

Modelling Reaction Systems with Petri Nets (Extended Abstract)

Jetty Kleijn¹, Maciej Koutny², and Grzegorz Rozenberg^{1,3}

¹ LIACS, Leiden University, 2300 RA, The Netherlands

² School of Computing Science, Newcastle University, NE1 7RU, UK

³ Department of Computer Science, University of Colorado at Boulder
430 UCB Boulder, CO 80309-0430, U.S.A.

Abstract. We investigate how Petri nets could be used to provide a faithful semantics of reaction systems, a formal framework for the investigation of processes carried by biochemical reactions. We propose and discuss possible approaches to this problem using some existing Petri net classes and concurrency concepts, such as maximal parallelism. After that we introduce a new class of Petri nets, called SET-nets, which provide a computational model matching very closely that exhibited by reaction systems. The key difference between standard Petri nets and SET-nets is that the former support multiset-based token arithmetic, whereas the latter support set-based operations on tokens.

Keywords: reaction system, Petri net, living cell, natural computing, SET-net, model translation

1 Introduction

The investigation of the computational nature of biochemical reactions is a research topic of Natural Computing. One of the goals of this research is to contribute to a computational understanding of the functioning of the living cell.

Reaction systems [2, 3, 7–10] are a formal framework for the investigation of processes carried out by biochemical reactions in living cells. The central idea of this framework is that the functioning of a living cell is based on interactions between (a large number of) individual reactions, and moreover these interactions are regulated by two main mechanisms: facilitation/acceleration and inhibition/retardation. These interactions determine the dynamic processes taking place in living cells, and reaction systems form a formal framework for developing an abstract theory of these processes.

The model of reaction systems is based on principles remarkably different from those underlying other *existing models of computation*. The aim of this paper is to develop a faithful Petri net model of reaction systems. The main motivation behind this is to establish whether Petri net based concepts (such as causal processes) and methods (such as synthesis of nets from a specification of their behaviour) could be used to provide analytical tools for reaction systems. It is not the intention of this paper to provide direct feedback to the area of

biological applications, but to establish bridges between biology and Petri nets through the connection provided by reaction systems.

As a first step, we propose and discuss four different approaches to the modeling of reaction systems by using existing Petri net classes and concurrency concepts. However, as it turns out, in order to obtain a good match between reaction systems and Petri nets, it is necessary to re-evaluate one of the basic net principles, namely, token counting. This leads us to the introduction of a new class of Petri nets, called SET-nets, which provide a net based computational model matching very closely the computations exhibited by reaction systems. The main difference between SET-nets and standard Petri nets is that the latter support multiset-based token arithmetic, whereas the former support set-based (boolean) operations on tokens. Thus, the computational ‘intuition’ originating from reaction systems provides the inspiration to introduce a new class of nets with intriguing and yet to discover properties. Consequently, the main contribution of this paper is more than just providing a bridge between reaction systems and the world of Petri nets. In the future, after fully understanding and mastering the properties of the new SET-nets, one would hope to provide also a new set of tools and analyses for biological applications.

The paper is organised in the following way. In the next section, we describe basic notions of reaction systems. Section 3 describes two methods of modelling reaction system using low-level Petri nets, and the next one does the same using high-level Petri nets. The new class of SET-nets is introduced in Section 5, and in Section 6 we explain why this new class of nets can faithfully and elegantly model reaction systems. Comparison with related work is presented in Section 7. Proofs of the results presented in this paper can be found in [16].

Notation We use the standard mathematical notions and notation. A multiset over a set X is a function $\mu : X \rightarrow \mathbb{N} = \{0, 1, 2, \dots\}$, and its support is $\|\mu\| = \{x \in X \mid \mu(x) > 0\}$. The empty multiset \emptyset satisfies $\|\emptyset\| = \emptyset$. A multiset may be represented, somewhat informally, by listing its elements with repetitions, e.g., $\mu = \{y, y, z\}$ is such that $\mu(y) = 2$, $\mu(z) = 1$, and $\mu(x) = 0$ otherwise. We treat sets as multisets without repetitions.

2 Reaction systems

In this section, we explain some notions relevant to reaction systems. It is our intention to introduce enough concepts to allow one to follow the subsequent discussion on the relationship between reaction systems and Petri nets. For a comprehensive description of reaction systems, including motivations, applications and examples, the reader is referred to [7–9].

Definition 1 (reaction system [7–9]). *A reaction system is a pair: $\mathcal{A} = (S, A)$, where S is a finite background set comprising the entities of \mathcal{A} , and A is the set of reactions of \mathcal{A} . Each reaction is a triplet of the form: $a = (R, I, P)$,*

where the three components are finite non-empty sets: $R \subseteq S$ is the set of reactants, $I \subseteq S$ is the set of inhibitors, and $P \subseteq S$ is the set of products.

The components of a reaction $a = (R, I, P)$ are denoted by R_a , I_a and P_a , respectively. Definition 1 describes the *static* structure of a reaction system. To capture the *dynamic* behaviour of reaction systems, we need additional notions.

Definition 2 (state of reaction system). *A state of a reaction system is any set C of its entities. Then an initialised reaction system is a triplet $\mathcal{A} = (S, A, C_0)$, where (S, A) is a reaction system and $C_0 \subseteq S$ is the initial state.*

In this and in the next section, we will consider as a running example the initialised reaction system $\mathcal{A}_0 = (\{w, x, y, z\}, \{a, b, c\}, \{x, z\})$, with background set $\{w, x, y, z\}$, initial state $\{x, z\}$, and three reactions:

$$a = (\{x\}, \{y\}, \{y, z\}) \quad b = (\{y\}, \{x\}, \{x, z\}) \quad c = (\{z\}, \{w\}, \{z\}).$$

A reaction system with background set S has exactly $2^{|S|}$ potential states. To describe possible transitions between these states, we need to say what is meant by an occurrence of a reaction or a set of reactions.

Definition 3 (state change). *A reaction a is enabled at a state $C \subseteq S$ if $R_a \subseteq C$ and $I_a \cap C = \emptyset$; the result of a reaction a at C is defined by $res_a(C) = P_a$ if a is enabled at C and $res_a(C) = \emptyset$ otherwise. The result of A on C , denoted by $res_{\mathcal{A}}(C)$ consists of the products of all reactions from A enabled at C , that is*

$$res_{\mathcal{A}}(C) = \bigcup_{a \in A} res_a(C).$$

This state change is denoted by $C \longrightarrow res_{\mathcal{A}}(C)$.

Note that the state changes captured by Definition 3 are deterministic. Moreover, all entities in $C \setminus \bigcup_{a \in A} res_a(C)$ disappear. As a result, and unlike in other formal models of dynamic systems, there is no persistency in a reaction system in the sense that an entity present in a state disappears unless it is sustained by at least one reaction.

For the example reaction system \mathcal{A}_0 , we have:

$$\{x, z\} \longrightarrow \{y, z\} \quad \text{and} \quad \{y, z\} \longrightarrow \{x, z\} \quad \text{and} \quad \{w, x, y\} \longrightarrow \emptyset.$$

One may observe that there is no conflict between reactions in the ‘classic’ sense that the occurrence of one reaction might imply that another reaction which is also enabled at the current state, cannot occur. This, again, is a feature not found in most other formal models of dynamic systems. In particular, it is worthwhile to point explicitly to the ‘non-counting’ features of reaction systems: entities are either present or not, and produced or not, and reactions can or cannot occur based only on the presence or absence of certain entities. There is no representation of multiple instances of entities or multiple occurrences of

reactions. Thus reaction systems are a *qualitative* rather than a quantitative model.

We also note that there is an alternative notion of conflict-freeness for a set of reactions, called consistency. A set of reactions \mathcal{R} is *consistent* if for any two reactions $a, b \in \mathcal{R}$, $R_a \cap I_b = R_b \cap I_a = \emptyset$. Clearly, if a set of reactions is *not* consistent, then the reactions it comprises cannot be executed simultaneously.

Although the goal of this paper is a faithful ‘translation’ of reaction systems into Petri nets, we conclude this section with a number of comments about research on reaction systems. This research happens in the *framework* of reaction systems where a reaction system constitutes the basic technical notion. Depending on the goal of a specific research theme, many other constructs are introduced and studied (see, e.g., [2, 9, 10]) — they form various extensions of the basic notion of reaction system. For example, there are many biological situations where one needs to assign quantitative parameters (time, concentrations, ...) to states of a biochemical system. Although reaction systems are a qualitative model (they cannot ‘count’), they can be extended so that such quantitative parameters can be accommodated. This is done through the use of *measurement functions* which lead to *reaction systems with measurements* (see [2, 3, 9, 10]), where various numerical parameters can be assigned to (calculated for) consecutive states of dynamic processes.

Finally, we want to point out that (because living cells are open systems) reaction systems have an environment and they operate/evolve within a changing context (with entities coming from the environment influencing the transitions of dynamic processes). In this paper, however, we will consider only *context-independent* processes defined by a reaction system with an initial state, where each next state is obtained solely as the result of reactions taking place in the previous state (thus assuming that the environment does not influence state transitions).

3 Reaction systems and low-level Petri nets

In this section, we discuss two possible ways of modelling context-independent processes of reaction systems using low-level Petri nets (PT-nets extended with with inhibitor and activator arcs).

In addition to the standard notions of reaction systems, in order to better explain how they relate to Petri nets, throughout the rest of this paper we will say that a set $\mathcal{R} \subseteq A$ is enabled at C if each reaction of \mathcal{R} is enabled at C . If $\mathcal{R} \subseteq A$ is enabled at C , then

$$C \xrightarrow{\mathcal{R}} res_{\mathcal{R}}(C) = \bigcup_{a \in \mathcal{R}} P_a .$$

denotes the effect of \mathcal{R} at C .

Definition 4 (PT-nets with inhibitor and activator arcs [14]). A PT-net with inhibitor and activator arcs (or PTIA-net) $N = (Pl, Tr, Flw, Inh, Act, M_0)$

is a tuple such that Pl and Tr are finite, disjoint sets of respectively places and transitions, and: $Flw \subseteq (Pl \times Tr) \cup (Tr \times Pl)$, $Inh \subseteq Pl \times Tr$, $Act \subseteq Pl \times Tr$ are respectively the sets of flow, inhibitor and activator arcs. Moreover, M_0 is a multiset of places, the initial marking of N ; in general, any multiset of places is called a marking.

In diagrams, places are drawn as circles and transitions as rectangles. Markings are the possible global configurations (states) of N . We say that a place q is *marked* under a marking M if $M(q) > 0$, where $M(q)$ denotes the number of occurrences of q in M . In diagrams, markings are indicated by putting $M(q)$ *tokens* inside the circle representing q . If $(x, y) \in Flw$, then (x, y) is an *arc* leading from *node* x to *node* y . A double headed arrow between q and t indicates that $(q, t), (t, q) \in Flw$. An inhibitor arc ends with a small open circle, while an activator arc ends with a small black circle.

Given a node x , we denote by $\bullet x$ the set of *input nodes* of x , i.e., those y for which $(y, x) \in Flw$, and by x^\bullet the set of *output nodes* of x , i.e., those y for which $(x, y) \in Flw$. For a transition t we use: ${}^\circ t = \{q \mid (q, t) \in Inh\}$ and $\blacklozenge t = \{q \mid (q, t) \in Act\}$ to denote the inhibitor and activator places of t . All four notations extend in the usual way to sets of nodes. As in the case of reaction systems, we now formalise the notion of marking (state) change.


Definition 5 (marking change). *A multiset of transitions U (also called a step) is enabled at a marking M if ${}^\circ U \cap \|M\| = \emptyset$, $\blacklozenge U \subseteq \|M\|$ and, for every place q , $M(q) \geq \sum_{t \in q^\bullet} U(t)$ (recall that $\|M\|$ is the set of q which occur in M , and $U(t)$ is the number of occurrences of t in U).*

In such a case, U can be fired with its effect on M being given by the resulting marking M' such that, for every place q : $M'(q) = M(q) - \sum_{t \in q^\bullet} U(t) + \sum_{t \in \bullet q} U(t)$. We denote this by $M[U]M'$. Moreover, if U is a maximal (w.r.t. multiset inclusion) step of transitions enabled at M , then we may denote this marking change also by $M[U]_{max}M'$.

Note that whenever a step U is enabled at marking M it must be the case that all activator places of transitions in $\|U\|$ are marked (are in $\|M\|$) and none of the inhibitor places of transitions in $\|U\|$ are marked.

We now make some general observations and assumptions about the relationship between reaction systems and nets.

- Entities can be represented by places, and reactions by net transitions.
- Since there are no conflicts between reactions, activator arcs can be used to test for the presence of reactants (rather than claiming resources for the exclusive use as with ordinary arcs and input places).
- All reactions that can occur in a reaction system do occur, and the only entities left after a state change are the newly generated products. In the Petri net framework, these features correspond to *maximal parallelism* described at the end of Definition 5, and *place resetting* [6] described later on.


 Fig. 1. Method I and II representations of the reaction system \mathcal{A}_0 .

Method I. The first attempt is illustrated in Figure 1(a) for the example reaction system \mathcal{A}_0 . Method I produces a PTIA-net $N_I(\mathcal{A}_0)$ such that:

- Transitions a , b and c use activator arcs and inhibitor arcs to test respectively for the presence and absence of tokens in the places w , x , y and z .
- Places q_a , q_b and q_c ensure that the three transitions modelling reactions, i.e., a , b and c , fire at most once in any step. This corresponds to the ‘non-counting’ of occurrence instances of the same reaction in a reaction system.
- Transitions r_w , r_x , r_y and r_z (in a maximal step) empty the four places modelling entities w , x , y and z . This does not have any influence on the firing of the transitions a , b and c .
- In a single maximal step, $M[U]_{max}M'$, the net fires a maximal multiset of transitions U enabled at marking M and then produces a new marking M' . For the net in Figure 1(a), such a firing rule gives:

$$\{x, z, q_a, q_b, q_c\} [\{r_x, r_z, a, c\}]_{max} \{y, z, z, q_a, q_b, q_c\} \\ \{x, x, x, z, q_a, q_b, q_c\} [\{r_x, r_x, r_x, r_z, a, c\}]_{max} \{y, z, z, q_a, q_b, q_c\} .$$

Formally, given an initialised reaction system $\mathcal{A} = (S, A)$, Method I yields a PTIA-net $N_I(\mathcal{A})$ such that the places, transitions and the initial marking are, respectively: $Pl = \{q_a \mid a \in A\} \cup S$, $Tr = \{r_s \mid s \in S\} \cup A$ and $M_0 = \{q_a \mid a \in A\} + C_0$. Moreover, the sets of flow, inhibitor and activator arcs are, respectively:

$$Flw = \{(s, r_s) \mid s \in S\} \cup \{(a, q_a), (q_a, a) \mid a \in A\} \cup \{(a, s) \mid a \in A \wedge s \in P_a\} \\ Inh = \{(s, a) \mid a \in A \wedge s \in I_a\} \quad Act = \{(s, a) \mid a \in A \wedge s \in R_a\} .$$

Note that this kind of modelling in combination with the ‘resetting’ of places w , x , y and z in each fired step, implemented by the auxiliary transitions r_w , r_x , r_y and r_z , means that the resulting Petri net is bounded (in every reachable marking the multiplicity of each place is never more than the number of reactions of \mathcal{A} if \mathcal{A} has at least one reaction).

In order to relate the behaviour of the original reaction system \mathcal{A} and its PTIA-net representation $N_I(\mathcal{A})$ just introduced, we need two mappings. The first

one takes a marking M of $N_I(\mathcal{A})$ and returns a state of \mathcal{A} , and the other takes a step U of transitions of $N_I(\mathcal{A})$ and returns a set of reactions of \mathcal{A} , as follows $\nu_I(M) = S \cap \|M\|$ and $\varphi_I(U) = A \cap \|U\|$. It is then possible to show a number of results, where a marking M of the PTIA-net $N_I(\mathcal{A})$ is called *well-formed* if $M(q_a) = 1$, for every $a \in A$.

First, M_0 is a well-formed marking satisfying $\nu(M_0) = C_0$, and if M is a well-formed marking and $M[U]M'$, then M' is also well-formed. Second, if M is a well-formed marking, then for every reaction $a \in A$, a is enabled at M iff $\{a\}$ is enabled at state $\nu_I(M)$. We then can show that the translation is sound.

Theorem 1. *If M is a well-formed marking then:*

1. $M[U]M'$ implies $\nu_I(M) \xrightarrow{\varphi_I(U)} \nu_I(M')$. Moreover, if $M[U]_{max}M'$, then $\varphi_I(U)$ comprises all reactions enabled at $\nu_I(M)$.
2. $\nu_I(M) \xrightarrow{\mathcal{R}} C$ implies $M[U]M'$ for some U and M' satisfying: $\varphi_I(U) = \mathcal{R}$ and $\nu_I(M') = C$. Moreover, if \mathcal{R} comprises all reactions enabled at $\nu_I(M)$, then $M[U]_{max}M'$.

Thus, each maximal computational step in the Petri net corresponds to a unique execution of the reaction system, and each execution in the reaction system corresponds to at least one maximal step in the Petri net. For example, the two executions given above for the Petri net in Figure 1(a) both correspond to $\{x, z\} \xrightarrow{\{a, c\}} \{y, z\}$ in the reaction system \mathcal{A}_0 .

Note that in Figure 1(a) one cannot simply delete the auxiliary places of the form q_r as then each of the transitions representing reactions could be unboundedly enabled. To address this problem one could change the activator arcs from places representing entities into flow arcs. Then, however, it would be necessary to add weights $|R|$ to the arcs corresponding to the production of new entities in order to avoid conflicts on the places representing the reactants.

Method II. The first attempt to model context-independent reaction systems provides a sound translation, but it is not simple as it employs features which can make formal analysis and verification far from easy. One way of improving Method I could be to replace multisets of fired transitions by sets of fired transitions leading to a *maximal set-semantics*. This can be achieved by using *reset arcs* [6], connecting places to transitions and indicated by \star 's in the diagrams, which always empty their source place. Formally, reset arcs $Reset \subseteq Pl \times Tr$ do not have any influence on the enabledness of a step U , but the calculation of the marking of a place q after the firing of U (now a set) at marking M changes to:

$$M'(q) = \begin{cases} M(q) - |q^\bullet \cap U| + |\bullet q \cap U| & \text{if } (\{q\} \times U) \cap Reset = \emptyset \\ |\bullet q \cap U| & \text{otherwise.} \end{cases}$$

The resulting PTIA-net with reset arcs $N_{II}(\mathcal{A}_0)$ is shown in Figure 1(b). Transition r is always enabled and, when fired, removes all the tokens from the places modelling the entities. For the net in Figure 1(b), the new firing rule gives

$\{x, z\} [\{r, a, c\}]_{max} \{y, z, z\}$ and $\{x, x, x, z\} [\{r, a, c\}]_{max} \{y, z, z\}$. One can then show that a counterpart of Theorem 1 holds also in this case, with ν_{II} defined as ν_I before and $\varphi_{II}(U) = U \setminus \{r\}$. As transition r is always enabled, we now have a one-to-one correspondence between groups of executed reactions and transitions, at the price of introducing non-standard reset arcs.

To remove the need to have reset arcs or, equivalently, to obtain a one-to-one correspondence between states and markings, one could change the rules for inserting tokens into places, by basically applying an OR-treatment for arriving tokens. This would, of course, be a radical departure from the standard Petri net approach, but one worth investigating. The resulting model of SET-nets will be described in Section 5.

4 Reaction systems and high-level Petri nets

The two translations described in the previous section use low-level PT-nets extended with reset arcs in addition to inhibitor and activator arcs as well as maximal parallelism. Reset arcs are a non-standard mechanism and, in particular, they do not as yet support a causal process semantics. Moreover, the effect of a reset arc depends on the current marking rather than on a fixed input/output relation with its neighbourhood. To cope with this problem, we will now outline two translations from context-independent reaction systems to high-level Petri nets. We assume familiarity with the basic concepts of high-level nets [13], in particular, arc inscriptions, activator and inhibitor arcs, and simple transition guards.

Method III. The first translation is illustrated by the high-level net $N_{III}(\mathcal{A}_0)$ shown in Figure 2(a). In this case, tokens are positive integers acting as though they were time-stamps. Intuitively, a token n is active only in the n -th execution cycle of the reaction system. Because the same token cannot be accessed more than once in a step sequence evolution, reset arcs are not needed anymore. Since the \checkmark transition fires in each maximal step, the cycle number n held in the ‘clock’ place clk is known to all transitions representing reactions. In the places representing entities, they check only for tokens n , ignoring all the other tokens produced in previous cycles, and then produce tokens with value $n+1$ to be used in the next cycle. The initial marking M_0 is formed by inserting a single token 1 into place clk and all the places s such that $s \in C_0$. Note that the resulting net may be unbounded as the tokens in places representing entities are not ‘garbage collected’. For the high-level net $N_{III}(\mathcal{A}_0)$ in Figure 2(b), we have:

$$\begin{aligned} & \{x \mapsto \{1\}, y \mapsto \emptyset, z \mapsto \{1\}, w \mapsto \emptyset, clk \mapsto \{1\}\} \\ & \quad [\{a_{n \mapsto 1}, c_{n \mapsto 1}, \checkmark_{n \mapsto 1}\}]_{max} \\ & \{x \mapsto \{1\}, y \mapsto \{2\}, z \mapsto \{1, 2, 2\}, w \mapsto \emptyset, clk \mapsto \{2\}\} \\ & \quad [\{b_{n \mapsto 2}, c_{n \mapsto 2}, \checkmark_{n \mapsto 2}\}]_{max} \\ & \{x \mapsto \{1, 3\}, y \mapsto \{2\}, z \mapsto \{1, 2, 2, 3, 3\}, w \mapsto \emptyset, clk \mapsto \{3\}\}. \end{aligned}$$

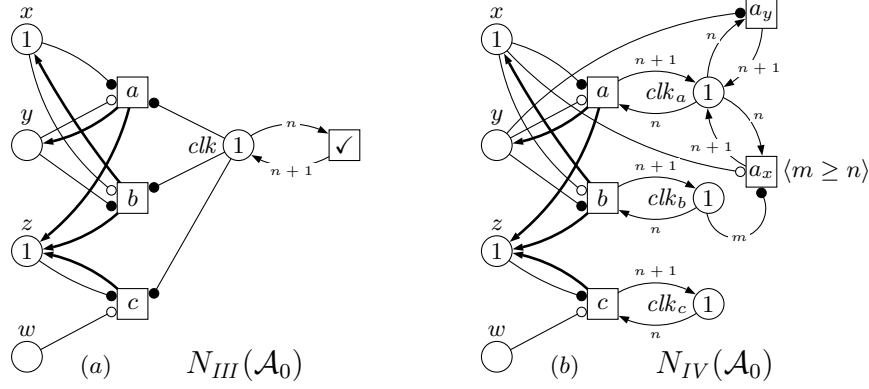


Fig. 2. Method III and IV representations of reaction system \mathcal{A}_0 . Note that n and m are net variables, and that to avoid clutter not all arcs have been annotated: all the flow (thicker) arcs to places x, y, z are in fact annotated with $n + 1$, and all the unannotated inhibitor and activator arcs are annotated with n . In (b), the auxiliary places for transitions b and c are omitted. Note that $\langle m \geq n \rangle$ is the guard of transition a_x , and all other transitions have the trivial *true* guard.

As in the case of Method I, not every marking M of $N_{III}(\mathcal{A})$ can represent a valid state of the reaction system \mathcal{A} . We say that M is *clock-consistent* if there is a single token k in place clk , and all the tokens l in other places satisfy $l \leq k$.

Relating the resulting net and the original reaction system can be done using the following two mappings: $\nu_{III}(M) = \{s \in S \mid \|M(clk)\| \cap \|M(s)\| \neq \emptyset\}$ and $\varphi_{III}(U) = U \setminus \{\checkmark\}$. One can show that M_0 is a clock-consistent marking satisfying $\nu(M_0) = C_0$, and if M is a clock-consistent marking and $M[U]_{max} M'$ then M' is also clock-consistent.

Theorem 2. *If M is a clock-consistent marking then:*

1. $M[U]_{max} M'$ implies $\nu_{III}(M) \xrightarrow{\varphi_{III}(U)} \nu_{III}(M')$.
2. $\nu_{III}(M) \xrightarrow{\mathcal{R}} C$ implies $M[U]_{max} M'$ for some U and M' satisfying: $\varphi_{III}(U) = \mathcal{R}$ and $\nu_{III}(M') = C$.

Method IV. In the second high-level net construction the aim is to eliminate the need for maximal parallelism using information present in the time-stamped tokens. We replace the global clk place by individual clk_a places, which are incremented by transitions a representing reactions. Moreover, whenever a is blocked from firing in a certain cycle one of the auxiliary transitions corresponding to the possible ‘reasons’ for the blocking a is fired to increment the token in clk_a . This results in an increment of the cycle number for this transition (in case

there is more than one reason for blocking, an auxiliary transition is chosen non-deterministically).

There are two possible reasons why a might be blocked in cycle n . One is the presence of a token n in the place representing an inhibitor of a , and to check for this we use a transition with an activator arc, e.g., a_y in Figure 2(b). The other is more complicated as it is a lack of token n in the place representing a reactant s for a , and to check for this we use a transition with an inhibitor arc. However, we also need to ensure that all transitions which feed tokens to s have already had a chance to do so, and we check this using extra activator arcs together with a transition guard which evaluates to true if all such feeding transitions have their local cycle sufficiently high, e.g., transition a_x in Figure 2(b). The overall result for the reaction system \mathcal{A}_0 is a high-level net $N_{IV}(\mathcal{A}_0)$ shown in Figure 2(b).

The resulting high-level net is executed according to the standard sequential (interleaving) firing rule and its behaviour closely simulates that of the net obtained by Method III, and so also the behaviour of the original reaction system. We skip the full description of the relationship between these two nets. Intuitively, a marking M of the second translation corresponds directly to a marking of the first one if all the places of the form clk_a contain the same single token k , and all the tokens l in other places satisfy $l \leq k$. (Note that from each reachable marking of the second translation one can execute a sequence of transitions leading to a marking with this property.)

5 Set-nets

In our attempts to obtain a direct and elegant translation from reaction systems into Petri nets, a major and as far as we can tell insurmountable problem was the fact that several transitions may insert tokens into a place representing the presence of a single entity. In this section, we introduce SET-nets, a model that resulted from closer investigations into the possibilities of an OR-treatment of arriving tokens representing the production of entities by reactions. Note that OR-treatment of causality has been considered in [20], but the underlying principle there was completely different from what we are going to propose.

The main idea is that in a SET-net there is no concept of counting. Places are marked or not marked and arcs have no weights. Set-nets resemble elementary net systems (EN-systems) [19] which is a fundamental model to study basic features of concurrent systems, including conflict, causality and independence. However, their execution semantics is different. In SET-nets, a marked place indicates the presence of a resource without any quantification. Hence any number of transitions that take input from this place can be fired at the same time. Moreover, firing a transition empties all its input places. Thus there are no conflicts over tokens in SET-nets, unlike in EN-systems or PT-nets. Similarly, places do not count the tokens, and the firing of a transition simply marks each of its output places (whether or not they were already marked). We will build up the new model in two stages, introducing first SET-nets with only flow arcs.




Fig. 3. A SET-net representing reaction system \mathcal{A}_1 (a); and an occurrence net constructed for its step sequence $\{b, q_\downarrow, s_\downarrow\}\{a, b, r_\downarrow, q_\downarrow\}$ (b).

Definition 6 (basic SET-net). A tuple $SN = (Pl, Tr, Flw, M_0)$ is a (basic) SET-net if the first three components are as in Definition 4, and $M_0 \subseteq Pl$ is the initial marking (in general, any set of places is a marking).

The graphical representation of SET-nets is the same as in the case of Petri nets. We now formalise the firing rule for SET-nets.

Definition 7 (marking change). A set of transitions U (also called a step) is enabled at a marking M if $\bullet U \subseteq M$. In such a case, U can be fired with its effect on M being given by the resulting marking $M' = (M \setminus \bullet U) \cup U^\bullet$. We denote this by $M[U]M'$. Moreover, if U is the set of all transitions enabled at M (i.e., all transitions t satisfying $\bullet t \subseteq M$), then we may write $M[U]_{max}M'$.

Hence a step U enabled at a marking M may contain two distinct transitions t and u for which $\bullet t \cap \bullet u \neq \emptyset$ or $t^\bullet \cap u^\bullet \neq \emptyset$ and yet the common places will never contain more than one token. Since *tokens are manipulated using set-based arithmetic* we have chosen the name ‘SET-nets’ for the new class of Petri nets.

We have introduced first basic SET-nets (without inhibitor and activator arcs), as it seems that one can attempt to develop for them a counterpart of ‘structure theory’ of PT-nets. To illustrate our point, let us consider a basic SET-net $SN = (Pl, Tr, Flw, M_0)$ with at least one transition. A non-empty set of places $Sphn \subseteq Pl$ is called a *siphon* if $\bullet Sphn \subseteq Sphn^\bullet$. Similarly, a non-empty set of places $Trap \subseteq Pl$ is called a *trap* if $Trap^\bullet \subseteq \bullet Trap$. It can be easily seen that an empty siphon cannot acquire a token by firing any transition, and a marked trap cannot become empty by firing any transition. Both type of sets of places can be used to provide a sufficient condition for deadlock-freeness in PT-nets which was a major motivation behind the development of their structure theory. As it turns out, the same can be done in case of SET-nets.

Theorem 3. *If in the initial marking, every siphon contains a marked trap, then the SET-net is deadlock free.*

We next introduce SET-nets with inhibitor and activator arcs.

Definition 8 (SET-net). *A tuple $SNIA = (Pl, Tr, Flw, Inh, Act, M_0)$ is a SET-net if the first five components are as in Definition 4, and the last one as in Definition 6.*

The definitions and notations concerning the marking change in $SNIA$ are the same as for SN in Definition 7 with one exception, namely a set of transitions U is enabled at a marking M if $\bullet U \cup \blacklozenge U \subseteq M$ and $\circ U \cap M = \emptyset$. It is interesting to observe that an enabled step U is always *consistent* in the sense that $(\bullet U \cup \blacklozenge U) \cap \circ U = \emptyset$. Such a property has a natural and direct (as we will see) connection with the notion of consistency introduced for reaction systems.

As before, given a transition t representing a reaction, the sets $\bullet t$, $\circ t$ and $\blacklozenge t$ correspond to the reactants, inhibitors and products of this reaction. However, we do not require that these sets be non-empty in a SET-net (at least at this point) as such an assumption is not necessary.

6 Reaction systems and SET-nets

Reaction systems and SET-nets fit together well in the sense that both do not count tokens and both change states on the basis of the presence/absence of resources, represented by sets. Moreover, under the SET-net semantics, ordinary arcs (transitions) can be used to empty places. In this semantics, reset arcs with their effect depending on the current number of tokens in a place are meaningless. Finally, following the assumption that all reactions that can take place do take place, the maximal set-semantics can be employed.

Figure 3(a) depicts a SET-net corresponding to a context-independent initialised reaction system $\mathcal{A}_1 = (\{r, q, s\}, \{a, b\}, \{q, s\})$, where $a = (\{r, q\}, \emptyset, \{r\})$ and $b = (\{q\}, \emptyset, \{r, q\})$. (For reasons of clarity, we allow in this section reactions without any inhibitors.) As before, places represent entities. Transitions r_\downarrow , q_\downarrow and s_\downarrow ensure that once the SET-net is active only tokens produced in the last maximal step are present in the current marking. For example, we have:

$$\{q, s\} [\{b, q_\downarrow, s_\downarrow\}]_{max} \{r, q\} [\{a, b, r_\downarrow, q_\downarrow\}]_{max} \{r, q\},$$

and so $\sigma = \{b, q_\downarrow, s_\downarrow\}\{a, b, r_\downarrow, q_\downarrow\}$ is a max-step sequence. Relating the behaviour of the SET-net model and the original reaction system is easy and we obtain a counterpart of Theorem 1 with $\nu(M) = M$ and $\nu(U) = U \setminus \{s_\downarrow \mid s \in S\}$.

For a SET-net without inhibitor and activator arcs as in Figure 3(a), one can investigate the causality semantics of reaction systems based on the unfoldings of the corresponding SET-nets. Figure 3(b) shows how such an occurrence net could be derived for the SET-net in Figure 3(a) and its step sequence $\{b, q_\downarrow, s_\downarrow\}\{a, b, r_\downarrow, q_\downarrow\}$ which corresponds of the state sequence $\{b\}\{a, b\}$ of the original reaction system. It is worth observing that the process has branching places which is not possible, in the case of processes of EN-systems or PT-nets. This, however, is fully consistent with the execution semantics of SET-nets.

Modelling inhibition aspects of reactions is rather straightforward using inhibitor arcs, as illustrated by the SET-net in Figure 4(a), representing the context-independent initialised reaction system $\mathcal{A}_2 = (\{r, q, s\}, \{a, b\}, \{q\})$, where:

$$a = (\{r, q\}, \emptyset, \{r\}) \quad \text{and} \quad b = (\{q\}, \{s\}, \{r, q\}) \quad \text{and} \quad c = (\{q\}, \emptyset, \{s\}).$$

Using inhibitor arcs gives a compact translation of reaction systems which is in a sense minimal w.r.t. the number of places, arcs and transitions. Moreover, relating the behaviour of the resulting SET-nets and the original reaction systems can be done as before. Formally, the places, transitions and initial marking of the translation are given by: $Pl = S$, $Tr = A \cup \{s_\downarrow \mid s \in S\}$ and $M_0 = C_0$. There are no activator arcs, and the flow and inhibitor arcs are as follows:

$$\begin{aligned} Flw &= \{(s, s_\downarrow) \mid s \in S\} \cup \{(s, a) \mid a \in A \wedge s \in R_a\} \cup \{(a, s) \mid a \in A \wedge s \in P_a\} \\ Inh &= \{(s, a) \mid a \in A \wedge s \in I_a\}. \end{aligned}$$

The development of a causal process semantics of SET-nets with inhibitor arcs is more difficult. It is therefore interesting to consider models of reaction systems using SET-nets without any inhibitor arcs, as outlined next.

Figure 4(b) shows a SET-net without inhibitor arcs modelling \mathcal{A}_2 . The way in which it does it is now more involved. More precisely, each execution step of the reaction system is simulated in two phases by the SET-net operating according to the maximal parallelism execution semantics. To keep these two phases clearly separated, they are controlled by an additional cyclic subnet with two places. The key aspect of the construction is the use of a ‘complement’ s^{cpl} of the ‘regular’ place s which at the time of checking whether s is empty by reaction b contains a token iff s is empty.

Figure 4(c) provides a generic picture of how, in the proposed construction, a SET-net (without inhibitor arcs) handles an entity r in its role as a reactant, inhibitor, and product. Note that r is represented by two places, r and r^{cpl} , and if r^{cpl} is marked then the entity r is absent in the current state. Moreover, each reaction d is represented by two transitions, d and d' . The first corresponds to the enabling stage of d , and the second to the generation of its products.

The first phase of the simulation always starts in a *consistent* marking M in which there is a token in place phI ; for every $s \in S$, $s \in M \Leftrightarrow s^{cpl} \notin M$, and otherwise all places are empty. In this phase transitions corresponding to reactions become active on the basis of the presence and absence of their reactants and inhibitors. Simultaneously, transitions of the form r_\downarrow and r_\uparrow take care that all the entities present in the current state cease to exist (their corresponding places are emptied and the complement places filled). In the second phase, each enabled transition d' finishes the execution of the corresponding reaction, and marks the places corresponding to the entities produced by reaction d and empties their complements.

Relating the behaviour of the SET-net model and the original reaction system is more complicated, using the following two mappings:

$$\nu(M) = M \setminus (\{phI\} \cup \{s^{cpl} \mid s \in S\}) \quad \varphi(U) = U \setminus (\{I\} \cup \{s_\downarrow \mid s \in S\} \cup \{s_\uparrow \mid s \in S\}).$$

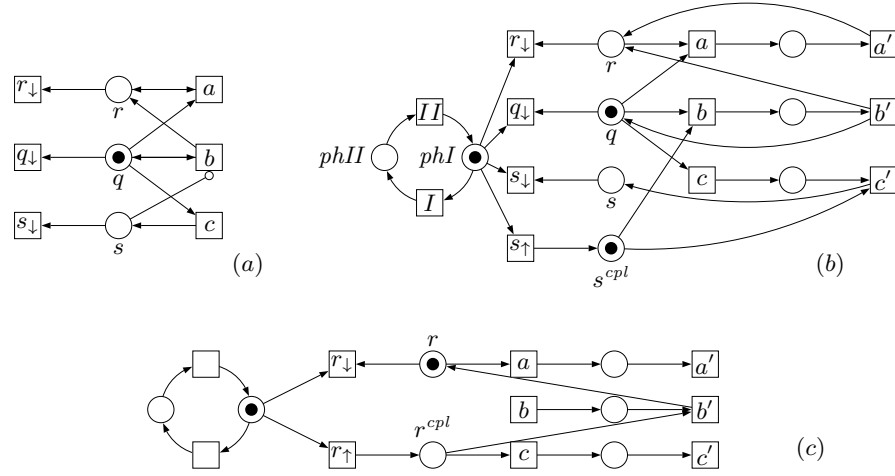


Fig. 4. Two SET-nets representing $\mathcal{A}_2(a, b)$. Generic translation without inhibitor arcs: here r is a reactant for reaction a , product for b , and inhibitor for c (c). Note that not all places and arcs are shown; in particular, each reaction has at least one reactant and hence transitions like c can only fire in the first phase.

One can then show that M_0 is consistent and satisfying $\nu(M_0) = C_0$, and if M is a consistent marking and $M[U]_{max}M''[U']_{max}M'$ then M' is also consistent.

Theorem 4. *If M is a consistent marking then:*

1. $M[U]M''[U']M'$ implies $\nu(M) \xrightarrow{\varphi(U)} \nu(M')$.
2. $\nu(M) \xrightarrow{\mathcal{R}} C$ implies $M[U]M''[U']M'$ for some U, U', M' and M'' satisfying: $\varphi(U) = \mathcal{R}$ and $\nu(M') = C$

7 Related work and concluding remarks

When introducing a new class of Petri nets, especially a fundamental one, it is necessary to put it in the context of existing formalisations. To make comparison fair, we will now drop the assumption about maximal parallelism in the execution of SET-nets (which is implied by the execution mode of reaction systems), and consider semantics which allows any set of enabled transitions to be fired.

Set-nets are so simple when it comes to their definition, that it is reasonable to expect that there were in the past net classes with similar features. Indeed, the fundamental class of EN-systems [19] extended with inhibitor as well as activator arcs [12, 17, 18] basically have the same static structure as SET-nets. However, their treatment of conflicts between transitions accessing the same token, as well blocking a transition which could add a token to a marked place, are totally




Fig. 5. Boolean net (SET-net with sequential semantics) (a), and 1-safe PT-net simulating its (sequential) behaviour (b).

different. The latter issue has been noted in the past, and the constraint relaxed. For example, there are variations of Petri nets, such as Boolean Petri nets, where adding a token to an already marked place does not add another token [4, 5, 11]. Also, behaviour of this kind was mentioned in [1] in the context of net synthesis. Having said that, the semantics considered in prior works known to us was based on single transition firings, rather than (maximal) steps as is the case for SET-nets. Therefore, the previous models were not concerned with multiple inputs of tokens to a single place something which is essential if one wants to faithfully model reaction systems. Furthermore, by aiming at a set-semantics, we had to introduce the non-conflict feature on the flow arcs consuming the tokens. Therefore, as far as we are aware, the model of SET-nets is an original contribution to the field of Petri nets.

As we already mentioned, SET-nets with interleaving semantics are nothing but Boolean nets used, for instance, in [5]. In such a case, the lack of conflict when firing two transitions sharing an input place is an irrelevant issue, and the only non-standard aspects is that firing a transition with a marked output place does not increase the token count in that place. Such a feature, moreover, can easily be modelled using ordinary 1-safe PT-nets, according to the following idea. First, one splits each place q into places q^0 and q^1 , respectively representing the lack and presence of a token in q . Then, each transition t adding tokens to place q is split into t^0 and t^1 to account for two different states the place q can be in represented by q^0 and q^1 . Figure 5 illustrates this construction. It can be easily seen that both nets generate the same *sequential* reachability graphs assuming that a^0 and a^1 are instances of a , and b^0 and b^1 are instances of a . However, once we start treating the net in Figure 5(a) as SET-net, the situation changes radically. The reason is that we then have three firings of the following form: $\emptyset[\{a\}]\{q\}$, $\emptyset[\{b\}]\{q\}$ and $\emptyset[\{a, b\}]\{q\}$. Now, the standard classes of Petri nets enjoy the so-called *subset* property which means that if a step U is enabled at marking M , then also any of its subsets is enabled as well. Suppose, then, that there is a Petri net N satisfying this property and such that its step reachability graph is the same as that of the SET-net in Figure 5(a), perhaps after renaming λ being applied to the transitions of the former. Then we have to have two transitions, t and u , in N such that $\lambda(t) = a$, $\lambda(u) = b$ and $M_0[\{t, u\}]M$. Then, by the subset closure property, we also have $M_0[\{t\}]M'$ and $M_0[\{u\}]M''$. Hence, by the reachability graph isomorphism, we must have

$M = M' = M''$ as well as $M_0 \neq M$. Hence we have: $M_0[\{t, u\}]M$ and $M_0[\{t\}]M$ and $M_0[\{u\}]M$ and $M_0 \neq M$. In the standard Petri nets, including various extensions of PT-nets, $M_0[\{t, u\}]M$ and $M_0[\{t\}]M$ would imply that u does not change the current marking. Similarly, $M_0[\{t, u\}]M$ and $M_0[\{u\}]M$ would imply that t does not change the current marking. Yet the simultaneous firing of t and u does change the marking as $M_0 \neq M$. This would produce a contradiction. What we just presented is intuition rather than proof, however, we expect that detailed arguments can be developed for any of the standard net classes. An important consequence, however, is that SET-nets are semantically different from the existing net classes and therefore deserve to be recognised as an original contribution.

8 Conclusions

The main initial motivation of our investigation was to see how Petri net based concepts could be deployed to analyse reaction systems. In particular, we wanted to discover methods for checking properties of reaction systems by relating them to the properties of the corresponding Petri nets and causal processes.

We proposed modelling methods resulting both in low-level and high-level nets. In all four cases, we established a close correspondence between the markings of Petri nets and states of the original reaction systems. The same was true of the evolutions of two corresponding models. In fact, we established that they have essentially isomorphic state spaces. All these net models, however, exhibited deficiencies w.r.t. simplicity and/or elegance and/or tractability of the translation. For example, both high-level net models are intrinsically unbounded, and the second of the low-level translations uses reset arcs. We therefore proposed a new class of Petri nets, called SET-nets, which we feel provide a strong match with the reaction systems and their semantics.

In this way we think we derived new interesting notions and contributions to Petri net theory based on our experiences with reaction systems in a similar way as the concepts of localities and locally maximal concurrency were derived from our previous investigation of a Petri net semantics of membrane systems [15].

Acknowledgement We would like to thank the anonymous reviewers for their suggestions and comments. This research was supported by the Pascal Chair award from Leiden University and the EPSRC VERDAD project.

References

1. E.Badouel and P.Darondeau: Theory of regions. Lecture Notes in Computer Science 1491 (1998) 529–586
2. R.Brijder, A.Ehrenfeucht, M.G.Main and G.Rozenberg: Reaction systems with duration. Lecture Notes in Computer Science 6610 (2011) 191–202
3. R.Brijder, A.Ehrenfeucht, M.G.Main and G.Rozenberg: A Tour of Reaction Systems. Int. Journal of Foundations of Computer Science (2011)

4. L.Czaja: A Calculus of Nets. *Cybernetics and Systems Analysis* 29 (1993) 185-193
5. P.De Bra, G.J.Houben and Y.Kornatzky: A Formal Approach to Analyzing the Browsing Semantics of Hypertext. *Proc. CSN-94 Conference* (1994) 78-89
6. C.Dufourd, A.Finkel, and Ph.Schoebelen: Reset Nets Between Decidability and Undecidability *Lecture Notes in Computer Science* 1443 (1998) 103-115
7. A.Ehrenfeucht, M.Main and G.Rozenberg: Combinatorics of Life and Death for Reaction Systems. *Int. J. of Foundations of Computer Science* 22 (2009) 345-356
8. A.Ehrenfeucht and G.Rozenberg: Reaction Systems. *Fundamenta Informaticae* 76 (2006) 1-18
9. A.Ehrenfeucht and G.Rozenberg: Events and Modules in Reaction Systems. *Theoretical Computer Science* 376 (2007) 3-16
10. A.Ehrenfeucht and G.Rozenberg: Introducing Time in Reaction Systems. *Theoretical Computer Science* 410 (2009) 310-322
11. M.Heiner, D.Gilbert and R.Donaldson: Petri Nets for Systems and Synthetic Biology. *Lecture Notes in Computer Science* 5016 (2008) 215-264
12. R.Janicki and M.Koutny: Semantics of Inhibitor Nets. *Information and Computation* 123 (1995) 1-16
13. K.Jensen: Coloured Petri Nets and the Invariant-Method. *Theoretical Computer Science* 14 (1981) 317-336
14. J.Kleijn and M.Koutny: Processes of Petri Nets with Range Testing. *Fundamenta Informaticae* 80 (2007) 199-219
15. J.Kleijn, M.Koutny and G.Rozenberg: Process Semantics for Membrane Systems. *Journal of Automata, Languages and Combinatorics* 11 (2006) 321-340
16. J.Kleijn, M.Koutny and G.Rozenberg: Modelling Reaction Systems with Petri Nets. *Technical Report CS-1244*. Newcastle University (2011)
17. M.Koutny and M.Pietkiewicz-Koutny: Synthesis of Elementary Net Systems with Context Arcs and Localities. *Fundamenta Informaticae* 88 (2008) 307-328
18. U.Montanari and F.Rossi: Contextual Nets. *Acta Informatica* 32 (1995) 545-596
19. G.Rozenberg and J.Engelfriet: Elementary Net Systems. *Lecture Notes in Computer Science* 1491 (1998) 12-121
20. A.Yakovlev, M.Kishinevsky, A.Kondratyev and L.Lavagno: et al: OR Causality: Modelling and Hardware Implementation. *Lecture Notes in Computer Science* 815 (1994) 568-587