set of labelled transitions and there are $|\mathcal{P}(LT_k)| = 2^{|LT_k|}$ subsets of LT_k, we have at most $2^{O\left(n^{(4\beta)}\right)}$ different object nets. qed.

Thus the of different object nets is only single-exponential for fan-bounded EHORNETS – and not double-exponential as in the general case.

Note, that the set of transitions is a polynomial in the number of places m where the degree is given by the fan-parameter $\beta = \beta(m) \leq m$. Of course if we have transitions that use all the places in the pre- or postset, i.e. $\beta = m$ we have an exponential number as before, since:

$$|T_k| = \left(\binom{|P_k|}{0} + \binom{|P_k|}{1} + \cdots + \binom{|P_k|}{m}\right)^2 = \left(2^{|P_k|}\right)^2 = 2^{(2|P_k|)}$$

So, the general analysis is just the special case where the fan-parameter β equals the numer of places m.

For safe, *β-fan-bounded* EHORNET we can give an upper bound for the number of reachable markings. The number of reachable markings is in $2^{O\left(n^{(4\beta+1)}\right)}$, i.e. exponential, where the exponent is a polynomial in the number of places n where the degree is given by the fan-parameter β.

Proposition 3. *A safe, β-*fan-bounded EHORNET *has a finite reachability set.*

The number of reachable markings is bounded by $2^{O\left(n^{(4\beta+1)}\right)}$ where n is the maximum of all $|P_k|$ and $|\widehat{P}|$.

Proof. Analogously Prop. 1 we have at most $(1 + U(m) \cdot 2^m)^{|\widehat{P}|}$ different markings in the safe HORNET.

For a *β-fan-bounded* EHORNET we have obtained in Prop. 2 a bound for the number of possible object-nets: $|\mathcal{U}_k| \leq U(m) = 2^{\left(const \cdot m^{(4\beta)}\right)}$. Thus the number of different markings in the safe, β-fan-bounded EHORNET is:

$$\begin{aligned}
(1 + U(m) \cdot 2^m)^{|\widehat{P}|} &\leq \left(1 + 2^{\left(const \cdot m^{(4\beta)}\right)} \cdot 2^m\right)^{|\widehat{P}|} \\
&\leq \left(2^{\left(const \cdot m^{(4\beta)} + m\right)}\right)^{|\widehat{P}|} \\
&\leq \left(2^{\left(const \cdot m^{(4\beta)}\right)}\right)^{|\widehat{P}|} \\
&= 2^{\left(const \cdot m^{(4\beta)} \cdot |\widehat{P}|\right)}
\end{aligned}$$

With $n := \max(m, |\widehat{P}|)$ the bound simplifies to:

$$2^{\left(const \cdot m^{(4\beta)} \cdot |\widehat{P}|\right)} \leq 2^{\left(const \cdot n^{(4\beta)} \cdot n\right)} = 2^{\left(const \cdot n^{(4\beta+1)}\right)}$$

The number of reachable markings is in $2^{O\left(n^{(4\beta+1)}\right)}$, i.e. exponential, where the exponent is a polynomial. qed.

The most extreme restriction is to forbid forks and joins at all. In this case we consider elementary HORNETS that have *state-machines* as object-nets only, i.e.

$$\forall k \in K : \forall N \in \mathcal{U}_k : \forall t \in T_N : |{}^\bullet t| \leq 1 \wedge |t^\bullet| \leq 1$$

An elementary Hornet with this restriction is called ESMHORNET (elementary state-machine HORNET) for short. An ESMHORNET is 1-fan-bounded EHORNET by definition. Therefore, we obtain the following corollary of Lemma 3:

Corollary 1. *A safe* ESMHORNET *has a finite reachability set. The number of reachable markings is bounded by* $2^{O\left(n^5\right)}$ *where* n *is the maximum of all* $|P_k|$ *and* $|\widehat{P}|$.

4 Complexity of the Reachability Problem

As p/t nets are a special subcase of fan-bounded EHORNETS (we simply restrict the system net to a single place with the p/t net of interest as the unique net-token) the reachability problem for safe, fan-bounded EHORNETS cannot be simpler than for p/t nets, i.e. it is at least PSPACE-hard, since the reachability problem for safe p/t nets is PSPACE-complete. In the following we show that the reachability problem for lies within PSPACE.

Lemma 1. *For safe,* β-fan-bounded EHORNETS *there exists a non-deterministic algorithm that decides the reachability problem within polynomial space:*
$\mathit{Reach}_{\beta\text{-}seH} \in \mathit{NSpace}\left(O\left(n^{(4\beta+1)}\right)\right)$ *where* n *is the maximum of all* $|P_k|$ *and* $|\widehat{P}|$.

Proof. Whenever μ^* is reachable it is reachable by a firing sequence without loops. The main idea of the algorithm is to guess a firing sequence $\mu_0 \xrightarrow{\theta_1} \mu_1 \xrightarrow{\theta_2} \cdots \xrightarrow{\theta_{max}} \mu_{max} = \mu^*$, where μ^* is the marking to be tested for reachability.

By Prop. 3 we know we have at most $max = 2^{O\left(n^{(4\beta+1)}\right)}$ different markings. Therefore, we can safely cut off the computation after max steps.

For each step $\mu_i \xrightarrow{\theta_i} \mu_{i+1}$ we choose non-deterministically some event θ_i. For a given marking μ_i we guess an event θ_i and a marking μ_{i+1} and test whether $\mu_i \xrightarrow{\theta_i} \mu_{i+1}$ holds.

- The markings μ_i and μ_{i+1} can be stored in $O\left(n^{(4\beta+1)}\right)$ bits, i.e. polynomial space.
- The event θ_i can be stored in polynomial bits: The choice for the system net transition $\widehat{t}$ fits in polynomial space. The variable binding α selects for each variable in the arc inscriptions some object net from the universe $\mathcal{U}_k$. Since we have always have a finite number of variables polynomial space is sufficient. For each kind k we have the multiset of channels $\widehat{l}^k_\alpha(\widehat{t})$ to synchronise with. In the proof of Prop 2 we have seen that we have at most $const \cdot n^{(4\beta)}$ labelled transitions in the object nets, i.e. polynomial space is sufficient. We guess a mode (λ, ρ). (Since we consider safe HORNETS the choice for λ is unique, since there is at most one net-token on each

place, which implies that whenever an event $\widehat{t^{\alpha}}[\vartheta]$ is enabled, then λ is uniquely determined. For the multiset ρ, we use the fact that for each place in the object net N we have at most one token to distribute over all generated net-tokes. For each net $N \in \mathcal{U}_k$ we select one of the net-tokens generated in the postset. All these choices need at most polynomial many bits.

– To check whether the event is enabled we have to whether test $\alpha \sqsubseteq \mu$. This holds iff $\exists \mu'' : \mu = \alpha + \mu''$. Since α and μ are known, this can be tested in-place by 'tagging' all addends from α in μ.
  Finally we check whether the successor marking $\mu' = \mu - \lambda + \rho$ is equal to μ_{i+1}. This can be done in-place as $\mu - \lambda$ are those addends, that have not been tagged.

After each step $\mu_i \xrightarrow{\theta_i} \mu_{i+1}$ we decrement a counter (which has been initialised with the maximal sequence length max), forget μ_i, and repeat the procedure with μ_{i+1} again until either the marking of interested is reached or the counter reaches zero.

As each step can be tested in polynomial space, the whole algorithm needs at most polynomial many bits to decide reachability. qed.

Now we use the technique of Savitch to construct a deterministic algorithm from the non-deterministic algorithm above.

Proposition 4. *The reachability problem* $\mathsf{Reach}_{\beta\text{-seH}}$ *for safe* β-fan-bounded EHORNETS *can be solved within polynomial space.*

Proof. We known by Lemma 1 that $\mathsf{Reach}_{\beta\text{-seH}} \in NSpace(O\left(n^{(4\beta+1)}\right))$. From Savitch's Theorem we obtain:

$$\mathsf{Reach}_{\beta\text{-seH}} \in DSpace\left(O\left(n^{(4\beta+1)}\right)^2\right) = DSpace\left(O\left(n^{2(4\beta+1)}\right)\right)$$

This proves our central result that the reachability problem for safe, fan-bounded EHORNETS is PSPACE-complete, while it requires exponential space in the general case of safe EHORNETS.

References

1. Köhler, M., Rölke, H.: Properties of Object Petri Nets. In Cortadella, J., Reisig, W., eds.: International Conference on Application and Theory of Petri Nets 2004. Volume 3099 of LNCS, Springer-Verlag (2004) 278–297
2. Köhler-Bußmeier, M., Heitmann, F.: On the expressiveness of communication channels for object nets. Fundamenta Informaticae **93** (2009) 205–219
3. Valk, R.: Object Petri nets: Using the nets-within-nets paradigm. In Desel, J., Reisig, W., Rozenberg, G., eds.: Advanced Course on Petri Nets 2003. Volume 3098 of LNCS, Springer-Verlag (2003) 819–848
4. Reisig, W.: Petri nets and algebraic specifications. Theoretical Computer Science **80** (1991) 1–34
5. Hoffmann, K., Ehrig, H., Mossakowski, T.: High-level nets with nets and rules as tokens. In: Application and Theory of Petri Nets and Other Models of Concurrency. Volume 3536 of LNCS, Springer-Verlag (2005) 268 – 288

6. Lomazova, I.: Nested petri nets for adaptive process modeling. In Avron, A., Dershowitz, N., Rabinovich, A., eds.: Pillars of Computer Science. Volume 4800 of LNCS, Springer-Verlag (2008) 460–474
7. Köhler-Bußmeier, M.: Hornets: Nets within nets combined with net algebra. In Wolf, K., Franceschinis, G., eds.: International Conference on Application and Theory of Petri Nets (ICATPN'2009). Volume 5606 of LNCS, Springer-Verlag (2009) 243–262
8. Köhler-Bußmeier, M., Heitmann, F.: Safeness for object nets. Fundamenta Informaticae **101** (2010) 29–43
9. Köhler-Bußmeier, M., Heitmann, F.: Liveness of safe object nets. Fundamenta Informaticae **112** (2011) 73–87
10. Köhler-Bußmeier, M.: A survey on decidability results for elementary object systems. Fundamenta Informaticae **130** (2014) 99–123
11. Köhler-Bußmeier, M.: On the complexity of the reachability problem for safe, elementary Hornets. Fundamenta Informaticae **129** (2014) 101–116
12. Lipton, R.J.: The reachability problem requires exponential space. Research Report 62, Dept. of Computer science (1976)
13. Köhler-Bußmeier, M., Heitmann, F.: An upper bound for the reachability problem of safe, elementary hornets. In Popova-Zeugmann, L., ed.: Proceedings of the 23rd International Workshop on Concurrency, Specification and Programming (CS&P). Number 1269 in CEUR Workshop Proceedings (2014) 101–112
14. Ehrig, H., Mahr, B.: Fundamentals of algebraic Specification. EATCS Monographs on TCS. Springer-Verlag (1985)