

A Maude Implementation of Rewritable Petri Nets: a Feasible Model for Dynamically Reconfigurable Systems

Lorenzo Capra

Dipartimento di Informatica

Università degli Studi di Milano, Milan, Italy

capra@di.unimi.it

Petri Nets (PN) are a central, theoretically sound model for concurrent or distributed systems but, at least in their classical definition, not expressive enough to represent dynamic reconfiguration capabilities. On the other side, Rewriting Logic has proved to be a natural semantic framework for several formal models of concurrent/distributed systems. We propose a compact, efficient Maude formalization of dynamically reconfigurable PT nets (with inhibitor arcs), using as a running example the specification of a simple, fault-tolerant manufacturing system. We discuss the advantages of such a combined approach, as well as some concerns that it raises.

1 Introduction

Modern distributed systems operate in highly dynamic environments and have to frequently face changing operational conditions. Some system components may become temporarily or permanently unavailable, may appear/disappear, e.g., due to failures and/or dynamic balance policies. Self-adaptation is an effective approach to deal with the increasing complexity and dynamism of such systems. Reconfiguration capabilities are highly required by self-adaptive distributed systems as well as by other types of system, e.g., automated systems, which function without manual intervention, and whose operations are a part of the system control whose processes are completely automated with the help of control loops and special logic.

Formal methods providing the ability to reason about such complex systems are highly demanded. In particular, formal models representing both the system's base structure/functionality and the (usually complex) reconfiguration strategies are extremely important (e.g., to validate early design choices or to verify the system behavior at run-time).

Petri nets (PNs) are a sound, central model of concurrent or distributed systems, but not expressive enough to specify, in a natural way, the ability of some systems to modify their behavior/structure at run-time. Several dynamic PN extensions have been proposed in which the enhanced expressivity is not always supported by adequate analysis techniques. A representative of this category is the “nets within nets” paradigm, introduced by Valk [21], that gave rise to special High-Level PNs such as [11]. As for PNs with indistinguishable tokens, we have to mention Reconfigurable PN, a family of PN-based formalisms consisting of a marked net and a separated set of net-transformation rules specified according to the classical, algebraic style of Graph Transformations Systems (as a double pushout) [9, 12, 8, 10, 18]. Most research has focused on trying to formulate these models as \mathcal{M} -adhesive categories. See [16] for a survey of dynamic PN extensions.

In this paper, we provide a formalization of “rewritable” PT nets with inhibitor arcs (a Turing-complete PN class) in terms of rewriting-logic [15, 4], using Maude as a specification language [7]. We focus on operational, modelling aspects, by proposing an expressive and efficient framework for the

specification/analysis of (automated) distributed systems with reconfiguration capability. With respect to similar approaches [2, 17], in which (reconfigurable) PN classes are converted in Maude modules to exploit the model-checking tools of Maude, our encoding provides more data abstraction to ease the modeller task, is more compact and efficient, and promotes the definition of rewrite rules with a higher level of flexibility. Throughout the paper, we use as a benchmark a simple, yet tricky, model of fault-tolerant Manufacturing System. We conduct some experiments of formal verification of properties and briefly discuss of the advantages of using a somehow hybrid modelling approach like that we propose.

The presented work should be deemed as a preliminary, however, encouraging, step towards a Maude-based tool-set for the specification and the analysis of dynamically reconfigurable PT systems.

2 PT nets

This section collects a few definitions used in the paper. We refer to [19] for a complete description.

2.1 Multisets

A *multiset* (or *bag*) on a domain D is a map $b : D \rightarrow \mathbb{N}$, where $b(d)$ is the *multiplicity* of d in b . We say $d \in b$ if and only if $b(d) > 0$. We represent a multiset as a weighted “sum” of its elements. nil_D (*nil*, when D is implicit) denotes the *empty* multiset. Let $Bag[D]$ denote the set of bags over D , and $b_1, b_2 \in Bag[D]$. The *sum* $b_1 + b_2$ and the *difference* $b_1 - b_2$ are $Bag[D]$ elements such that $\forall d \in D$: $b_1 + b_2(d) = b_1(d) + b_2(d)$; $b_1 - b_2(d) = b_1(d) - b_2(d)$ if $b_1(d) \geq b_2(d)$, 0 otherwise. Note that $-$ is not associative. Similarly, relational bag-operators are defined component-wise: e.g., $b_1 < b_2$ if and only if $\forall d \in D : b_1(d) < b_2(d)$. With $b_1 <' b_2$ we mean the *restriction* of $<$ to $\{d | d \in b_1\}$.

2.2 Place/Transition (PT) Nets with Inhibitor Arcs

A PT net is a 5-tuple (P, T, I, O, H) , where:

P, T are non-empty, finite sets such that $P \cap T = \emptyset$

I, O, H are maps $T \rightarrow Bag[P]$ such that $\forall t \in T : I(t) \neq O(t)$,

P and T hold the *places* and the *transitions*, respectively. The former, drawn as circles, represent system state-variables, whereas the latter, drawn as bars, represent events causing local *state changes*¹. A (distributed) state of a PT net, called *marking*, is defined as a bag $m \in Bag[P]$.

A PT net is a kind of directed, bipartite multi-graph whose nodes are P and T . Maps I, O, H describe the *input*, *output*, and *inhibitor* edges, respectively. Graphically: $\bigcirc \rightarrow \square$ $\square \rightarrow \bigcirc$ $\bigcirc \dashrightarrow \square$

Let $f \in \{I, O, H\}$: if $k = f(t)(p) > 0$, then a weight- k edge of corresponding type links p to t . The behavior of a PT net is specified by the *firing rule*. We say that $t \in T$ is *enabled* in a marking m if and only if:

$$I(t) \leq m \wedge H(t) >' m$$

If t is enabled in m it may fire, leading to m' (we denote this $m[t]m'$), where:

$$m' = m + O(t) - I(t)$$

We call a pair (N, m) , where N is a PT net and m is a marking of N , a PT-system. The interleaving semantics of (N, m_0) , where m_0 denotes the PT system’s initial state, is specified by the *reachability*

¹The condition set on transitions ensure this.

graph (RG), an edge-labelled, directed graph (V, E) whose nodes are reachable markings. The RG is defined inductively: $m_0 \in V$; if $m \in V$ and $m[t]m'$ then $m' \in V$, $m \xrightarrow{t} m' \in E$.

A desirable property of a system (N, m_0) may be the absence of isolated places, i.e., places which are present in m_0 but not in N . In a highly dynamic context, however, we do not consider this as a strict requirement.

3 Maude and Rewriting Logic

Maude [7] is an expressive, strictly declarative language with a sound semantics in rewriting logic [15, 4]. Maude's statements are (possibly conditional) *equations* (keyword `eq`) and *rules* (keyword `r1`). Both sides of a rule/equation are terms of a given *kind* and may contain typed variables. Both rules and equations have a simple rewriting semantics in which instances of the lefthand side pattern are replaced by corresponding instances of the righthand side.

A Maude *functional* module (keyword `fmod`) contains only equations and is a functional program defining one or more operations through equations, used as simplifications. A Maude *system module* (keyword `mod`) contains rules and possibly equations. Rules are also computed by rewriting terms from left to right, but represent local transitions in a (concurrent) system. Although declarative in style and with a clear logical semantics, system modules are non-functional.

In Maude, a distributed system's state is usually represented as a kind of associative "multiset". Rules apply *concurrently* to different portions of a system, leading to a new state. Unlike for equations, there is no assumption on the confluence of rewrites. Reactive systems may not even have any final states.

Maude features expressivity, simplicity and performance. A wide range of systems are naturally expressible, from sequential, deterministic to highly concurrent non-deterministic ones. Maude may be used as a formal specification language, as a programming language and as a meta-language in which other formalisms, languages and logics can be naturally expressed. In this paper, we use Maude to specify adaptable PT systems and their operational semantics.

Maude's expressivity is achieved through: equational pattern matching modulo operator equational attributes; user-definable operator syntax/evaluation strategy; sub-typing(sorting) and partiality (kinds); generic types; reflection.

A Maude program is a logical theory, and a Maude computation is a deduction according to the axioms specified in the program. Under certain executability conditions of modules, the mathematical and the operational semantics coincide. We refer to [3] and [22] for functional and system modules, respectively.

A functional module specifies an *equational theory* in membership equational logic [14, 3]. Formally, such a theory is a pair $(\Sigma, E \cup A)$, where Σ is the signature, that is, the specification of all the sort, subsort, kind², and (overloaded) operator declarations in a functional module (plus the imported functional modules if any); E is the collection of (conditional) equations and memberships declared in the module(s), and A is the collection of equational attributes (`assoc`, `comm`, and so on) of operators (treated as predefined equations).

The models underlying functional modules are algebras, i.e., sets of data with related operations. For example, the model for (instantiated generic) module `BAG[A]` is the multisets over A with their standard operations. The family of Σ -ground terms T_Σ defines a model called Σ -algebra ($T_{\Sigma(X)}$ denotes the whole set of terms). According to Goguen and Burstall [5], the best model of $(\Sigma, E \cup A)$ is one that satisfies

²A *kind* is an implicit equivalence class gathering all sorts connected by the subsort partial-order relation; terms having a kind but not a sort may be considered as *undefined* or *errors*.

$E \cup A$ and is both junk-free (all elements can be denoted by T_Σ terms) and confusion-free (only elements that are forced to be equal by $E \cup A$ are identified). This model, called the *initial algebra* of $E \cup A$ and denoted $T_{\Sigma/E \cup A}$, does exist [3] and provides the denotational semantics of the Maude functional module specifying $(\Sigma, E \cup A)$. Formally, $T_{\Sigma/E \cup A}$ is the quotient of T_Σ in which the equivalence classes hold terms that prove equal using $E \cup A$.

If the axioms E are Church-Rosser and terminating modulo A (each ground term is thus simplified in a unique way regardless of the order in which equations apply) there is an intuitive, equivalent description for $T_{\Sigma/E \cup A}$. The final values (*canonical forms*) of all ground terms form an algebra called the *canonical term algebra*, denoted $CAN_{\Sigma/E \cup A}$. By definition, the `reduce` command of Maude interpreter reduces operators to their values in this algebra. The coincidence of the denotational and operational semantics is expressed by $T_{\Sigma/E \cup A} \cong CAN_{\Sigma/E \cup A}$.

A Maude's system module, including the imported submodules, specifies a generalized *rewrite theory* [15, 4], that is, a four-tuple $\mathcal{R} = (\Sigma, E \cup A, \phi, R)$ where $(\Sigma, E \cup A)$ is the membership equational theory specified by the signature, equational attributes, and equation statements in the module; ϕ is a map specifying, for each operator in Σ , its frozen arguments; and R is a set of rewrite rules³.

Intuitively, a rewrite theory specifies a concurrent system. The equational part $(\Sigma, E \cup A)$ specifies the algebraic structure of the states, formalized by the initial algebra $T_{\Sigma/E \cup A}$. The rules R (and ϕ) specify the system's dynamics, that is, the possible concurrent transitions of the system. In our context, they represent changes to the state/structure of a PT system. In rewriting logic, concurrent transitions become rewrite proofs; since several proofs may correspond to the same computation (because of different, equivalent interleavings), rewriting logic has an equational theory of proof equivalence [4, 15].

The initial model $\mathcal{T}_{\mathcal{R}}$ of \mathcal{R} associates to each kind k a labeled transition system (a category) whose states are $T_{\Sigma/E \cup A, k}$, and whose transitions take the form: $[t] \xrightarrow{[\alpha]} [t']$, with $[t], [t'] \in T_{\Sigma/E \cup A, k}$, and $[\alpha]$ an equivalence class of rewrites modulo the equational theory of proof-equivalence. Different $[\alpha]$ represent different "truly concurrent" computations of the system specified by \mathcal{R} .

The executability conditions for system modules match the notion of ground *coherence* between rules and equations [22]. Assuming that $E \cup A$ is Church-Rosser and terminating, an efficient strategy (the one adopted by the Maude `rewrite` command) is to first apply the equations to get a canonical form, then apply a rewrite rule in R . Coherence ensures that this strategy is complete, i.e., any rewrite of $t \in T_\Sigma$ with R is also possible with t 's canonical form. Coherence is crucial because rewriting modulo an equational theory is in general undecidable. It reduces rewriting with R modulo $E \cup A$ to rewriting with E and R modulo A , which is decidable given an A -matching algorithm.

Checking the confluence/termination and coherence properties of Maude modules is under the user's responsibility. In many cases, it is possible to prove these properties using the tools of Maude Formal Environment (MFE), available at <https://github.com/maude-team/MFE>, most of which are currently integrated with the Maude's interpreter. The Maude modules listed in this paper have been proven to be executable. Their functional part also meets two other desirable properties: i) each (equivalence class of a) ground term has a least sort (kind); ii) the canonical form of a well-defined ground term is only built of constructors (characterized by the `ctor` operator attribute). A syntactical condition (called term pre-regularity) and Maude's Sufficient Completeness Checker (SCC) have been used to verify properties i) and ii), respectively.

³ R rules don't apply to frozen arguments; in the paper we do not use frozen arguments.

4 Running example: a reconfigurable, fault-tolerant MS

As an example of “rewritable” PT system, we use the model of a manufacturing system (MS) equipped with dynamic reconfiguration capabilities to face failures affecting the production lines. Despite its simplicity, this is an interesting benchmark for any formalism intended to specify highly dynamic systems. We use a small variant (as for the reconfiguration steps) of the model introduced in [6].

The MS (Fig. 1, top) is composed of two symmetric production lines that (globally) work an even number of raw pieces. Pairs of worked pieces are assembled to get the final artefacts. Either line may get broken occasionally, in which case the system reconfigures itself so that it can continue working using the line left. The reconfiguration involves some changes to the system’s topology and the transfer of residual raw pieces staying on the faulty component to the available one (Fig. 1, bottom). If another fault affects the reconfigured system then (after a hypothetical repair phase, not represented) the MS goes back to its nominal configuration. The model’s parameter $M \in \mathbb{N}^+$ defines the number of raw pieces ($2 \cdot M$) worked during an entire production cycle. Each final artefact is immediately replaced by a pair of pieces to be worked, what makes the system’s nominal behaviour cyclic.

Tokens flowing through the PT places represent either raw or worked pieces. The two production lines are modelled by the symmetrical subnets $\{ p_2, t_1, p_4 \}$ (line 1) and $\{ p_3, t_2, p_5 \}$ (line 2). The assembly of raw pieces is represented by transition t_3 . Transition t_0 models the loader component, that (initially) picks up two raw pieces at a time from a storage represented by place p_1 putting them onto the lines. Transition t_4 models the immediate reload of the MS. The model includes also a very simple specification of fault occurrence, represented by transitions t_5 (for line 1) and t_6 (line 2). A fault on a line causes its immediate block, which is modelled, for example, by the inhibitor arc linking p_7 to t_1 . We assume that a simultaneous failure of both lines is not possible. For this reason, the transitions t_5 and t_6 are in symmetric structural conflict.

The PT net at the bottom of Fig. 1 clearly illustrates the changes to the MS layout that are required so that it can work using the only available line upon a fault, preserving the expected system behaviour. The marking of the PT system at the top represents the initial state, whereas that of the PT system at the bottom reflects the situation just after a fault occurrence: we may imagine (as explained later on) that M residual raw pieces on line 1 (place p_2) have been moved to line 2 (place p_3). Two major concerns that will be addressed in the next section are how to formally specify the MS reconfiguration(s) described in Fig. 1 and the condition(s) under which it occurs.

Although PT nets with inhibitor arcs are a Turing-powerful formalism, their limited expressivity doesn’t help in representing scenarios like that described (unless you enormously complicate the model). In the rest of the paper, we show how Maude can be used as a formal framework (in combination with PT) to overcome these drawbacks.

5 Encoding Rewritable PT systems in Maude

In this section, we review the Maude formalization of PT systems with dynamic reconfiguration capabilities. The full list of Maude source files is available at <https://github.com/lgcapra/rewpt>. Only one (system) module refers to a specific model, all the others should be seen as reusable components. The reusable part is built of a set of functional modules specifying the PT system signature and some structural equivalences (the signature’s main module is PT-SYS) and a system module (PT-EMU) which describes the system dynamics (firing rule). The main Maude modules are listed in the Appendix A.

The Maude formalization relies on two generic functional modules, BAG- $\{X\}$ and MAP+ $\{X, Y\}$, listed




Figure 1: The MS and its adaptation upon a fault on line 1 (a fault on line 2 is managed analogously).

in the Appendix B. The parameters of these generic types (all satisfying the built-in Maude trivial theory `TRIV`⁴) correspond to the bag's support and to the map's domain and codomain.

Differently from [20, 17], bags on a set are not merely represented as the free commutative monoid on that set. Despite its elegance and simplicity, such a formalization is unsuitable for bags with many repetitions. Our formalization is more convenient and practical, providing the $\{ _ \cdot _ , _ + _ , _ [_] _ - _ , _ < = _ , _ > _ , \text{set} \}$ operators (the first two *constructors*). In particular, the overloaded, commutative/associative sum makes it possible to intuitively represent a bag in a compact way, e.g., $3 \cdot a + 10 \cdot b$. Note the use of Maude advanced features, such as operator equational attributes, precedence and gathering (a term like $3 \cdot a - 1 \cdot a + 1 \cdot a$ is parsed as: $3 \cdot a - (1 \cdot a + 1 \cdot a) \equiv 1 \cdot a$), and the otherwise equation attribute, which permits a straightforward, efficient definition of difference, relational and lookup ($_ [_]$) operators. The `set` operator sets up a multiplicity for an element overwriting the older one. Setting a zero-multiplicity corresponds to removing an element.

The module `MAP+` defines a map as a “set” of entries, built with the associative/commutative operator $_ \cdot _$. Sort `Entry`, whose only constructor is the operator $_ | \rightarrow _$, is a subsort of `Map`. The constant `undefined` of the kind `[Y$Elt]` (corresponding to the sort `Y$Elt`) represents an undefined association. The overloaded operator `insert` adds a new key-value pair to a map, consistently updating the map

⁴Maude uses views to instantiate the formal type-parameters (Maude theories) of a generic module to concrete types (modules). We only report an example of theory and view, because they are very intuitive.

in the event the key is already present. Notice the use of matching and of the `otherwise` attribute to distinguish these two cases easily. The auxiliary operation `hasMapping` checks whether a key actually has an associated entry in a map. The lookup operator is represented by `_[_]`. With respect to the built-in MAP module of Maude, MAP+ has two significant extensions: as in the built-in module, the constructor `_,_` doesn't guarantee (for the sake of efficiency) the uniqueness of keys in a map; a conditional membership equation (using an auxiliary predicate), however, assigns a term of kind `[Map]` the sort `Map` only if it doesn't contain duplicated keys. The overloaded operator `remove`, instead, withdraws a given key from the map. We have defined bags apart, rather than as particular maps, for efficiency and readability.

PT system signature. The Maude signature of PT systems strictly follows the definition given in Sect. 2, so that passing from a formalization to the other is straightforward. The nodes of a PT net are trivially defined (see module TRAN) as terms with a subscript, e.g., $\tau(1)$. The PLACE module is analogously defined. A richer definition is possible, as discussed later. The module IMATRIX provides a compact representation for the incidence matrix of a PT net transition, syntactically represented as a triplet (according to the three types of edges) of bags on the set of places through the `[_ , _ , _]` constructor. IMATRIX imports BAG in a protected way (i.e., without introducing junks/confusion) by renaming some sorts/operators for readability. Two auxiliary operators are defined, to remove a given place from the transition incidence matrix and to check the adjacency of the transition to a place (the predicate `in`).

The central module PT-SYS defines the PT system signature. It builds on MAP+ and IMATRIX, both imported in a protected way, with some sort/operator renaming. In particular the parameterized sort `Map{Tran, Imatrix}` is renamed `Net`. That is, a PT net is syntactically expressed as a set of entries built of transitions and (local) incidence matrices. This formalization is very intuitive and allows structural changes to a system to be efficiently and consistently specified. A PT system term is a simple juxtaposition (empty constructor `__`) of a `Net` term and a `BagP` term (the latter representing a PT net marking). For convenience, the module provides three operators on the sort `Net` that correspond to the maps I, O, H of the PT net formal definition, and three predicates matching well-known conditions: `enabled` checks for the enabling condition of a PT net transition in a certain marking; `dead` (which builds on `enabled`) checks that a PT system is dead, i.e., no transition is enabled in the associated marking; `in` checks that a place belongs to a PT system (to the net and/or the associated marking). Note that the co-domains of `__`, `enabled`, and `In`, `Out`, `Inh` are *kinds* (`[System]`, `[BagP]`, and `[Bool]`, respectively), i.e., these are somehow *partial* operations. As for `__`, the reason is that the net of a PT system must not be empty (constant `emptyN`): a membership axiom specifies those terms having sort `System`. As for the others, the subtle reason is that they are defined only for `Tran` operands belonging to the specified PT net (system). Thus, `enabled(net m0, $\tau(1)$)` results in `false` if $\tau(1)$ belongs to `net` and is not enabled in `m0`.

The module PT-SYS includes two equations specifying structural equivalences between PT nets. One eliminates transitions having a null effect, according to the definition in Sect. 2. The other eliminates transitions structurally dead, because linked to a place by an input edge and by an inhibitor edge of non lesser weight. The commented equation would remove isolated places from a PT system (as we said, this is not mandatory).

PT system dynamics. The generic (system) module PT-EMU specifies the PT system operational semantics. The module's parameter must satisfy the Maude theory PTSYSTH, that requires two zero-arity operators describing a marking and a net. A conditional rewrite rule, syntactically much like the definition in Sect. 2, specifies the PT system firing rule. Notice the use (in the rule's condition) of the conjunction associative connective `/\` and of a matching equation (`$\tau := \tau'$`) which makes the rule com-

pact and efficient to apply. The free variables I, O, H, N' (two of which used in the rule's right-hand side) are instantiated by matching the left-hand side of the matching equation against the canonical ground term bound to variable N (that occurs on the rule's left-hand side).

Rewritable PT systems. The model-dependent part consists of a single Maude (system) module importing PT-SYS and satisfying the Maude theory PTSYSTH, which is (in part) mechanically derived from a PT system (N, m_0) . The module RWPT-FMS, e.g., comes from the PT system in Fig. 1-top, describing the FMS nominal behavior. Two equations assign terms `net`, `m0` the expressions encoding the PT net N and the initial marking m_0 , respectively. These expressions are obtained by arbitrarily defining two bijections $\phi_P : P \rightarrow \{0, \dots, |P| - 1\}$, $\phi_T : T \rightarrow \{0, \dots, |T| - 1\}$ (in our example, implicitly defined by subscripts of PT nodes). For the rest, the encoding of N (which consists of directly encoding the functions I, O, H into a net term, i.e., a set of entries $T : \text{Tran} \rightarrow Q : \text{Imatrix}$) and m_0 (as a BagP term) turns out to be straightforward (we skip the details).

Note the use of zero-arity, sort Nat operator `M` in RWPT-FMS, that allows the module to be syntactically parametric in the initial marking of place p_1 .

The remaining part of the module contains a (possibly empty) set R of rewrite rules specifying structural/state changes that may occur during the evolution of the system initially matched by the value of term `net m0`. With respect to similar approaches [17], where rewrite rules strictly follow, both syntactically and semantically, the paradigm of algebraic GTS (based, e.g., on DPO rules) and require quite complex consistency checks (e.g., to verify the respect of gluing conditions), our Maude-based specification fosters a much more flexible, free-style, though rigorous, formalization of (rule-based) system transformations.

We can classify rewrite rules according to the kind of $T_{\Sigma(X),k}$ terms in rule left- and right-hand sides: `Place`, `Tran`, `BagP`, `Imatrix`, `Net`, `System`. Except for rules of `System` type, all the other rules are local, that is, they act on one or more portions of a term describing a PT system. This is perfectly coherent with the operational semantics, intrinsically distributed, of rewriting logic, ensuring a lot of flexibility, but may have undesirable effects. A rule of `BagP` type, e.g., may touch both the marking and some local incidence matrix of a `System`.

The system module RWPT-FMS specifies two `System`-level (conditional) rewrite rules that model the system transformations described in Fig. 1. Both rule `r1` (`nominal=>faulty`) and rule `r2` (`faulty=>nominal`) are *symmetric*, i.e., refer to any of the two lines. In other words, each rule folds two ordinary rules. In rule `r1`, this is achieved by the matching of rule's left-hand pattern(s), which contains some variables (by convention, denoted by capital letters): for example, variable `P3` is instantiated to either ground value `p(2)` or `p(3)`, depending on whether a fault occurred on line 1 or line 2 (the opposite, as for variable `P2`). In rule `r2`, the two matching equations in the rule's condition make the free variables `P2, P3` be instantiated one to `p(2)` and the other to `p(3)`, depending on which line is broken. The condition of rule `r1` ensures that the transformation `nominal=>faulty` may take place when there are no residual "worked" pieces on the faulty line awaiting the assembly. Residual raw pieces on that line are then contextually moved to the working line (the marking of place bound to `P3` is cleared and that of place bound to `P2` increased correspondingly). The last clause of rule `r2`'s condition makes the back-transformation `faulty=>nominal` may take place when the system (which operates in a degraded way) enters a deadlock, upon another fault breaking the only line left: this happens when in the (broken) line there is at most one piece "worked" and all the other pieces still raw (the total of pieces is $2 \cdot M$). In order for the system to safely restart (after a global repair), M raw pieces are moved from that line to the one newly (re-)introduced (M tokens are withdrawn from the place bound to `P2` and added to the place bound to `P3`).

A library of base transformations. In the running example, we use two monolithic, non-trivial rewrite rules, each implementing a transformation of the system as a whole. Beside being complex, such an approach may be impractical when dealing with realistic adaptive systems, where changes are gradual, local and follow a strategy (procedure). Using our formal framework, you can define rewrite rules of varying granularity and locality, with both simplicity and great flexibility. To further ease the modeling task, however, we provide a minimal (extensible) set of base net-transformations specified by simple operators. The (functional) module PT-RWLIB collects two main operators (others are implicitly provided by modules BAG, MAP+, IMATRIX). `setw` allows one to update the weight of any edge in a PT net. For convenience, the edge type is specified as fourth argument of the operator: three constants of sort `Atype` are available for that. New edges/nodes may be added to a net using `setw`. Edges may be also removed with `setw`, using a zero-weight. This may indirectly cause the erasure of nodes that removed edges are incident to. The operator `remove` withdraws a place from a net. The homonym operator of MAP+ may be used to remove a transition from a net. Using BAG operators (e.g., `set`) we may freely modify the marking of a PT system and (local) incidence matrices, add/remove places, and so on. Module PT-RWLIB defines also an operator `w` to easily introspect the edge weight. This operator implements a partial function, as some of the operators of module PT-SYS.

For example, assume that the MS switches from one of the configurations shown in Fig. 1 to the other as a result of a sequence of small, local changes to the system's structure/state, carried out upon a fault occurrence. In this scenario, a thing to eventually do is to stop loading raw pieces into the MS. This is implemented by the (commented) rewrite rule `r3` in RWPT-FMS (requiring module PT-RWLIB). Using the operator `setw`, a new empty input place (`p(9)`) is linked to transition `t(0)` (the loader) to temporarily disable it. The rule's condition ensures that place `p(0)` is empty (i.e., a fault has occurred) and `p(9)` is not yet in the system. This rewrite step may be used in both the transformations described in Fig. 1.

5.1 Base notions and properties of rewritable PT systems

In this section, we provide an essential theoretical basis for the Maude formalization of rewritable PT systems through a few intuitive notions and properties. We refer to the *canonical* form of ground terms, that does exist and, if a term is well-defined (i.e., has an associated least *sort*), is built of constructors, since Maude modules satisfy the executability conditions (end of Section 2).

Property 1 (correspondence between PT systems and well-defined terms). *A PT system $S = (N, m)$ has an associated ground term of sort `System`, vice-versa, a ground term of sort `System` represents a PT system (up to isomorphism⁵).*

We have described how to get a `System` term from $S = (N, m)$. Vice versa, we observe that a canonical term of sort `System` is built (by definition of PT-SYS and MAP+) of $(n \ m)$, with n being a non-empty set of entries $T:Tran \rightarrow Q:Imatrix$ (without duplicate-keys) with the first two components (in, out) of each local incidence matrix unequal, and m a bag of places. By the way, the property above sets a bijection, denoted ψ (a pair ψ_P, ψ_T), between (classes of isomorphic) PT systems and well-defined `System` terms.

Let r be a rewrite rule, t, t' two ground terms of kind k . The notation $t \xrightarrow{r(\sigma)} t'$ means that 1) the rule's lefthand side $u \in T_{\Sigma(X),k}$ matches t (i.e., there is a ground substitution σ such that $\sigma(u) = t$)⁶, 2) t is rewritten to t' using r, σ .

⁵ S and S' are isomorphic iff there are two bijections $\phi_p: P \rightarrow P', \phi_t: T \rightarrow T'$, preserving the edges and the initial markings.

⁶ σ may be empty if u is a ground term; if r is a conditional rule σ may involve free variables introduced by matching equations used in the rule's condition.

Given a PT system $S = (N, m_0)$, let $\text{RWPT-}S$ represent a system module (satisfying PTSYSTH) in which the term $(\text{net } m_0)$ encodes S , R be the set of rewrite rules defined in $\text{RWPT-}S$, and $S\text{-EMU}$ the module PT-EMU whose parameter is instantiated (via an obvious view) to $\text{RWPT-}S$.

The interleaving semantics of a rewritable PT system specified by module $\text{RWPT-}S$ is expressed naturally by the labelled transition system, denoted RWLT_S , which is built from the initial term/state $(\text{net } m_0)$.

Definition 1 (State-transition system of $\text{RWPT-}S$). *Let $R' = R \cup \{\text{firing}\}$. RWLT_S is an edge-labelled, directed graph $(V_{\text{RW}_S}, E_{\text{RW}_S})$ inductively defined:*

$(\text{net } m_0) \in V_{\text{RW}_S}$; if $s \in V_{\text{RW}_S}$ and $s \xrightarrow{r(\sigma)} s'$ then: $s' \in V_{\text{RW}_S}$, $s \xrightarrow{r(\sigma)} s' \in E_{\text{RW}_S}$.

By default, the Maude interpreter's search command explores the state-space associated with an initial term by executing one-step rewrite rules in a fair, breadth-first way, therefore, coherent with the definition above.

$\text{RWPT-}S$ includes the ordinary behaviour of the PT system S .

Property 2 (RG inclusion). *RWLT_S contains a sub-graph isomorphic to RG_S .*

It directly follows from the fact that, by definition, for any transition t and for any marking m of S : $m[t]m'$ if and only if $\psi_P(m) \xrightarrow{\text{firing}(\sigma)} \psi_P(m')$, with $\sigma(T) = \psi_T(t)$ (T is the variable used in rule firing of module PT-EMU).

Notice that, in the event of badly defined/used rules, we may reach undefined (error) states, despite the initial system $(\text{net } m_0)$ is well-defined. For example, the rule

```
cr1 : (N, T |-> [I, 1 . P, nilP], T' |-> [1 . P, 0, nilP]) S =>
(N, T |-> [ I, 0, nilP]) S if I[P] = 0 /\ 0[P] = 0 /\ S[P] = 0 .
```

that aggregates two transitions connected by an intermediate empty place not linked to any other transitions, rewrites the `System` ground term

```
(t(1) |-> [1 . p(1), 1 . p(2), nilP], t(2) |-> [1 . p(2), 1 . p(1), nilP]) nilP
```

into: $(\text{emptyN } \text{nilP})$, an undefined term of kind `[System]` (due to the equations of functional module PT-SYS).

Definition 2 (Well-defined $\text{RWPT-}S$). *$\text{RWPT-}S$ specifies a well-defined rewritable PT system if and only if all reachable states in V_{RW_S} are terms of sort `System`.*

Rule validation. The module RWPT-FMS is well-defined. In general, however, ensuring the well-definiteness of Maude system modules specifying rewritable systems may not be simple. There are two approaches, shortly discussed in the following, each with different possible implementations.

One consists of defining (structurally) valid rewrite rules, and works also in the event the system state-space is infinite.

Definition 3 (Valid rewrite rule). *$r \in R$ is valid if and only if, for any ground term s of sort `System`, if $s \xrightarrow{r(\sigma)} s'$ then s' is of sort `System`.*

Each Maude rewrite rule `cr1 [r] : s => s' if cond`, where `cond` is the (possibly empty) rule's condition and $s, s' \in T_{\Sigma(X), \text{System}}$, may be rephrased as a valid rule using the built-in *sort predicate*

```
cr1 [vr] s => s' if cond /\ s' :: System
```

The weak spot of this elegant and efficient solution is that it may shadow bad design choices. As an alternative, we might define rewrite rules exclusively composed of *safe* net-operators. For example, the operator `setw`, defined in module PT-RWLIB , always results in a term of sort `Net`, possibly `emptyN` (for the sake of flexibility). The operator `setwS`, which builds on `setw`, also guarantees that the resulting term is a non-empty PT net (note the mixed use of *owise* equation attribute and a matching equation).

6 Property Verification

In this section, we briefly address some tools that are available to formally verify the properties of a Maude-based specification of rewritable PT systems, and we present a few experimental data. The goal is not to contend with well-established model-checkers based on PN and/or related formalisms, but to show some potential benefits of the methodology proposed to specify dynamically reconfigurable systems.

A Maude system module (which specifies a rewrite theory) provides an executable formal model of a distributed system. Under appropriate conditions, we can check that this model satisfies some properties, or obtain counterexamples. This kind of model-checking analysis is quite general and builds on the `search` command, which allows one to explore (following a breadth-first strategy) the reachable state-space in different ways. For example, using bounded model-checking, if the system state-space is huge (or even infinite), or model-checking of infinite-state systems through abstractions. Under finite reachability assumptions, we might more efficiently model check any linear time temporal logic (LTL) property of a system module using the LTL Maude modules. We here focus on a simple, yet helpful, model-checking capability, namely, the model checking of invariants using the `search` command.

One invariant we might like to verify about a rewritable PT system is deadlock freedom. A straightforward way to check this property is to issue the `search` command below, which searches for *any* final states of our running example starting from the initial PT system configuration.

```
Maude> search in FMS-EMU : net m0 =>! X:System .
```

It gives no solution, meaning that the self-adaptive MS is deadlock-free⁷. Table 1 reports some data about the performance of this command, as the system size varies. The experiments were carried out on a Intel Core i7-6700 equipped with 32 GB of RAM.

Another interesting search concerns the existence of dead states inside the different configurations (the nominal one and the two symmetric, faulty ones) that the system enters during its evolution. The command to issue is (`*` means “in zero or more steps”, `dead` is the predicate defined in module `PT-SYS`):

```
Maude> search in FMS-EMU : net m0 =>* X:System such that dead(X:System) .
```

There are six solutions (for any M), two for each configuration of the system. For example, for $M = 50$, we get (for readability, in the following excerpt of the command’s output, we use the term `net` to denote the system’s nominal configuration instead of the much longer canonical form)

Solution 1 (state 133541)

```
states: 133542 rewrites: 16647633 in 65130ms cpu
X:System --> (t(0) |-> [1 . p(1),1 . p(3),nilP], t(2) |-> [1 . p(3),1 . p(5),
1 . p(8)], t(3) |-> [2 . p(5),1 . p(6), nilP], t(4) |-> [1 . p(6),2 . p(1),
nilP], t(6) |-> [1 . p(0),1 . p(8),nilP]) 1 . p(8) + 100 . p(3)
```

Solution 2 (state 133542)

```
states: 133543 rewrites: 16647811 in 65140ms cpu
X:System --> (t(0) |-> [1 . p(1),1 . p(2),nilP], t(1) |-> [1 . p(2),1 . p(4),
1 . p(7)], t(3) |-> [2 . p(4),1 . p(6), nilP], t(4) |-> [1 . p(6),2 . p(1),
nilP], t(5) |-> [1 . p(0),1 . p(7),nilP]) 1 . p(7) + 100 . p(2)
```

Solution 3 (state 142860)

```
states: 142861 rewrites: 17841539 in 68140ms cpu
X:System --> (t(0) |-> [1 . p(1),1 . p(2),nilP], t(1) |-> [1 . p(2),1 . p(4),
```

⁷Using the MAUDE LTL modules we might even check that the initial state is a home-state, i.e., the system is cyclic.

Table 1: Performance of search command as M varies

M	# states	# rewrites	time (ms)
3	350	35730	44
5	1232	130539	153
10	8932	995724	1192
20	84007	9737494	36090
50	2186132	261564504	862536

```
1 . p(7)], t(3) |-> [2 . p(4),1 . p(6), nilP], t(4) |-> [1 . p(6),2 . p(1),
nilP], t(5) |-> [1 . p(0),1 . p(7),nilP]) 1 . p(4) + 1 . p(7) + 99 . p(2)
```

Solution 4 (state 142865)

states: 142866 rewrites: 17842109 in 68140ms cpu

```
X:System --> (t(0) |-> [1 . p(1),1 . p(3),nilP], t(2) |-> [1 . p(3),1 . p(5),
1 . p(8)], t(3) |-> [2 . p(5),1 . p(6),nilP], t(4) |-> [1 . p(6),2 . p(1),
nilP], t(6) |-> [1 . p(0),1 . p(8),nilP]) 1 . p(5) + 1 . p(8) + 99 . p(3)
```

Solution 5 (state 1083022)

states: 1083023 rewrites: 143322169 in 781220ms cpu

```
X:System --> net 1 . p(8) + 50 . p(3) + 50 . p(4)
```

Solution 6 (state 1084477)

states: 1084478 rewrites: 143524908 in 782130ms cpu

```
X:System --> net 1 . p(7) + 50 . p(2) + 50 . p(5)
```

We may suppose a generic form for the six states above in which expressions M , $2 \cdot M$, $2 \cdot M - 1$ should be used in place of values 50, 100, 99, respectively.

As a last example of use of search, we check whether the system correctly evolves from an inner deadlock: the command below lists the system states reachable in one-step from a dead state which refers to the system's nominal configuration after line 2 gets faulty (for $M = 10$).

```
Maude> search net 1 . p(8) + 10 . p(3) + 10 . p(4) =>1 X:System .
```

Solution 1 (state 1)

states: 2 rewrites: 23 in 0ms cpu

```
X:System --> (t(0) |-> [1 . p(1),1 . p(2),nilP], t(1) |-> [1 . p(2),1 . p(4),
1 . p(7)], t(2) |-> [1 . p(3),1 . p(5),1 . p(8)], t(3) |-> [2 . p(4),1 . p(6),
nilP], t(4) |-> [1 . p(6),2 . p(1),nilP], t(5) |-> [1 . p(0),1 . p(7),nilP])
1 . p(0) + 10 . p(2) + 10 . p(4)
```

As you can see, the system enters a degraded configuration in which only line 1 is working and the raw pieces left on the faulty line (a half of the total) have been moved on line 1.

6.1 Exploiting PT nets structural analysis

In a somehow hybrid modelling framework like that proposed, we can benefit from the tools/techniques available for both formalisms. For example, structural analysis of Petri nets may be an interesting, really

Table 2: P- and T-semiflows of the PT nets specifying the MS (Fig. 1)

nominal behavior		faulty behavior	
pin ₁	$p_1 + 2 \cdot p_6 + 2 \cdot p_2 + 2 \cdot p_4$	pin' ₁	$p_1 + 2 \cdot p_6 + p_3 + p_5$
pin ₂	$p_1 + 2 \cdot p_6 + 2 \cdot p_3 + 2 \cdot p_5$	pin' ₂	$p_0 + p_8$
pin ₃	$p_0 + p_7 + p_8$		
tin ₁	$t_0 + t_1 + t_2 + t_3 + t_4$	tin' ₁	$2 \cdot t_0 + 2 \cdot t_2 + t_3 + t_4$

efficient alternative/complement to state-space inspection techniques, when the system state-space is huge and performances degrade (as shown in Table 1). Furthermore, structural analysis doesn't depend on the initial marking of a PT system, so we may use it to get some kind of parametric outcomes. Let us show an application of semiflows on this direction to our running example.

Let \mathbf{Q} be the $|P| \cdot |T|$ matrix such that is $\mathbf{Q}_{[p,t]} = 0(t)(p) - I(t)(p) \in \mathbb{Z}$. Any P -vector \mathbf{p} which is a non-null, positive integer solution of the product $\mathbf{p} \cdot \mathbf{Q} = \mathbf{0}$, called P -semiflow, expresses a conservative law for the marking of places corresponding to non-zero entries of \mathbf{p} . Any T -vector \mathbf{t} which is a non-null, positive integer solution of the product $\mathbf{Q} \cdot \mathbf{t} = \mathbf{0}$, called T -semiflow, expresses a cyclicity effect for firing sequences matching the semiflow.

By inspecting the Maude module RWPT-FMS one can (also formally) check that the system's structure may only be one of those described in Fig. 1, more a symmetrical faulty one.

Table 2 shows the semiflows⁸ of the PT nets in Fig. 1 (there are symmetric semiflows for the other faulty configuration). We observe that the PT nets are covered by P -semiflows, therefore, the whole system is structurally bounded. The T -semiflows represent a base production cycle in the nominal and faulty configurations of the MS. Consider the two P -semiflows of the faulty configuration, and the initial marking immediately after a system switch. We derive these invariant marking-expressions: $\forall m : m(p_0) + m(p_8) = 1$, $m(p_1) + 2 \cdot m(p_6) + m(p_3) + m(p_5) = 2 \cdot M$. With simple arguments we can show that if $m(p_8) = 1$ then we *eventually* reach either of the two dead states: $m' = 1 * p_8 + 2 \cdot M * p_3$, $m'' = 1 * p_8 + 1 * p_5 + (2 \cdot M - 1) * p_3$, corresponding to matches 1,4 found with the `search` command (for $M = 50$). We have used this parametric outcome in rule `r2` of module RWPT-FMS.

Even if, in general, a static prediction of all the possible structural changes of a rewritable PT system specified with Maude may be more complex, or even impossible, the opportunity to exploit the structural analysis capabilities of Petri nets together with the formal analysis tools of Maude looks promising.

7 Conclusion, open issues and ongoing work

We have presented a Maude formalization of “rewritable” PT nets, a framework for the specification/analysis of distributed system with dynamic reconfiguration capabilities. With respect to similar approaches, the proposed encoding provides much more data abstraction to ease the modeller task, is more compact and efficient, and fosters the definition of rewrite rules with a high degree of flexibility. We have used as a (simple but tricky) running example throughout the paper a reconfigurable, fault-tolerant Manufacturing System. We have reported some experiments of formal verification of properties and discussed about the possible advantages of such a hybrid modelling approach.

⁸Semiflows have been computed with the GreatSPN tool [1].

Ongoing work & open issues We first plan to enrich the modular and intuitive Maude specification with structural extensions (e.g., test/flush arcs, transition priorities) that further enhance the model's expressivity. The idea is to use "decorated" terms representing PT nodes (e.g, $t(1, "line", 0)$, where there are a label and a value indicating the priority) and to update the firing rule in module PT-EMU accordingly.

A more complex, really interesting extension is the Maude encoding of rewritable High-Level PN (e.g., algebraic PN) possibly using the meta-level and term-unification modules.

The Labelled Transition System of a Maude specification of rewritable PT nets should be defined up to isomorphism of PT systems (in that case, the state-space of the running example should be reduced by half). A possible solution is to define a syntactical normal form for PT systems using equations. This is generally complex, but some heuristics could help dramatically reduce the inefficiency in most practical cases, e.g., by defining classes of "similar" PT nodes (using labels) so that an isomorphism preserves these symmetry classes. This topic, however, deserves special attention.

Rule-based transformations are simple to specify and elegant, but sometimes they are shown to be flawed and ill-equipped for representing realistic situations. Passing to more complex and suitable forms should be easy using the *reflection* capability of Maude and/or the Maude *strategy language* [13] (a specification layer above those of equations and rules).

These advanced features of Maude might be also used to control the rewriting process and break the construction of the transition system of a rewritable PT net when an undefined/error state is reached. The strategy language provides a cleaner way, respecting the separation of concerns.

References

- [1] Elvio Gilberto Amparore, Gianfranco Balbo, Marco Beccuti, Susanna Donatelli & Giuliana Franceschinis (2016): *30 years of GreatSPN*. In: *Principles of Performance and Reliability Modeling and Evaluation*, Springer, pp. 227–254, doi:10.1007/978-3-319-30599-8_9.
- [2] Paulo E. S. Barbosa, João Paulo Barros, Franklin Ramalho, Luís Gomes, Jorge Figueiredo, Filipe Moutinho, Anikó Costa & André Aranha (2011): *SysVeritas: A Framework for Verifying IOPT Nets and Execution Semantics within Embedded Systems Design*. In Luis M. Camarinha-Matos, editor: *Technological Innovation for Sustainability - Second IFIP WG 5.5/SOCOLNET Doctoral Conference on Computing, Electrical and Industrial Systems, DoCEIS 2011, Costa de Caparica, Portugal, February 21-23, 2011. Proceedings, IFIP Advances in Information and Communication Technology 349*, Springer, pp. 256–265, doi:10.1007/978-3-642-19170-1_28.
- [3] Adel Bouhoula, Jean-Pierre Jouannaud & José Meseguer (2000): *Specification and proof in membership equational logic*. *Theoretical Computer Science* 236(1), pp. 35–132, doi:10.1016/S0304-3975(99)00206-6.
- [4] Roberto Bruni & José Meseguer (2003): *Generalized Rewrite Theories*. In Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow & Gerhard J. Woeginger, editors: *Automata, Languages and Programming*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 252–266, doi:10.1007/3-540-45061-0_22.
- [5] R. M. Burstall & J. A. Goguen (1982): *Algebras, Theories and Freeness: An Introduction for Computer Scientists*. In Manfred Broy & Gunther Schmidt, editors: *Theoretical Foundations of Programming Methodology: Lecture Notes of an International Summer School, directed by F. L. Bauer, E. W. Dijkstra and C. A. R. Hoare*, Springer Netherlands, Dordrecht, pp. 329–349, doi:10.1007/978-94-009-7893-5_11.
- [6] Matteo Camilli, Carlo Bellettini & Lorenzo Capra (2018): *A High-Level Petri Net-Based Formal Model of Distributed Self-Adaptive Systems*. In: *Proceedings of the 12th European Conference on Software Architecture: Companion Proceedings, ECSA'18*, Association for Computing Machinery, New York, NY, USA, pp. 1–7, doi:10.1145/3241403.3241445.

- [7] Manuel Clavel, Francisco Dur'an, Steven Eker, Patrick Lincoln, Narciso M. Oliet, Jos'e Meseguer & Carolyn Talcott (2007): *All About Maude - A High-Performance Logical Framework: How to Specify, Program, and Verify Systems in Rewriting Logic*. Lecture Notes in Computer Science, Springer, doi:10.1007/978-3-540-71999-1.
- [8] Hartmut Ehrig, Kathrin Hoffmann, Julia Padberg, Ulrike Prange & Claudia Ermel (2007): *Independence of Net Transformations and Token Firing in Reconfigurable Place/Transition Systems*. In Jetty Kleijn & Alex Yakovlev, editors: *Petri Nets and Other Models of Concurrency – ICATPN 2007*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 104–123, doi:10.1007/978-3-540-73094-1_9.
- [9] Hartmut Ehrig & Julia Padberg (2004): *Graph Grammars and Petri Net Transformations*. In Jörg Desel, Wolfgang Reisig & Grzegorz Rozenberg, editors: *Lectures on Concurrency and Petri Nets: Advances in Petri Nets*, Springer, Berlin, Heidelberg, pp. 496–536, doi:10.1007/978-3-540-27755-2_14.
- [10] Laid Kahloul, Allaoua Chaoui & Karim Djouani (2010): *Modeling and Analysis of Reconfigurable Systems Using Flexible Petri Nets*. In Filip Zavoral, Jakub Yaghob, Pit Pichappan & Eyas El-Qawasmeh, editors: *Networked Digital Technologies*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 343–357, doi:10.1109/TASE.2010.28.
- [11] Michael Köhler-Bußmeier (2009): *Hornets: Nets within Nets Combined with Net Algebra*. In Giuliana Franceschinis & Karsten Wolf, editors: *Applications and Theory of Petri Nets*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 243–262, doi:10.1007/978-3-642-02424-5_15.
- [12] M. Llorens & J. Oliver (2004): *Structural and dynamic changes in concurrent systems: reconfigurable Petri nets*. *IEEE Transactions on Computers* 53(9), pp. 1147–1158, doi:10.1109/TC.2004.66.
- [13] Narciso Martí-Oliet, José Meseguer & Alberto Verdejo (2009): *A Rewriting Semantics for Maude Strategies*. *Electronic Notes in Theoretical Computer Science* 238(3), pp. 227–247, doi:10.1016/j.entcs.2009.05.022. Proceedings of the Seventh International Workshop on Rewriting Logic and its Applications (WRLA 2008).
- [14] José Meseguer (1998): *Membership algebra as a logical framework for equational specification*. In Francesco Parisi Presicce, editor: *Recent Trends in Algebraic Development Techniques*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 18–61, doi:10.1007/3-540-64299-4_26.
- [15] José Meseguer (1992): *Conditional rewriting logic as a unified model of concurrency*. *Theoretical Computer Science* 96(1), pp. 73–155, doi:10.1016/0304-3975(92)90182-F.
- [16] Julia Padberg & Laid Kahloul (2018): *Overview of Reconfigurable Petri Nets*. In Reiko Heckel & Gabriele Taentzer, editors: *Graph Transformation, Specifications, and Nets: In Memory of Hartmut Ehrig*, Springer International Publishing, Cham, pp. 201–222, doi:10.1007/978-3-319-75396-6_11.
- [17] Julia Padberg & Alexander Schulz (2016): *Model Checking Reconfigurable Petri Nets with Maude*. In Rachid Echahed & Mark Minas, editors: *Graph Transformation*, Springer International Publishing, Cham, pp. 54–70, doi:10.1007/978-3-319-40530-8_4.
- [18] Ulrike Prange, Hartmut Ehrig, Kathrin Hoffmann & Julia Padberg (2008): *Transformations in Reconfigurable Place/Transition Systems*. In Pierpaolo Degano, Rocco De Nicola & José Meseguer, editors: *Concurrency, Graphs and Models: Essays Dedicated to Ugo Montanari on the Occasion of His 65th Birthday*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 96–113, doi:10.1007/978-3-540-68679-8_7.
- [19] W. Reisig (1985): *Petri Nets: An Introduction*. Springer-Verlag New York, Inc., New York, NY, USA, doi:10.1007/978-3-642-69968-9.
- [20] Mark-Oliver Stehr, José Meseguer & Peter Csaba Ölveczky (2001): *Rewriting Logic as a Unifying Framework for Petri Nets*, pp. 250–303. Springer, Berlin, Heidelberg, doi:10.1007/3-540-45541-8_9.
- [21] Rüdiger Valk (2004): *Object Petri Nets*. In Jörg Desel, Wolfgang Reisig & Grzegorz Rozenberg, editors: *Lectures on Concurrency and Petri Nets: Advances in Petri Nets*, Springer, Berlin, Heidelberg, pp. 819–848, doi:10.1007/978-3-540-27755-2_23.
- [22] Emanuele Viola (2001): *E-unifiability via Narrowing*. In: *Theoretical Computer Science*, Springer, Berlin, Heidelberg, pp. 426–438, doi:10.1007/3-540-45446-2_27.

A Appendix - Maude specification of rewritable PT nets

```

fmod TRAN is
  protecting NAT .
  sort Tran .
  op t : Nat -> Tran [ctor] .
  op subscript : Tran -> Nat .
  vars N : Nat .
  eq subscript (t(N)) = N .
endfm

fmod IMATRIX is
  pr BAG{Place} * (sort Bag{Place} to BagP, sort NeBag{Place} to NeBagP, op nil to nilP) .
  pr EXT-BOOL .
  sort Imatrix .
  op [_,_,_] : BagP BagP BagP -> Imatrix [ctor] .
  op remove : Imatrix Place -> Imatrix .
  op in : Imatrix Place -> Bool .
  vars X Y Z : BagP .
  var P : Place .
  eq remove([X,Y,Z], P) = [set(X,P,0),set(Y,P,0),set(Z,P,0)] .
  eq in([X,Y,Z], P) = X[P] /= 0 or-else Y[P] /= 0 or-else Z[P] /= 0 .
endfm

fmod PT-SYS is
  pr MAP+{Tran, Imatrix} * (sort Map{Tran, Imatrix} to Net, op empty to emptyN) .
  sort System .
  op __ : Net BagP -> [System] [ctor] .
  ops In Out Inh : Net Tran -> [BagP] .
  op enabled : System Tran -> [Bool] .
  op dead : System -> Bool .
  op in : Net Place -> Bool . *** test the existence of a place
  var N : Net .
  var T : Tran .
  var P : Place .
  var I O H S : BagP .
  var Q : Imatrix .
  var K K' : NzNat .
  eq In((T |-> [I,O,H], N), T) = I .
  eq Out((T |-> [I,O,H], N), T) = O .
  eq Inh((T |-> [I,O,H], N), T) = H .
  eq in((T |-> Q, N), P) = in(Q, P) or-else in(N, P) .
  eq in(emptyN, P) = false .
  eq T |-> [I,I,H] = emptyN [metadata "null t"] .
  ceq T |-> [K . P + I, O, K' . P + H] = emptyN if K >= K' [metadata "dead t"] .
  *** ceq N K . P + S = N S if in(N, P) = false [metadata "isolated p"] .
  eq enabled( (T |-> [I,O,H], N) S, T) = I <= S and-then H >' S .
  ceq dead((T |-> Q, N) S) = false if enabled( (T |-> Q, N) S, T) .
  eq dead(N S) = true [owise] .
  cmb N S : System if N /= emptyN . *** a PT system cannot be "empty"
endfm

```



```

fth PTSYSTH is
  protecting PT-SYS .
  op m0 : -> BagP .
  op net : -> Net .
endfth

```

```

mod PT-EMU{X :: PTSYSTH} is
  var T : Tran .
  var I O H S : BagP .
  var N N' : Net .
  crl [firing] : N S => N S + O - I if T |-> [I,O,H], N' := N \ I <= S \ H >' S .
endm

```

```

mod RWPT-FMS is
  *** protecting PT-SYS .
  protecting PT-RWLIB .
  op net : -> Net .
  op m0 : -> BagP .
  op M : -> Nat . *** model's parameter
  vars N N' : Net .
  vars TL TF : Tran .
  vars P2 P3 P4 P5 PF : Place .
  var S : BagP .
  var K : NzNat .
  eq M = 50 .
  eq net = t(0) |-> [2 . p(1), 1 . p(2) + 1 . p(3), nilP], t(1) |-> [1 . p(2), 1 . p(4), 1 . p(7)],
    t(2) |-> [1 . p(3), 1 . p(5), 1 . p(8)], t(3) |-> [1 . p(4) + 1 . p(5), 1 . p(6), nilP],
    t(4) |-> [1 . p(6), 2 . p(1), nilP], t(5) |-> [1 . p(0), 1 . p(7), nilP], t(6) |-> [1 . p(0), 1 . p(8), nilP].
  eq m0 = 2 * M . p(1) + 1 . p(0) .

  crl [r1] : (N, t(0) |-> [2 . p(1), 1 . P2 + 1 . P3, nilP], t(3) |-> [1 . P4 + 1 . P5, 1 . p(6), nilP],
    TF |-> [1 . p(0), 1 . PF, nilP], TL |-> [1 . P3, 1 . P5, 1 . PF]) S + 1 . PF =>
    (N, t(0) |-> [1 . p(1), 1 . P2, nilP], t(3) |-> [2 . P4, 1 . p(6), nilP]) set(S, P3, 0) + S[P3] . P2 + 1 . p(0)
    if S[P5] = 0 .
  crl [r2] : N S => net S + 1 . p(0) + M . P3 - M . P2 - 1 . p(7) - 1 . p(8) if 1 . P2 := Out(N, t(0))
    \ 1 . P2 + 1 . P3 := Out(net, t(0)) \ dead(N S) .
  *** crl [r3] : N S => setw(N, t(0), p(9), i, 1) S if S[p(0)] = 0 \ S[p(9)] = 0 \ in(N, p(9)) = false .
endm

```

```

view Fms from PTSYSTH to RWPT-FMS is

```

```

  op m0 to m0 .
  op net to net .
endv

```

```

mod FMS-EMU is
  including PT-EMU{Fms} .
endm

```

```

fmod PT-RWLIB is
  protecting PT-SYS .
  sort Atype . *** arc type
  ops i o h : -> Atype [ctor] .
  op w : Net Tran Place Atype -> [Nat] . ***get an arc's weight
  op setw : Net Tran Place Atype Nat -> Net . ***set an arc
  op setwS : Net Tran Place Atype Nat -> Net . ***set an arc in a safe way
  ops remove : Place Net -> Net .
  var N N' : Net .
  var T : Tran .
  var P : Place .
  var I O H M : BagP .
  var Q : Imatrix .
  var K K' : NzNat .
  var Y : Nat .
  var A : Atype .
  eq w(N, T, P, i) = In(N,T)[P] .
  eq w(N, T, P, o) = Out(N,T)[P] .
  eq w(N, T, P, h) = Inh(N,T)[P] .
  ceq setw(N, T, P, i, Y) = N, T |-> [Y . P, nilP, nilP] if N[T] = undefined .
  eq setw((N, T |-> [I,O,H]), T, P, i, Y) = N, T |-> [set(I, P, Y), O, H] .
  ceq setw(N, T, P, o, Y) = N, T |-> [nilP, Y . P, nilP] if N[T] = undefined .
  eq setw((N, T |-> [I,O,H]), T, P, o, Y) = N, T |-> [I, set(O, P, Y), H] .
  ceq setw(N, T, P, h, Y) = N if N[T] = undefined .
  eq setw((N, T |-> [I,O,H]), T, P, h, Y) = N, T |-> [I, O, set(H, P, Y)] .
  eq remove(P, (T |-> Q, N) ) = T |-> remove(Q, P), remove(P, N) .
  eq remove(P, emptyN) = emptyN .
  ceq setwS(N, T, P, A, Y) = N' if N' := setw(N, T, P, A, Y) /\ N' /= emptyN .
  eq setwS(N, T, P, A, Y) = N [owise] .
endfm

```

B Appendix - Maude generic modules

```

fmod BAG{X :: TRIV} is
  protecting INT .
  sorts Bag{X} NeBag{X} .
  subsorts NeBag{X} < Bag{X} .
  op nil : -> Bag{X} [ctor] .
  op __ : Nat X$Elt -> Bag{X} [prec 35 ctor] .
  op __ : NzNat X$Elt -> NeBag{X} [ctor ditto] .
  op _['_] : Bag{X} X$Elt -> Nat [prec 23] .
  op _+_ : Bag{X} Bag{X} -> Bag{X} [prec 39 ctor assoc comm id: nil] .
  op _+_ : NeBag{X} Bag{X} -> NeBag{X} [ctor ditto] .
  op _-_ : Bag{X} Bag{X} -> Bag{X} [prec 41 gather (E e)] .
  op _<=_ : Bag{X} Bag{X} -> Bool [prec 43] .
  op _>'_ : Bag{X} Bag{X} -> Bool [prec 43] .
  op set : Bag{X} X$Elt Nat -> Bag{X} .
  vars X Y : X$Elt .
  vars N M : NzNat .
  var K : Nat .

```

```

var B B' : Bag{X} .
eq (N . X + B)[X] = N .
eq B[X] = 0 [owise] .

eq 0 . X = nil .
eq N . X + M . X = (N + M) . X .

eq B - nil = B .
eq N . X + B - M . X + B' = if N > M then N - M . X + (B - B') else B - B' fi .
eq B - M . X + B' = B - B' [owise] .

ceq N . X + B <= B' = false if N > B'[X] .
eq B <= B' = true [owise] .

ceq N . X + B >' B' = false if N <= B'[X] .
eq B >' B' = true [owise] .

eq set(B + N . X, X, K) = B + K . X .
eq set(B, X, K) = B + K . X [owise] .
endfm

fmod MAP+{X :: TRIV, Y :: TRIV} is
  sorts Entry{X,Y} Map{X,Y} .
  subsorts Entry{X,Y} < Map{X,Y} .
  op _|->_ : X$Elt Y$Elt -> Entry{X,Y} [ctor] .
  op empty : -> Map{X,Y} [ctor] .
  op _'[_] : [Map{X,Y}] [Map{X,Y}] -> [Map{X,Y}] [assoc comm id: empty ctor prec 121] .
  op undefined : -> [Y$Elt] [ctor] .
  op insert : X$Elt Y$Elt [Map{X,Y}] -> [Map{X,Y}] .
  op insert : X$Elt Y$Elt Map{X,Y} -> Map{X,Y} .
  op remove : X$Elt [Map{X,Y}] -> [Map{X,Y}] .
  op remove : X$Elt Map{X,Y} -> Map{X,Y} .
  op _'[_] : [Map{X,Y}] X$Elt -> [Y$Elt] [prec 23] .
  op $hasMapping : [Map{X,Y}] X$Elt -> Bool .
  op hasduplicate : [Map{X,Y}] -> Bool .

var D : X$Elt .
vars R R' : Y$Elt .
var M : [Map{X,Y}] .

eq $hasMapping((M, D |-> R), D) = true .
eq $hasMapping(M, D) = false [owise] .
eq insert(D, R, (M, D |-> R')) = if $hasMapping(M, D) then insert(D, R, M) else (M, D |-> R) fi .
eq insert(D, R, M) = (M, D |-> R) [owise] .
eq remove(D, (M, D |-> R)) = remove(D, M) .
eq remove(D, M) = M [owise] .
eq (M, D |-> R) [D] = if $hasMapping(M, D) then undefined else R fi .
eq M [D] = undefined [owise] .
eq hasduplicate((D |-> R, D |-> R', M)) = true .
eq hasduplicate(M) = false [owise] .
cmb M : Map{X,Y} if hasduplicate(M) = false .
endfm

```