A Classification of the Expressive Power of Well-structured Transition Systems

Parosh Aziz Abdulla[1], Giorgio Delzanno[2], and Laurent Van Begin[3][*]

[1] Uppsala University, Sweden, `parosh@it.uu.se`
[2] Università di Genova, Italy, `giorgio@disi.unige.it`
[3] Université Libre de Bruxelles, Belgium, `lvbegin@ulb.ac.be`

Abstract. We compare the expressive power of a class of well-structured transition systems that includes relational automata, Petri nets, lossy channel systems, constrained multiset rewriting systems, and data nets. For each one of these models we study the class of languages generated by labelled transition systems describing their semantics. We consider here two types of accepting conditions: coverability and reachability of a fixed a priori configuration. In both cases we obtain a strict hierarchy in which constrained multiset rewriting systems is the the most expressive model.

Keywords. Expressiveness, well-structured systems, language theory.

1 Introduction

The theory of well-structured transition systems [1, 13] is a powerful tool for studying the decidability of verification problems of infinite-state systems. A system is well-structured when its transition relation is monotonic with respect to a well-quasi ordering defined over configurations. A well-known example of well-structured system is that of *Petri nets* [21] equipped with marking inclusion [1, 13]. For a well-structured transition system, the *coverability problem* can be decided by the symbolic backward reachability algorithm scheme proposed in [1]. Since checking safety properties can be translated into instances of the coverability problem, an algorithm for coverability like that proposed in [1] can be used for automatic verification of an infinite-state system. This connection has been exploited in order to develop automatic verification procedures for several infinite-state models:

- *relational automata* (RA) [8], an abstract models of imperative programs with integer valued variables;
- Reset and transfer nets [11, 10], i.e. Petri nets extended with *whole-place operations* that atomically operate on the whole set of tokens in a place;
- *lossy (FIFO) channel systems* (LCS) [4, 7], an abstract models of unreliable communication systems;

- *constrained multiset rewriting systems* (CMRS) [2], an extension of Petri nets in which tokens are colored with natural numbers and in which transitions are guarded by conditions on colors;
- *affine well-structured nets* (aWSNs) [18] a generalization of Petri nets and transfer/resetnets in which the firing of a transition is split into three steps: subtraction, multiplication, and addition of black tokens. Multiplication is a whole-place operation that generalizes transfer and reset arcs;
- *data nets* (DNs) [19], a generalization of aWSNs in which subtraction, multiplication and addition are defined on tokens that carry data taken from an infinite, ordered domain. Conditions on data values can be used here to restrict the type of tokens on which apply whole-place operations. DNs are a natural extension of CMRS with whole-place operations on colored tokens.

Although several efforts have been spent on studying the expressive power of extensions of Petri nets like reset and transfer nets [11, 12, 14]), a comparison of the relative expressiveness of the class of well-structured transition systems is still missing. Such a comparison is a challenging research problem with a possible practical impact. Indeed, it can be useful to extend the applicability of a verification method (e.g. a particular instance of the scheme of [1]) to an entire class of models.

In this paper we apply tools of language theory to formally compare the expressive power of a large class of well-structured infinite-state systems that includes extensions of Petri nets, constrained multiset rewriting systems, lossy channel systems, relational automata, and data nets. To achieve the goal, for each one of these models we study the class of languages generated by labeled transition systems describing their semantics. We consider here two types of accepting conditions: coverability (with respect to a fixed ordering) and reachability of a given configuration. Two models are considered to be equivalent if they generate the same class of languages.

For coverability accepting conditions, we obtain the following classification.

- We show that, differently from nets with indistinguishable tokens, whole-place operations do not augment the expressive power of models in which tokens carry data taken from an ordered domain. The proof is based on a weak, effectively constructible encoding of data nets into CMRS that can be used to reduce the coverability problem from one model to the other. As a corollary, we have that the symbolic backward reachability algorithm for solving the coverability problem in CMRS described in [2] can also be applied to data nets.

  By slightly extending the CMRS encoding, we also prove that the coverability problem remains decidable for different extensions of data nets. In particular we consider data net transitions that select data that *must* be fresh (in [19] a transition selects values that *may* be fresh).
- We prove that lossy channel systems are equivalent to a syntactic fragment of constrained multiset rewriting, we named Γ_0. The fragment Γ_0 is obtained by restricting conditions of a rule in such a way that equalities cannot be

used as guards. Furthermore, we prove that lossy channel systems are strictly less expressive than the full model of constrained multiset rewriting systems. We then show that Petri nets are equivalent to a syntactic fragment of constrained multiset rewriting systems, we named Γ_1, obtained by considering nullary predicates only.

– We prove that aWSNs are strictly more expressive than Petri nets and strictly less expressive than LCSs, thus separating Petri nets from LCSs with respect to their relative expressive power. Furthermore, we prove that $aWSNs$ are as expressive as transfer/reset nets. This result show that the inclusion between the coverability languages of transfer/reset nets and LCS proved in [3] is strict.

– We prove that relational automata are equivalent to a syntactic fragment of constrained multiset rewriting, we named Γ_2, obtained by imposing an upper bound on the size (number of predicates) of reachable configurations.

– Finally, we prove that Γ_2 generates the class of regular languages. This implies that relational automata are strictly less expressive than Petri nets.

For reachability accepting conditions, we obtain a slightly different classification. First, we prove that Γ_0 is equivalent to constrained multiset rewriting systems and *two counter machines*. Thus, with reachability acceptance, Γ_0 and constrained multiset rewriting systems turn out to be strictly more expressive than lossy channel systems. On the contrary, Γ_1 is still equivalent to Petri nets and strictly less expressive than Γ_0 and Γ_2 is still equivalent to relational automata and to finite automata. Finally, we show that lossy channel systems and Petri nets define incomparable classes of languages.

Concerning related work, the relative expressiveness of well-structured systems has been investigated for a limited number of extensions of Petri nets with reset, transfer, and non-blocking arcs in [12, 14]. Classical results on finite and infinite languages generated by Petri nets can be found, e.g., in [15]. A classification of infinite-state systems in terms of structural properties and decidable verification problems is presented in [16]. The classification is extended to well-structured systems in [6]. A classification of the complexity of the decision procedures for coverability is studied in [19]. In contrast with the aforementioned work, we provide here a strict classification of the expressive power of several well-structured transition systems built with the help of tools of language theory. An extended version of the present paper is available as technical report [3].

Outline In Section 2, we give some preliminary notions on well-structured transition systems. In Section 3, we give some first results on the class of languages accepted by constrained multiset rewriting systems. In Section 6, 7, and 8, we compare the class of languages recognized by constrained multiset rewriting systems and, respectively, lossy channel systems, Petri nets, and relational automata. Finally, in Section 9 we discuss some final remarks.

2 Wsts and Languages with Coverability Acceptance

In this section we recall some definitions taken from [1]. A *transition system* is a tuple $T = (S, R)$ where S is a (possibly infinite) set of configurations, R is a finite set of transitions where each $\xrightarrow{\sigma} \in R$ is a binary relation over S, i.e. $\xrightarrow{\sigma} \subseteq S \times S$. We use $\gamma \xrightarrow{\sigma} \gamma'$ to denote $(\gamma, \gamma') \in \xrightarrow{\sigma}$, and $\gamma \xrightarrow{\rho_1 \ldots \rho_k} \gamma'$ to denote that there exist $\gamma_1, \ldots, \gamma_{k-1}$ such that $\gamma \xrightarrow{\rho_1} \gamma_1 \ldots \xrightarrow{\rho_{k-1}} \gamma_{k-1} \xrightarrow{\rho_k} \gamma'$. A quasi-ordering $(S, \preceq)$ is a *well-quasi ordering* if for any infinite sequence $s_1 s_2 \ldots s_i \ldots$ there exist indexes $i < j$ such that $s_i \preceq s_j$. A transition system $T = (S, R)$ is *well-structured* with respect to a quasi-order $\preceq$ on S iff: (i) $\preceq$ is a well-quasi ordering; (ii) for any $\xrightarrow{\sigma} \in R$ and $\gamma_1, \gamma_1', \gamma_2$ s.t. $\gamma_1 \preceq \gamma_1'$ and $\gamma_1 \xrightarrow{\sigma} \gamma_2$, there exists γ_2' s.t. $\gamma_1' \xrightarrow{\sigma} \gamma_2'$ and $\gamma_2 \preceq \gamma_2'$, i.e., T is *monotonic*. We use $T = (S, R, \preceq)$ to indicate a *well-structured transition system* (wsts for short).

To formalize the comparison between models, a wsts $T = (S, R, \preceq)$ can be viewed as a language acceptor. For this purpose, we assume a finite alphabet Σ and a labelling function $\lambda : R \mapsto \Sigma$ that associates to each transition of R a symbol of $\Sigma \cup \{\epsilon\}$, where ϵ denotes the empty sequence ($w \cdot \epsilon = \epsilon \cdot w = w$ for any $w \in \Sigma^*$). In the following, we use $\gamma_1 \xrightarrow{w} \gamma_2$ with $w \in \Sigma^*$ to denote that $\gamma_1 \xrightarrow{\rho_1 \ldots \rho_k} \gamma_2$ and $\lambda(\xrightarrow{\rho_1}) \cdots \lambda(\xrightarrow{\rho_k}) = w$. Furthermore, we associate to T an *initial* configuration $\gamma_{init} \in S$ and a *final* configuration $\gamma_{acc} \in S$ and assume an accepting relation $\bowtie : S \times S$. For a fixed accepting relation $\bowtie$, we define the language accepted (generated) by $T = (S, R, \preceq, \gamma_{init}, \gamma_{acc})$ as:

$$L(T) = \{w \in \Sigma^* \mid \gamma_{init} \xrightarrow{w} \gamma \ \text{ and } \gamma_{acc} \bowtie \gamma\}$$

In this paper we consider two types of accepting relations:

- *Coverability*: the accepting relation $\bowtie$ is defined as $\gamma_{acc} \preceq \gamma$.
- *Reachability*: the accepting relation $\bowtie$ is defined as $\gamma_{acc} = \gamma$.

Let $\mathcal{M}$ be a wsts model (e.g. Petri nets) and let T be one of its instances (i.e. a particular net). We define $L_c(T)$, resp $L_r(T)$, as the language accepted by T with accepting relation $\bowtie_c$, resp. $\bowtie_r$. We say that L is a c-language, resp. r-language, of $\mathcal{M}$ if there is an instance T of $\mathcal{M}$ such that $L = L_c(T)$, resp. $L = L_r(T)$. We use $L_c(\mathcal{M})$, resp. $L_r(\mathcal{M})$, to denote the class of c-languages, resp. r-languages, of $\mathcal{M}$. Finally, we use $\mathcal{L}_1 \not\sim \mathcal{L}_2$ to denote that $\mathcal{L}_1$ and $\mathcal{L}_2$ are incomparable classes of languages.

Given a wsts $T = (S, R, \preceq)$ with labels in $\Sigma \cup \{\epsilon\}$, a *lossy* version of T is a wsts $T' = (S, R', \preceq)$ for which there exists a bijection $h : R \mapsto R'$ such that $\xrightarrow{\rho} \in R$ and $\xrightarrow{h(\rho)}$ have the same label, $\xrightarrow{\rho} \subseteq \xrightarrow{h(\rho)}$ and if $\gamma \xrightarrow{h(\rho)} \gamma'$, then $\gamma \xrightarrow{\rho} \gamma''$ with $\gamma' \preceq \gamma''$. In other words, in a lossy version of a wsts the set of reachable configurations contains configurations that are smaller than those of the original model. The next lemma states an important property used in the remainder of the paper.

Lemma 1. *For any lossy version T' of a wsts T, we have that $L_c(T) = L_c(T')$.*

3 Constrained Multiset Rewriting Systems (CMRS)

In this section we recall the main definitions and prove the first results for *constrained multiset rewriting systems* [2]. Let us first give some preliminary definitions. We use $\mathbb{N}$ to denote the set of natural numbers and $\overline{n}$ to denote the interval $[0, \ldots, n]$ for any $n \in \mathbb{N}$. We assume a set $\mathbb{V}$ of variables which range over $\mathbb{N}$, and a set $\mathbb{P}$ of unary predicate symbols. For a set A, we use A^* and $A^\otimes$ to denote the sets of (finite) words and (finite) multisets over A respectively. Sometimes, we write multisets as lists, so $[1, 5, 5, 1, 1]$ represents a multiset with three occurrences of 1 and two occurrences of 5; $[\ ]$ represents the empty multiset. We use the usual relations and operations such as $\leq$ (inclusion), $+$ (union), and $-$ (difference) on multisets. For a set $V \subseteq \mathbb{V}$, a *valuation Val* of V is a mapping from V to $\mathbb{N}$. A *condition* is a finite conjunction of *gap order* formulas of the forms: $x <_c y$, $x \leq y$, $x = y$, $x < c$, $x > c$, $x = c$, where $x, y \in \mathbb{V}$ and $c \in \mathbb{N}$. Here $x <_c y$ stands for $x + c < y$. We often use $x < y$ instead of $x <_0 y$. Sometimes, we treat a condition ψ as a set, and write e.g. $(x <_c y) \in \psi$ to indicate that $x <_c y$ is one of the conjuncts in ψ. We use *true* to indicate an empty set of conditions. A *term* is of the form $p(x)$ where $p \in \mathbb{P}$ and $x \in \mathbb{V}$. A *ground term* is of the form $p(c)$ where $p \in \mathbb{P}$ and $c \in \mathbb{N}$. We sometimes say that a predicate symbol is *nullary* to mean that its parameter is not relevant (hence may be omitted).

A *constrained multiset rewriting system (CMRS) $\mathcal{S}$* consists of a finite set of *rules* each of the form $L \rightsquigarrow R : \psi$, where L and R are multisets of terms, and ψ is a condition. We assume that ψ is consistent (otherwise, the rule is never enabled). For a valuation *Val*, we use $Val(\psi)$ to denote the result of substituting each variable x in ψ by $Val(x)$. We use $Val \models \psi$ to denote that $Val(\psi)$ evaluates to *true*. For a multiset T of terms we define $Val(T)$ as the multiset of ground terms obtained from T by replacing each variable x by $Val(x)$. A *configuration* is a multiset of ground terms. Each rule $\rho = L \rightsquigarrow R : \psi \in \mathcal{S}$ defines a relation between configurations. More precisely, $\gamma \xrightarrow{\rho} \gamma'$ if and only if there is a valuation *Val* s.t. the following conditions are satisfied: (*i*) $Val \models \psi$, (*ii*) $\gamma \geq Val(L)$, and (*iii*) $\gamma' = \gamma - Val(L) + Val(R)$.

Example 1. Consider the CMRS rule:

$$\rho = [p(x) , q(y)] \rightsquigarrow [q(z) , r(x) , r(w)] \;\; : \;\; \{x + 2 < y , x + 4 < z , z < w\}$$

A valuation which satisfies the condition is $Val(x) = 1$, $Val(y) = 4$, $Val(z) = 8$, and $Val(w) = 10$.

A CMRS configuration is a multisets of ground terms, e.g., $[p(1), p(3), q(4)]$. Therefore, we have that $[p(1), p(3), q(4)] \xrightarrow{\rho} [p(3), q(8), r(1), r(10)]$.

Let us fix a CMRS $\mathcal{S}$ operating on a set of predicate symbols $\mathbb{P}$. Let *cmax* be the maximal constant which appears in the rules of $\mathcal{S}$; *cmax* is equal to 0 if there are no constant in $\mathcal{S}$. We now define an ordering $\preceq_c$ on configurations extracted from the ordering defined in [2] to solve the coverability problem.

Definition 1. *Given a configuration γ, we define the* index *of γ, index(γ), to be a word of the form $D_0 \cdots D_{cmax} \; d_0 \; B_0 \; d_1 \; B_1 \; d_2 \; \cdots \; d_n \; B_n$ where*

- $D_0, \ldots, D_{cmax}, B_0, \ldots, B_n \in \mathbb{P}^\otimes$ *and* $d_0, d_1, \ldots, d_n \in \mathbb{N} \setminus \{0\}$;
- B_i *must not be empty for* $0 \le i \le n$;
- *for each* $p \in \mathbb{P}$, D_i *contains* k *occurrences of predicate* p *iff* $p(i)$ *occurs* k *times in* γ *for* $0 \le i \le cmax$;
- *given* $v_0 = cmax + d_0$, *for each* $p \in \mathbb{P}$, B_0 *contains* k *occurrences of predicate* p *iff* $p(v_0)$ *occurs* k *times in* γ;
- *given* $v_{i+1} = v_i + d_{i+1}$, *for each* $p \in \mathbb{P}$, B_{i+1} *contains* k *occurrences of predicate* p *iff* $p(v_{i+1})$ *occurs* k *times in* γ *for all* $0 \le i < n$;
- *for all* $p(v) \in \gamma$ *with* $v > cmax$, *there exists* $i : 0 \le i \le n$ *such that* $v = cmax + d_0 + d_1 + \ldots + d_i$.

The ordering $\preceq_c$ is defined as follows.

Definition 2. *Let* $D_0 \ D_1 \ \cdots \ D_{cmax} \ d_0 \ B_0 \ d_1 \ B_1 \ d_2 \ \cdots \ d_n \ B_n$ *be the index of a configuration* γ_1 *and* $D_0' \ D_1' \ \cdots \ D_{cmax}' \ d_0' \ B_0' \ d_1' \ B_1' \ d_2' \ \cdots \ d_m' \ B_m'$ *be the index of a configuration* γ_2. *Then,* $\gamma_1 \preceq_c \gamma_2$ *iff* $D_i \le D_i'$ *for* $0 \le i \le cmax$ *and there exists a monotone injection* $h : \overline{n} \mapsto \overline{m}$ *such that* $B_0 \le B_{h(0)}'$, $B_i \le B_{h(i)}'$, $d_0 \le \sum_{k=0}^{h(0)} d_k'$, *and* $d_i \le \sum_{k=h(i-1)+1}^{h(i)} d_k'$ *for* $1 \le i \le n$.

In the rest of the paper we assume that the values appearing in the initial configuration γ_{init} and in the accepting configuration γ_{acc} are smaller or equal than *cmax*. The ordering $\preceq_c$ is obtained by composing string embedding and multiset inclusion. From standard properties of orderings, it follows that $\preceq_c$ is a well-quasi ordering. Furthermore, a CMRS is monotonic with respect to corresponding ordering $\preceq_c$. The following property then holds.

Proposition 1 ([2]). *A CMRS* $\mathcal{S}$ *equipped with* $\preceq_c$ *is well-structured.*

3.1 A Symbolic Algorithm for Testing Coverability

In this section we give an overview of the algorithm for solving the coverability problem based on the generic backward analysis algorithm presented in [1]. The difficult challenge in applying this methodology is to invent a symbolic representation (called *constraints*) which allows effective implementation of each step, and which guarantees termination of the algorithm.

The algorithm operates on *constraints*, where each constraint ϕ characterizes an infinite set $[\![\phi]\!]$ of configurations. A *constraint* ϕ is of the form $T : \psi$ where T is a multiset of terms and ψ is a condition. The constraint characterizes the (upward closed) set $[\![\phi]\!] = \{\gamma | \exists Val. (Val \models \psi) \wedge (Val(T) \le \gamma)\}$ of configurations. Notice that if ψ is inconsistent, then $[\![\phi]\!]$ is empty. Such a constraint can be safely discarded in the reachability algorithm presented below. Therefore, we assume in the sequel that all conditions in constraints are consistent. We define $Var(\phi) = Var(T) \cup Var(\psi)$. Observe that the coverability problem can be reduced to constraint reachability. More precisely, $\gamma_{init} \xrightarrow{*} p_{fin}$ is equivalent to $\gamma_{init} \xrightarrow{*} \phi_{fin}$ where ϕ_{fin} is the constraint $[p_{fin}(x) : true$.

For constraints ϕ_1, ϕ_2, we use $\phi_1 \sqsubseteq \phi_2$ to denote that ϕ_1 is *entailed* by ϕ_2, i.e., $[\![\phi_1]\!] \supseteq [\![\phi_2]\!]$. For a constraint ϕ, we define $Pre(\phi)$ to be a finite set

of constraints which characterize the configurations from which we can reach a configuration in ϕ through the application of a single rule. In other words $\bigcup_{\phi_1 \in Pre(\phi)} \llbracket \phi_1 \rrbracket = \{\gamma \mid \exists \gamma' \in \llbracket \phi \rrbracket. \, \gamma \longrightarrow \gamma'\}$.

For instance, given $\phi_1 = [p(x_1), q(x_2), q(x_3)] : \{x_1 <_2 x_2, x_2 <_1 x_3\}$, and the configurations $\gamma_1 = [p(2), q(8), q(5), p(1)]$ and $\gamma_2 = [p(2), q(2), q(5), p(1)]$. Then $\gamma_1 \in \llbracket \phi_1 \rrbracket$ and $\gamma_2 \notin \llbracket \phi_1 \rrbracket$. Consider now $\phi_2 = [p(y_1), \, q(y_2)] : \{y_1 < y_2\}$ and $\phi_3 = [p(y_1), \, q(y_2)] : \{y_1 <_4 y_2\}$. Then $\phi_2 \sqsubseteq \phi_1$ and $\phi_3 \not\sqsubseteq \phi_1$.

Given an instance of the coverability problem, defined by γ_{init} and the constraint ϕ_{fin} corresponding to p_{fin}, the symbolic algorithm performs a fixpoint iteration starting from ϕ_{fin} and repeatedly applying Pre on the generated constraints. The iteration stops if either (i) we generate a constraint ϕ with $\gamma_{init} \in \llbracket \phi \rrbracket$; or (ii) we reach a point where, for each newly generated constraint ϕ, there is a constraint ϕ' generated in a previous iteration with $\phi' \sqsubseteq \phi$. We give a positive answer to the coverability problem in the first case, while we give a negative answer in the second case.

In [2] we show computability of membership, entailment, and define an effective predecessor operator for constraints. To give an idea of these definition, let $\mathcal{S}$ be a CMRS and ϕ_2 be a constraint. We define $Pre(\phi_2) = \bigcup_{\rho \in \mathcal{S}} Pre_\rho(\phi_2)$, where $Pre_\rho(\phi_2)$ describes the effect of running the rule ρ backwards from the configurations in ϕ_2. Let $\rho = (L \rightsquigarrow R : \psi)$ and $\phi_2 = (T_2 : \psi_2)$. Let W be any set of variables such that $|W| = |Var(\phi_2) \cup Var(\rho)|$. We define $Pre_\rho(\phi_2)$ to be the set of constraints of the form $T_1 : \psi_1$, such that there are renamings Ren, Ren_2 of $Var(\rho)$ and $Var(\phi_2)$ respectively to W, and

$\bullet \; T_1 = Ren_2(T_2) \ominus Ren(R) + Ren(L)$ $\qquad \bullet \; \psi_1 = Ren(\psi) \wedge Ren_2(\psi_2)$

Example 2. For instance, consider the constraint $\phi = [q(x_1), \, s(x_2), \, r(x_2)] : \{x_1 < x_2\}$ and the rule $\rho = [p(y_1), \, p(y_3)] \rightsquigarrow [q(y_2), \, r(y_3)] : \{y_3 < y_2\}$. Fix $W = \{w_1, w_2, w_3, w_4, w_5\}$, and define $Ren_2 = (x_1 \mapsto w_1, x_2 \mapsto w_2)$, and $Ren = (y_1 \mapsto w_3, y_2 \mapsto w_1, y_3 \mapsto w_4)$.
Then one member of Pre_ρ is given by
$[s(w_2), r(w_2), p(w_3), p(w_4)] : \{w_1 < w_2 \, , \, w_4 < w_1\}$, which after normalization becomes $[s(w_2), r(w_2), p(w_3), p(w_4)] : \{w_4 <_1 w_2\}$.

The termination of the algorithm is obtained by a non-trivial application of a methodology based on the theory of well- and better-quasi orderings described in [2].

3.2 Three Interesting Fragments of CMRS: Γ_0, Γ_1 and Γ_2

In this section we defined three fragments of CMRS that we use as a technical tool for comparisons with other wsts.

The fragment Γ_0 In the fragment Γ_0 of CMRS every rule $L \rightsquigarrow R : \psi$ satisfies the following conditions: every variable x occurs at most once in L and at most once in R, and ψ does not contain equality constraints. As an example, $[p(x), r(y)] \rightsquigarrow [q(x), r(z)] : x < y, y < z$ is a rule in Γ_0, whereas $[p(x), q(x)] \rightsquigarrow [q(y)] : true$ and $[p(x)] \rightsquigarrow [q(y), r(y)] : true$ are not in Γ_0.

The fragment Γ_1 The fragment Γ_1 is obtained by restricting CMRS to nullary predicates only (i.e. predicates with no parameters).

The fragment Γ_2 The fragment Γ_2 is defined as follows. Let us first use $|B|$ to denote the cardinality of a multiset B. Γ_2 is the fragment of CMRS in which a rule $L \rightsquigarrow R : \psi$ satisfies the condition $|R| \leq |L|$. In other words, in Γ_2 the cardinality of a reachable configuration is always bounded by the cardinality of the initial configuration.

In the rest of the paper we show that these three fragments have the same expressive power resp. of Lossy FIFO Channel Systems, Petri nets, and Integral Relational Automata. To prove this statement, it is useful to isolate properties of CMRS and of these fragments with respect to coverability acceptance.

3.3 Properties of CMRS

In this section we prove some properties of CMRS needed in the rest of the paper.

We first introduce some new terminology. We say that a configuration γ with $index(\gamma) = D_0 \ldots D_{cmax} B_0 d_0 \ldots d_n B_n$ is *linear* if B_i is a singleton multiset for $i : 0 \leq i \leq n$. We also say that an execution $\gamma_0 \xrightarrow{\rho_1} \ldots \xrightarrow{\rho_k} \gamma_k$ is linear whenever γ_i is linear for any $i : 0 \leq i \leq k$.

Furthermore, we say that γ is *cmax-bounded* if $index(\gamma) = D_0 \ldots D_{cmax}$, i.e. all the natural numbers in γ are between 0 and $cmax$.

An important property of CMRS is related to the possibility of lifting an execution from an initial *cmax*-bounded configuration γ_{init} to a configuration γ to a new execution leading from γ_{init} to a configuration with larger "gaps" (for values greater than *cmax*) than those in γ. In the rest of the section we often refer to the *index* of a configuration γ, a notion defined in the first section of the appendix (definition of the wqo on CMRS configurations).

We first define a restriction $\prec$ of the relation $\preceq_c$ in which we require that the distribution of predicates in two configurations has the same structure but larger gaps.

Formally, $\gamma_1 \prec \gamma_2$ holds iff the following conditions are satisfied:

- $index(\gamma_1) = D_0 \ldots D_{cmax} d_0 B_0 d_1 \ldots d_n B_n$
- $index(\gamma_2) = D_0 \ldots D_{cmax} d_0' B_0 d_1' \ldots d_n' B_n$
- $d_i' \geq d_i$ for any $i : 0 \leq i \leq n$.

We say that an execution $\gamma_0 \xrightarrow{\rho_1} \ldots \gamma_1 \xrightarrow{\rho_k} \gamma_k$ subsumes an execution $\gamma_0' \xrightarrow{\rho_1'} \gamma_1' \ldots \xrightarrow{\rho_k'} \gamma_k'$ if for all $i : 0 \leq i \leq k, \gamma_i' \prec \gamma_i$ and for all $i : 1 \leq i \leq k, \rho_i = \rho_i'$.

The following property then holds. The proof is in appendix A.

Proposition 2. *Let a CMRS with initial cmax-bounded configuration γ_{init}. For any execution $e = \gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma$, for any configuration γ' such that $\gamma \prec \gamma'$, there exists an execution $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma'$ that subsumes e.*

Now, we introduce the notion of linearization of a configuration a technical definition we use to prove other interesting properties of the class of languages accepted by CMRS.

Given a configuration γ with

$$index(\gamma) = D_0 \ldots D_{cmax} d_0 B_0 \ldots d_i B_i + [p] d_{i+1} B_{i+1} \ldots d_n B_n$$

where B_i is not empty, we say that γ' is a linearization of γ if

$$index(\gamma') = D_0 \ldots D_{cmax} d_0' B_0 \ldots d_i' B_i d[p] d_{i+1}' B_{i+1} \ldots d_n' B_n$$

such that

- $\forall 0 \leq j \leq n : d_j' \geq d_j$;
- $d \geq 1$.

The following lemmas then hold. The proof is in appendix A.

Lemma 2. *Let $\mathcal{S}$ be a Γ_0 model with initial cmax-bounded configuration γ_{init}. Suppose there exists a linear execution $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma_k$, $\gamma_k \xrightarrow{\rho} \gamma$ and γ is not linear. Then, there exists a (possibly different) linear execution $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma_k'$ such that $\gamma_k' \xrightarrow{\rho} \gamma'$ and γ' is a linearisation of γ_k.*

Lemma 3. *For a Γ_0 $\mathcal{S}$, let γ_1 and γ_2 be two configurations such that γ_2 is a linearisation of γ_1, $\gamma_1 \xrightarrow{\rho_1 \ldots \rho_k} \gamma_3$ implies there exists γ_4 such that $\gamma_2 \xrightarrow{\rho_1 \ldots \rho_k} \gamma_4$.*

Given a CMRS $\mathcal{S}$ of Γ_0 with initial *cmax*-bounded linear configuration γ_{init} and accepting *cmax*-bounded linear configuration γ_{acc}, we define

$$L_c^{lin}(\mathcal{S}) = \{w \mid \gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma_k \text{ is linear}, \gamma_{acc} \preceq_c \gamma_k, \lambda(\rho_1) \cdots \lambda(\rho_k) = w\}$$

From Lemma 2 and 3, we obtain the following proposition.

Proposition 3. *For all CMRS $\mathcal{S}$ of Γ_0 with an initial cmax-bounded linear configuration γ_{init} and cmax-bounded linear accepting configuration γ_{acc}, we have $L_c(\mathcal{S}) = L_c^{lin}(\mathcal{S})$.*

Proof.
$\supseteq$: Immediate.

$\subseteq$: To simplify the presentation, let us assume that $\forall L \rightsquigarrow R : \psi \in \mathcal{S} :$ for each variable x that appears in $L + R$: either $(x = c) \in \psi$ $(0 \leq c \leq cmax)$ or $(x > cmax) \in \psi$. This assumption implies that the effect of a rule ρ is constant if we only consider ground terms $p(x)$ with $x : 0 \leq x \leq cmax$.

Suppose that $\gamma_{init} \xrightarrow{\rho_1} \gamma_1 \xrightarrow{\rho_2} \ldots \xrightarrow{\rho_k} \gamma_k$ with $\gamma_{acc} \preceq_c \gamma_k$.

Now suppose that γ_1 is not linear. Applying Lemma 2, we have a linearization γ_1' of γ_1 such that $\gamma_{init} \xrightarrow{\rho_1} \gamma_1'$. Furthermore, following Lemma 3, we have $\gamma_1' \xrightarrow{\rho_2 \ldots \rho_k} \gamma_k'$. Iterating the reasoning, we obtain a linear configuration γ_1'' such that $\gamma_{init} \xrightarrow{\rho_1} \gamma_1''$ and $\gamma_1'' \xrightarrow{\rho_2 \ldots \rho_k} \gamma_k''$.

Repeating the reasoning for the other intermediate configurations, we conclude that there exists a linear execution $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma_k'''$.

From our hypothesis, $\rho_1 \ldots \rho_k$ has constant effect if we only consider ground terms $p(x)$ with $x : 0 \leq x \leq cmax$. Hence, we have that

$$\sum_{p(n) \in \gamma_k, 0 \leq n \leq cmax} [p(n)] = \sum_{p(n) \in \gamma_k''', 0 \leq n \leq cmax} [p(n)]$$

and $\gamma_{acc} \preceq_c \gamma_k$ implies that $\gamma_{acc} \preceq_c \gamma_k'''$ since γ_{acc} is $cmax$-bounded. We conclude that $\lambda(\rho_1) \cdots \lambda(\rho_k) \in L_c(\mathcal{S})$ implies $\lambda(\rho_1) \cdots \lambda(\rho_k) \in L_c^{lin}(\mathcal{S})$. $\qquad\square$.

4 Expressive Power of CMRS

We are ready now to give a first characterization for the expressive power of CMRS. In [14, Prop. 4], the authors show that there exists a recursively enumerable (RE) language that cannot be recognized by any wsts with coverability acceptance. Hence, the following proposition holds.

Theorem 1. $L_c(CMRS) \subset RE$.

With reachability as accepting condition, CMRS recognize instead the class of recursively enumerable languages (RE).

Theorem 2. $L_r(CMRS) = RE$.

Proof. We prove that CMRS can weakly simulate 2-counter machines. In the proof we also show that repeated reachability is undecidable for CMRS. We recall the model of a *2-counter machine (CM)* which is pair (Q, δ), where Q is a finite set of states, and δ is the transition function. A transition is of the form (q_1, op, q_2), where $q_1, q_2 \in Q$, and op is either an *increment* (of the form cnt_1++ or cnt_2++); a *decrement* (of the form cnt_1-- or cnt_2--); or a *zero-testing* (of the form $cnt_1 = 0$? or $cnt_2 = 0$?). Operations and tests on counters have their usual semantics, assuming that the values of counters are natural values. In particular, decrement on a counter equal to zero is blocking. A 2-counter machine accepts an execution if it ends into the state q_{fin}. A *lossy 2-counter machine (LCM)* is of the same form as a counter machine. The difference in semantics is in the zero-testing operation. More precisely, the zero-testing of cnt_1 is simulated by resetting the value cnt_1 to zero, and decreasing the value of cnt_2 by an arbitrary natural number (possibly 0). The zero-testing of cnt_2 is performed in a similar manner.

Assume an LCM $\mathcal{M} = (Q, \delta)$. We shall construct a CMRS $\mathcal{S}$ which will simulate simulate $\mathcal{M}$. The simulation of $\mathcal{M}$ occurs in a sequence of *phases*. During each phase, $\mathcal{S}$ simulates increment and decrement transitions of $\mathcal{M}$. Each phase is indexed by a natural number which is incremented at the end of the phase. As soon as $\mathcal{M}$ performs a zero-testing of a counter, $\mathcal{S}$ enters an intermediate stage. After conclusion of the intermediate stage, a new phase is started and the index phase is increased.

The set of predicates symbols in $\mathcal{S}$ is divided into three groups:

- Two nullary predicate symbols q and q' for each $q \in Q$. We use q' during the intermediate stages of the simulation;
- Two predicate symbols cnt_1 and cnt_2, which encode the values of cnt_1 and cnt_2 respectively;
- A predicate *phase* whose argument carries the index of the current phase. Furthermore, we use a predicate symbol *phase'* to store the index of the previous phase during the intermediate stages of the simulation.

A configuration of $\mathcal{S}$ contains, during a given phase of the simulation, the following ground terms:

- A term of the form q which encodes the current state of $\mathcal{M}$.
- A term of the form $phase(c)$ where c is the index of the current phase.
- Terms of the form $cnt_1(c)$ where c is the index of the current phase. The number of such terms encodes the current value of cnt_1. There are also a number of terms of the form $cnt_1(d)$ where d is stricly lesser than the index of the current phase. Such terms are redundant and do not affect the encoding. Similar terms exist to encode cnt_2.

An increment transition $(q, cnt_1 + +, q_2) \in \delta$ labeled with a is simulated by a rule labeled with a of the form

$$[q_1 \;,\; phase(x)] \;\rightsquigarrow\; [q_2 \;,\; phase(x) \;,\; cnt_1(x)] \;\;:\;\; true$$

We increase the value counter cnt_1 by adding one more term whose predicate symbol is cnt_1 and whose argument is equal to the index of the current phase.

A decrement transition $(q, cnt_1 - -, q_2) \in \delta$ labeled with a is simulated by a rule labeled with a of the form

$$[q_1 \;,\; phase(x) \;,\; cnt_1(x)] \;\rightsquigarrow\; [q_2 \;,\; phase(x)] \;\;:\;\; true$$

We decrease the value counter cnt_1 by removing one of the corresponding terms from the configuration. Observe that terms whose arguments are less than index of the current phase are not used, and hence they do affect the encoding.

A transition $(q, cnt_1 = 0?, q_2) \in \delta$ labeled by a is simulated by the following three rules (the two first are labeled with ϵ and the last one with a):

$$[q_1 \;,\; phase(x) \;,\; phase'(x)] \;\rightsquigarrow\; [q_1' \;,\; phase(y) \;,\; phase'(x)] \;\;:\;\; x < y$$
$$[q_1' \;,\; cnt_2(x) \;,\; phase(y) \;,\; phase'(x)]$$
$$\rightsquigarrow\; [q_1' \;,\; cnt_2(y) \;,\; phase(y) \;,\; phase'(x)] \;\;:\;\; true$$
$$[q_1' \;,\; phase(y) \;,\; phase'(x)] \;\rightsquigarrow\; [q_2 \;,\; phase(y) \;,\; phase'(y)] \;\;:\;\; true$$

We enter the intermediate phase by changing from q_1 to q_1'. We store the current index using *phase'*, and generate a new index which is strictly larger than the current one. This resets counter cnt_1 since all terms in its encoding will now have too small arguments. Finally, we change the arguments of (some of) the terms encoding cnt_2 to the new index. Here, not all such terms may receive new arguments, and hence the value cnt_2 may "unintentionally" be reduced. We use

redundant terms to refer to terms which have either cnt_1 or cnt_2 as predicate symbol, and whose arguments are smaller than the current index,

Mayr shows in [20] undecidability of the *repeated state reachability* problem for LCM: given a lossy counter machine, and two states q_{init} and q_{fin}, check whether there is a computation starting from q_{init} (with both counter values being equal to zero) and visits q_{fin} infinitely often.

We can extend the proof to show Theorem 2 as follows. The key observation here is that redundant terms are not removed during the simulation procedure described above. This means that any reachable configuration which does not contain redundant terms are generated by computations of the perfect (non-lossy) counter machines. We add the nullary predicate q_{fin}'' and the following rules labeled by ϵ to our CMRS:

$$
\begin{array}{rcll}
[q_{fin} \ , \ cnt_1(x)] & \rightsquigarrow & [q_{fin}''] & : \quad \textit{true} \\
[q_{fin}'' \ , \ phase(x) \ , \ cnt_1(x)] & \rightsquigarrow & [q_{fin}'' \ , \ phase(x)] & : \quad \textit{true} \\
[q_{fin}'' \ , \ phase(x) \ , \ cnt_2(x)] & \rightsquigarrow & [q_{fin}'' \ , \ phase(x)] & : \quad \textit{true} \\
[q_{fin}'' \ , \ phase(x)] & \rightsquigarrow & [q_{fin}''] & : \quad \textit{true}
\end{array}
$$

In other words, if we reach a configuration where $\mathcal{M}$ is in q_{fin}, then $\mathcal{S}$ can initiate a procedure where it moves to q_{fin}'' and then erases the ground terms corresponding to the counters. However, we only remove such ground terms if there arguments are equal to the index of the current phase, i.e. redundant ground terms are not removed. This implies that there exists an execution where $\mathcal{S}$ recognizes a word w that reaches $[q_{fin}'']$ iff there exists an execution where the perfect MC recognizes the word w that reaches q_{fin}. $\qquad\square$

5   Data nets

Data nets [19] are an extension of Petri nets in which tokens are colored with data taken from an infinite domain D equipped with a linear and dense ordering $\prec$. Due to lack of space, we present here only the key concepts needed in the rest of the paper (see [19] for formal definitions). A data net consists of a finite set of places P and of a finite set of transitions. A data net marking s is a multiset of tokens that carry data in D. Formally, a marking s is a finite sequence of vectors in $\mathbb{N}^P \setminus \{\mathbf{0}\}$, where $\mathbf{0}$ is the vector that contains only 0's. Each index i in the sequence s corresponds to some $d_i \in D$ such that $i \leq j$ if and only if $d_i \prec d_j$. For each $p \in P$, $s(i)(p)$ is the number of tokens with data d_i in place p.

First of all, a data net transition t has an associated arity α_t (a natural number greater than zero). The arity $\alpha_t = k$ is used to non-deterministically select k distinct data $d_1 \prec \ldots \prec d_k$ from the current configuration s. Some of the selected data may not occur in s (they are fresh). This choice induces a finite and ordered partitioning of the data in s, namely $R(\alpha_t) = (R_0, S_1, R_1, \ldots, S_k, R_k)$, where R_0 contains all data $d : d \prec d_1$ in s, $S_i = \{d_i\}$ for $i : 1, \ldots, k$, R_i contains all $d : d_i \prec d \prec d_{i+1}$ in s for $i : 1, \ldots, k-1$, and R_k contains all $d : d_k \prec d$ in s. Clearly, $R(\alpha_t)$ also induces a natural partitioning of the multiset of tokens in s based on the attached data.

We define $\overline{k} = \{1,\ldots,k\}$ and $[k]_0 = \{0,\ldots,k\}$. Transition t operates on the regions in the partitioning $R(\alpha_t)$ in three steps defined resp. by three matrices $F_t, H_t \in \mathbb{N}^{R(\alpha_t) \times P}$, and $G_t \in \mathbb{N}^{R(\alpha_t) \times P \times R(\alpha_t) \times P}$.

(1) Subtraction: F_t specifies the number of tokens with data $d_1,\ldots,d_k$ that has to be removed from s. By definition, $F_t(R_i, p) = 0$ for $i \in [k]_0$ and $p \in P$. The transition t is enabled if the subtraction is possible on each place P. This step yields an intermediate configuration s_1 defined as follows:

For each $i \in \overline{k}$ and $p \in P$, $s_1(d_i)(p) = s(d_i)(p) - F_t(S_i, p)$

For each $j \in [k]_0$, $d \in R_j$, $p \in P$, $s_1(d)(p) = s(d)(p)$.

(2) Multiplication: G_t specifies how many tokens are transferred from one place to another with possible multiplication of their occurrences and modification of their data ($G_t(\pi, p, \pi', p') \geq 0$ and, by definition, $G_t(R_i, p, R_j, q) = 0$ for any $i \neq j \in [k]_0$ and any $p, q \in P$). This step yields an intermediate configuration s_2 defined as follows:

For each $i \in \overline{k}$ and $p \in P$:

$$s_2(d_i)(p) = \sum_{j \in \overline{k}} \sum_{q \in P} s_1(d_j)(q) \cdot G_t(S_j, q, S_i, p) + \sum_{j \in [k]_0} \sum_{d \in R_j} \sum_{q \in P} s_1(d)(q) \cdot G_t(R_j, q, S_i, p)$$

For each $i \in [k]_0$, $d \in R_i$, and $p \in P$:

$$s_2(d)(p) = \sum_{j \in \overline{k}, q \in P} s_1(d_j)(q) \cdot G_t(S_j, q, R_i, p) + \sum_{q \in P} s_1(d)(q) \cdot G_t(R_i, q, R_i, p)$$

Notice that transfers of tokens from region R_i to region R_j with $i \neq j$ are forbidden.

(3) Addition: Finally, H_t specifies the number of tokens that are added to each place in P. Its application yields the successor configuration s' such that:

For each $i \in \overline{k}$ and $p \in P$, $s'(d_i)(p) = s_2(d_i)(p) + H_t(S_i, p)$

For each $j \in [k]_0$, $d \in R_j$, $p \in P$, $s'(d)(p) = s(d)(p) + H_t(R_j, p)$.

As proved in [19], data nets are well-structured with respect to the well-quasi ordering $\preceq_d$ defined on markings as follows. Let $Data(s)$ be the set of data values that occur in a marking s. Then, $s_1 \preceq_d s_2$ iff there exists an injective function $h : Data(s_1) \mapsto Data(s_2)$ such that (i) h is monotonic and (ii) $s_1(d)(p) \leq s_2(h(d))(p)$ for each $d \in Data(s_1)$ and $p \in P$.

Example 3. Consider a data net with $P = \{p, q\}$ and the transition in Fig. 1. For a generic configuration s, the new configuration s' is such that:

- $s'(d_1)(p) = s(d_1)(p) - 1 + \Sigma_{d \in R_0} 3 * s(d)(p)$ and $s(d_1)(q)' = 1$;
- For each $d \prec d_1$, $s'(d)(p) = s(d)(p) + 3 * s(d)(q) + 2 * s(d_1)(q)$ and $s(d)(q)' = s(d)(q)$;
- For each $d \gg d_1$, $s(d)(p)' = s(d)(p)$ and $s(d)(q)' = s(d)(q)$.

For instance, let $e_1 \prec e_2 \prec e_3 \prec e_4 \in D$ and assume that the transition selects e_3 as index in S_1, then:

$$\begin{pmatrix} & e_1 & & e_2 & & \boxed{e_3} & & e_4 & \\ p & q & p & q & p & q & p & q \\ 3 & 2 & 5 & 1 & 2 & 10 & 2 & 2 \end{pmatrix} \rightarrow \begin{pmatrix} & e_1 & & e_2 & & \boxed{e_3} & & e_4 & \\ p & q & p & q & p & q & p & q \\ 29 & 2 & 28 & 1 & 25 & 1 & 2 & 2 \end{pmatrix}$$

$$F_t = \left(\begin{array}{cc|cc|cc} \multicolumn{2}{c}{R_0} & \multicolumn{2}{c}{S_1} & \multicolumn{2}{c}{R_1} \\ p & q & p & q & p & q \\ 0 & 0 & 1 & 0 & 0 & 0 \end{array} \right)$$

$$H_t = \left(\begin{array}{cc|cc|cc} \multicolumn{2}{c}{R_0} & \multicolumn{2}{c}{S_1} & \multicolumn{2}{c}{R_1} \\ p & q & p & q & p & q \\ 0 & 0 & 0 & 1 & 0 & 0 \end{array} \right)$$

$$G_t = \left(\begin{array}{cc|cc|cc} & \multicolumn{2}{c}{R_0} & \multicolumn{2}{c}{S_1} & \multicolumn{2}{c}{R_1} \\ & p & q & p & q & p & q \\ \hline R_0 \;\; p & 1 & 0 & 3 & 0 & 0 & 0 \\ q & 3 & 1 & 0 & 0 & 0 & 0 \\ \hline S_1 \;\; p & 0 & 0 & 1 & 0 & 0 & 0 \\ q & 2 & 0 & 0 & 0 & 0 & 0 \\ \hline R_1 \;\; p & 0 & 0 & 0 & 0 & 1 & 0 \\ q & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right)$$

Fig. 1. A data net transition with arity $\alpha_t = 1$.

5.1 Datanets vs CMRS

In [19] the authors mention that it is possible to define an encoding of CMRS in the fragment of data net without whole place operations (*Petri data nets*) that preserves coverability. From this observation, it follows that $L_c(CMRS) \subseteq L_c(Data\ nets)$.

In this section we tighten this relation and show that for each Data nets $\mathcal{D}$ we can effectively build a CMRS $\mathcal{S}$ such that $L_c(\mathcal{S}) = L_c(\mathcal{D})$. In the following, given a multi-set M with symbols in P and a value or variable x, we use M^x to denote the multi set of P-terms such that $M^x(p(x)) = M(p)$ (=number of occurrences of p in M) for each $p \in P$, and $M^x(p(y)) = 0$ for any $y \neq x$ and $p \in P$.

Configurations Assume an initial data net marking s_0 with data $d_1 \prec \ldots \prec d_n$. We build a CMRS representation of s_0 by non-deterministically selecting n natural numbers $v_1 < \ldots < v_n$ strictly included in some interval $[f, l]$. P-terms with parameter v_i represent tokens with data d_i in place p. Formally, we generate the representation of s_0 by adding to $\mathcal{S}$ a rule that rewrites an initial zerary term *init* as follows:[4]

$$[init] \rightsquigarrow [first(f), last(l)] + \sum_{i:1,\ldots,n} M_i^{x_i} \;\; : \;\; f < x_1 < \ldots < x_n < l \qquad (init)$$

where M_i is the multiset $s_0(d_i)$ for each $i \in \overline{n}$. The non-determinism in the choice of $f, l, x_1, \ldots, x_n$ make the CMRS representation of s_0 independent from specific parameters assumed by terms.

Transitions are encoded by CMRS rules that operate on the values in $[f, l]$ used in the representation of a marking. Most of the CMRS rule are based on left-to-right traversals of P-terms with parameters in $[f, l]$.

Subtraction Consider a transition t with $\alpha_t = k$. We first define a (silent) CMRS-rule that implements the subtraction step of t:

$$[first(f), last(l)] + F_t(S_1)^{x_1} + \ldots + F_t(S_k)^{x_k} \;\; \rightsquigarrow \qquad\qquad (subtract)$$
$$[\imath_0(f), \imath_1(x_1), \ldots, \imath_k(x_k), \imath_{k+1}(l), new_t] : \;\; f < x_1 < \ldots < x_k < l$$

---

[4] We recall that $[t_1, \ldots, t_n]$ denotes a multisets of terms. Furthermore, $\sum_{i:1,\ldots,k} M_i = M_1 + \ldots + M_k$, where $+$ is multiset union.

In the *subtract* rule we non-deterministically associate a value w_i to region S_i. The selection is performed by removing (from the current configuration) the multiset $F_t(S_i)^{x_i}$ that contains $F_t(S_i, p)$ occurrences of $p(x_i)$ for each $p \in P$. The association between value x_i and region S_i is maintained by storing x_i in a ι_i-term (introduced in the right-hand side of the rule). If $F_t(S_i, p) = 0$ for any $p \in P$, then v_i may be associated to a data d_i not occurring in the current marking (i.e. selection of fresh data is a special case). Furthermore, by removing both the *first*- and the *last*-term, we disable the firing of rules that encode other data net transitions. Fig. 9 in appendix shows an example of application of the *subtract* rule.

In the rest of the section we refer to The values $x_1, \ldots, x_k$ stored in ι_1-,$\ldots$,ι_k-terms play the role of pointers to the regions $S_1, \ldots, S_k$ We refer to them as to the set of α_t-*indexes*. The parameters of terms in $[f, l]$ associated to the other regions $R_0, \ldots, R_k$ are called *region-indexes*.

Multiplication To simulate the multiplication step we proceed as follows. We first make a copy of the multiset of P-terms with parameters $v_1, \ldots, v_n$ in $[f, l]$ by copying each p-term with parameter v_i in a $\overline{p}$-term with parameter w_i such that $f' < w_1 < \ldots < w_n < l'$ and $[f', l']$ is an interval to the right of $[f, l]$, i.e., $l < f'$. The new_t-term in the *subtract* rule is used to enable the set of (silent) CMRS rules in Fig. 2 in appendix that create the copy-configuration. During the copy we add a $\checkmark$-term for any *visited* region index. These terms are used to remember region indexes whose corresponding $\overline{P}$-terms are all removed in the multiplication step (e.g. when all tokens with data $d \in R_i$ are removed).

For instance, $[p(v_1), p(v_2), p(v_2), q(v_3)]$ with $f < v_1 < v_2 < v_3 < l$ is copied as $[\overline{p}(w_1), \checkmark(w_1), \overline{p}(w_2), \overline{p}(w_2), \checkmark(w_2), \overline{q}(w_3)\checkmark(w_3)]$ for some w_1, w_2, w_3 such that $f < l < f' < w_1 < w_2 < w_3 < l'$. The CMRS rules of Fig. 2 use a special term as a pointer scan the indexes in $[f, l]$ from left to right and create new $\overline{P}$-term with parameters in the interval $[f', l']$. The pointer is non-deterministically moved to the right. Thus during the traversal we may forget to copy some token. This is the first type of loss we find in our encoding. Notice that lost tokens have parameters strictly smaller that f'.

The simulation of the multiplication step operates on the copy-configuration only (that with $\overline{P}$-terms). The (silent) CMRS rules that implement this step are shown in Fig. 3 in appendix. The intuition behind their definition is as follows.

We first consider all α_t-indexes of $\overline{P}$-terms from left to right. For each α_t-index v_i, we proceed as follows. We first select and remove a term $\overline{p}(v_i)$ (encoding a given token). We compute then the effect of the whole-place operation on the entire set of α_t-indexes (including v_i itself). More specifically, for an α_t-index v_j we add $G_t(S_i, p, S_j, q)$ occurrences of the term $q(v_j)$ to the current CMRS configuration. The use of P- and $\overline{P}$-terms with parameters in the same interval allows us to keep track of tokens still to transfer ($\overline{P}$-terms) and tokens already transferred (P-terms). We then consider all remaining indexes by means of a left-to-right traversal of region-indexes in the current configuration. During the traversal, we add new P-terms with region-indexes as parameters as specified by G_t. During this step, we may forget to transfer some $\overline{P}$-term. This is the second

type of loss we find in the encoding. After this step we either consider the next token with α_t-index v_i or we move to the next α_t-index. Fig. 10(a) in appendix illustrates the simulation of this kind of transfers (i.e. from S_i to S_j/R_j).

After the termination of the whole-place operations for terms with α_t-indexes, we have to simulate the transfer of $\overline{P}$-terms with region-indexes. For each such an index, we transfer tokens within the same region-index or to an α_t-index. To simulate these operations we scan region-indexes from left-to-right to apply the matrix G_t. Furthermore, we mark visited region-indexes using $\checkmark$-terms. The $\checkmark$-terms are used in the simulation of the addition step. The (silent) CMRS rules that implement this step (enabled by the by term trR_t) are shown in Fig. 4 in appendix. Fig. 10(b) in appendix illustrates the simulation of this type of whole-place operation.

As a last step we add tokens to α_t-indexes and visited region indexes as specified by H_t. For α_t-indexes, we need a single rule that applies the matrix H_t. For region-indexes, we traverse from left-to-right the current configuration and apply H_t to each marked (with a $\checkmark$-term) region-index w. As mentioned before, the $\checkmark$-term allows us to apply H_t to regions emptied by the multiplication step. The rules for this step (associated to terms add_t and $addR_t$) are shown in Fig. 4. All the rules are silent except the last one whose label is the same as that of t. Fig. 10 in appendix shows an example of their application.

During the traversal, we may ignore some (marked) region-index. This is the last type of loss in our encoding. The new configuration is the final result of the simulation of the transition. Due to the possible losses in the different simulation steps, we may get a representation of a data net configuration smaller than the real successor configuration.

To formalize the relation between a data net $\mathcal{D}$ and its CMRS encoding $\mathcal{E}(\mathcal{D})$, for a configuration s with data $d_1 \prec \ldots \prec d_k$ we use $s^{\boldsymbol{v}}$ to denote the CMRS representation with indexes $\boldsymbol{v} = (v_1, \ldots, v_k)$.

Proposition 4. *For configurations s_0, s_1, s, s', the following properties hold:*

  i) If $s_0 \overset{w}{\Longrightarrow} s_1$ in $\mathcal{D}$, then there exists $\boldsymbol{v}$ such that $[init] \overset{w}{\Longrightarrow} s_1{}^{\boldsymbol{v}}$ in $\mathcal{E}(\mathcal{D})$;
  ii) Furthermore, if $[init] \overset{w}{\Longrightarrow} c$ in $\mathcal{E}(\mathcal{D})$ and $s^{\boldsymbol{v}} \preceq_c c$ for some $\boldsymbol{v}$, then there exists s_1 such that $s_0 \overset{w}{\Longrightarrow} s_1$ in $\mathcal{D}$ with $s \preceq_d s_1$.

Finally, suppose that the accepting data net marking is a sequence $M_1 \ldots M_k$ of k vectors (multi-sets) over $\mathbb{N}^P$. Then, we add a silent CMRS rule

$$[first(f), last(l)] + \sum_{i \in \{1,\ldots,k\}} M_i^{x_i} \rightsquigarrow [acc] \;:\; f < x_1 < x_2 < \ldots < x_k < l, x = 0$$

where acc is a fresh (with arity zero) predicate. By adding this rule, the accepting CMRS configuration can be defined as the singleton $[acc]$. From Lemma 1 and Prop. 4, we have the following result.

Theorem 3. $L_c(data\ nets) = L_c(CMRS)$

5.2 Extensions of data nets

An encoding into CMRS can also be applied to extend the decidability result for coverability to different extensions of data nets. As a first extension, we consider freshness of data values. Let us consider a data net transition t with $\alpha_t = k$. In the semantics of data nets, some of the k data values selected by t *may* be fresh, i.e., they don't have to occur in the current configuration. This definition can be extended by introducing the constraint that some of the selected data *must* be fresh for the transition to be fired. In the Petri net setting a similar operator has been considered in the ν-nets of [17] to create new, unused identifier.

For simplicity, we consider here the extension of data net transitions in which we require that only one of the α_t data value must be fresh. This kind of transition can be modelled by extending the CMRS encoding of the subtraction step as follows. Before selecting the α_t data, we make a copy (in a new interval) of the current configuration. In the new configuration we non-deterministically mark using predicate *new* a value x distinct from the values used to represent tokens. After this preliminary step, we apply the subtraction phase by requiring that the value x is one of the selected ones (i.e. we need α_t rules for this last step). Formally, we use the rules in Fig. 6 in appendix. This extension provides a direct way to model freshness without need of ordering identifiers and of maintaining in a special place the last used one (the natural way of modelling ν-nets in ordinary data nets).

Another possible extension concerns the relaxation of some of the restrictions in the definition of data nets in [19]. For instance, assume we allow transfers between regions R_i and R_j with $i \neq j$. The semantics of a transfer with $G_t(R_i, p, R_j, p') = m > 0$ is the following. For each $d \in R_i$, place $p \in P$, and each token with d in p, we add m tokens with data d' to p' for each $d' \in R_j$. Furthermore, we can also consider a new type of whole-place operation within the same region R_i in which we can multiply the tokens with data d for each data with value $d' \in R_i$ with $d \neq d'$. More formally, assume we add a new matrix $M_t(R_i, p, R_i, p')$ to specify, for each token in p with data $d \in R_i$, how many tokens to add to place p' with data d' for each $d' \in R_i$, $d' \neq d$. These extensions of data net transitions are still monotonicity w.r.t. $\preceq_d$. Furthermore, they can be weakly simulated in CMRS as shown in Fig.7 in appendix. Thus, we have that coverability remains decidable and, from Lemma 1, we have that $L_c(extended\ data\ nets) = L_c(CMRS)$.

6 Lossy FIFO Channel Systems

In this section we study the relationship between the fragment Γ_0 of CMRS defined in Section 3.2 and *lossy (FIFO) channel systems* (LCS) [4].

A *Lossy FIFO Channel System* (LCS) consists of an asynchronous parallel composition of finite-state machines that communicate through sending and receiving messages via a finite set of unbounded lossy FIFO channels (in the sense that they can non-deterministically lose messages). Formally, an LCS $\mathcal{F}$

For $k = \alpha_t$, $i \in \{0, \ldots, k\}$, and any $p \in P$:

Copy of indexes in α_t

$[\imath_0(x_0), \ldots, \imath_{k+1}(x_{k+1}), new_t] \rightsquigarrow$
$\quad [\imath_0(x_0), \ldots, \imath_{k+1}(x_{k+1}), \jmath_0(x'_0), \ldots, \jmath_{k+1}(x'_{k+1}), \uparrow(x_0), \Uparrow(x'_0)] :$
$\qquad x_{k+1} < x'_0 < \ldots < x'_{k+1}$

Copy p to $\overline{p}$ for $\alpha_t - indexes$

$[\uparrow(x), \imath_i(x), \Uparrow(y), \jmath_i(y), p(x)] \rightsquigarrow [\uparrow(x), \imath_i(x), \Uparrow(y), \jmath_i(y), \overline{p}(y)] : true$

Copy p to $\overline{p}$ for $region - indexes$

$[\imath_i(x), \uparrow(u), p(u), \imath_{i+1}(x'), \jmath_i(y), \Uparrow(v), \jmath_{i+1}(y')] \rightsquigarrow$
$\quad [\imath_i(x), \uparrow(u), \imath_{i+1}(x'), \jmath_i(y), \overline{p}(v), \Uparrow(v), \jmath_{i+1}(y')] : x < u < x', y < v < y'$

Move pointers to the right

$[\uparrow(u), p(u'), \imath_{k+1}(x), \Uparrow(v), \jmath_{k+1}(y)] \rightsquigarrow$
$\quad [\uparrow(u'), \imath_{k+1}(x), \Uparrow(v'), p(v'), \checkmark(v'), \jmath_{k+1}(y)] : u < u' < x, v < v' < y$

Terminate copy, replace current conf with new one

$[\imath_0(f), \imath_1(x_1), \ldots, \imath_k(x_k), \imath_{k+1}(l), \jmath_0(f'), \jmath_1(x'_1), \ldots, \jmath_k(x'_k), \jmath_{k+1}(l'), \uparrow(u), \Uparrow(v)] \rightsquigarrow$
$\quad [\imath_0(f'), \imath_1(x'_1), \ldots, \imath_k(x'_k), \imath_{k+1}(l'), tr_t] : true$

Fig. 2. Silent CMRS rules for new_t: generation of a new configuration with $\overline{p}$-terms.

is a tuple (Q, C, N, δ) where Q is a finite set of control states (the Cartesian product of those of each finite-state machine), C is a finite set of channels, M is a finite set of messages, δ is a finite set of transitions, each of which is of the form (q_1, Op, q_2) where $q_1, q_2 \in Q$, and Op is a mapping from channels to channel operations. For any $c \in C$ and $a \in M$, an operation $Op(c)$ is either a *send* operation $!a$, a *receive* operation $?a$, the *empty test* $\epsilon?$, or the *null* operation *nop*. A configuration γ is a pair (q, w) where $q \in Q$, and w is a mapping from C to M^* giving the content of each channel. The initial configuration γ_{init} of $\mathcal{F}$ is the pair (q_0, ε) where $q_0 \in Q$, and ε denotes the mapping that assigns the empty sequence ϵ to each channel. The (strong) transition relation (that defines the semantics of machines with *perfect* FIFO channels) is defined as follows: $(q_1, w_1) \xrightarrow{\sigma} (q_2, w_2)$ if and only if $\sigma = (q_1, Op, q_2) \in \delta$ such that if $Op(c) = !a$, then $w_2(c) = w_1(c) \cdot a$; if $Op(c) = ?a$, then $w_1(c) = a \cdot w_2(c)$; if $Op(c) = \epsilon?$ then $w_1(c) = \epsilon$ and $w_2(c) = \epsilon$; if $Op(c) = nop$, then $w_2(c) = w_1(c)$. Now let $\preceq_l$ be the quasi ordering on LCS configurations such that $(q_1, w_1) \preceq_l (q_2, w_2)$ iff $q_1 = q_2$ and $\forall c \in C : w_1(c) \preceq_w w_2(c)$ where $\preceq_w$ indicates the subword relation. By Higman's theorem, we know that $\preceq_l$ is a well-quasi ordering. We introduce then the weak transition relation $\xRightarrow{\sigma}$ that defines the semantics of LCS: we have $\gamma_1 \xRightarrow{\sigma} \gamma_2$ iff there exists γ'_1 and γ'_2 s.t. $\gamma'_1 \preceq_l \gamma_1$, $\gamma'_1 \xrightarrow{\sigma} \gamma'_2$, and $\gamma_2 \preceq_l \gamma'_2$. Thus, $\gamma_1 \xRightarrow{\sigma} \gamma_2$ means that γ_2 is reachable from γ_1 by first losing messages from the channels and reaching γ'_1, then performing a transition, and, thereafter losing again messages from channels. As shown in [4], an LCS is well-structured with respect to $\preceq_l$. Furthermore, as shown in [5], in presence of transitions labeled

For $k = \alpha_t$, $i \in \{1, \ldots, k\}, j \in \{0, \ldots, k\}$, and $p \in P$:

Start from first index

$[tr_t] \rightsquigarrow [tr_{t,1}]$: *true*

Select a token from an index in α_t, apply G_t to other indexes :

$[\imath_0(x_0), \imath_1(x_1), \ldots, \imath_i(x_i), \ldots, \imath_k(x_k), \imath_{k+1}(x_{k+1}), \overline{p}(x_i), tr_{t,i}] \rightsquigarrow$
$\quad [\imath_0(x_0), \imath_1(x_1), \ldots, \imath_k(x_k), \imath_{k+1}(x_{k+1}), apply_{t,i,p}(x)] + \Sigma_{j=1}^{k} G_t(S_i, p, S_j)^{x_j} : x_0 < x < x_{k+1}$

Apply G_t to indexes inside regions, move to the right

$[\imath_j(v), apply_{t,i,p}(u), \overline{p}(u), \imath_{j+1}(v')] \rightsquigarrow$
$\quad [\imath_j(v), apply_{t,i,p}(u'), \overline{p}(u), \imath_{j+1}(v')] + G_t(S_i, p, R_j)^u : v < u < v', u < u'$

Terminate visit continue with next token

$[apply_{t,i,p}(u)] \rightsquigarrow [tr_{t,i}]$: *true*

Move to next index

$[tr_{t,j}] \rightsquigarrow [tr_{t,j+1}]$: *true*

Terminate transfer of tokens for indexes in α_t, start transfer of tokens of regions

$[\imath_0(f), tr_{t,k}] \rightsquigarrow [trR_t(f)]$: *true*

Fig. 3. Silent CMRS rules for simulation of transfer: $G_t(S_i, p, \pi)^x$ is the multiset that, for each $q \in P$, contains $G_t(S_i, p, \pi, q)$ occurrences of the term $q(x)$.

with ϵ, we can restrict our attention to systems with only one channel. As a last remark, notice that for any model with lossy semantics like LCS, e.g. lossy vector addition systems [20], the class of c-languages coincide with the class of r-languages, i.e., $L_r(LCS) = L_c(LCS)$.

Our first result is that Γ_0 and LCS define the same class of c-languages.

Theorem 4. $L_c(\Gamma_0) = L_c(LCS)$.

To prove the previous result, we give separate proofs of the two inclusions.

Proposition 5. $L_c(LCS) \subseteq L_c(\Gamma_0)$

Proof. Assume an LCS $\mathcal{F}$. We build a Γ_0 $\mathcal{S}$ that simulates $\mathcal{F}$. The set of predicate symbols in $\mathcal{S}$ consists of the following: For each $q \in Q$, there is a nullary predicate symbol q in $\mathcal{S}$. For each channel c_i we use the function symbols $head^i$ and $tail^i$ as pointers to the head and tail of the queue c_i. For each channel c_i and each message $a \in M$ we have the predicate symbol a^i in $\mathcal{S}$.

If $C = \{c_1, \ldots, c_n\}$, then the initial configuration (s_0, ϵ) is represented as

$$M_0 = \left[s_0 \, , \, head^1(v_0) \, , \, tail^1(v_0 + 1) \, , \, \ldots \, , \, head^n(v_0) \, , \, tail^n(v_0 + 1) \right]$$

for some $v_0 \in \mathbb{N}$.

In order to represent the queue c_i containing the word $a_1 a_2 \ldots a_n$, we will use the multiset

$$\left[head^i(v_0) \, , \, a_1^i(v_1) \, , \, \ldots \, , \, a_n^i(v_n) \, , \, tail^i(v_{n+1}) \right]$$

For k $= \alpha_t$, i $\in \{0, \ldots, k\}$, and any p $\in$ P :

Remove token and apply G_t to indexes inside regions

$$[\iota_0(x_0), \iota_1(x_1), \ldots, \iota_i(x_i), \iota_{i+1}(x_{i+1}), \ldots, \iota_{k+1}(x_{k+1}), \overline{p}(u), trR_t(u)] \rightsquigarrow$$
$$[\iota_0(x_0), \iota_1(x_1), \ldots, \iota_i(x_i), \iota_{i+1}(x_{i+1}), \ldots, \iota_{k+1}(x_{k+1}), trR_t(u)]$$
$$+ G_t(R_i, p, R_i)^u + \Sigma_{j=1}^k G_t(R_i, p, S_j)^{x_j} : x_i < u < x_{i+1}$$

Move pointers to the right

$$[trR_t(u), \iota_{k+1}(l)] \rightsquigarrow [trR_t(u'), \iota_{k+1}(l),] : u < u' < l$$

Terminate visit, move to addition step

$$[trR_t(u)] \rightsquigarrow [add_t] : true$$

Fig. 4. Silent CMRS rules for trR: transfer inside a region-index and from a region-index to α_t-indexes: $G_t(R_i, p, \pi)^x$ is the multiset that, for each $q \in P$, contains $G_t(R_i, p, \pi, q)$ occurrences of the term $q(x)$.

For $k = \alpha_t$, $i \in \{0, \ldots, k\}$, and any $p \in P$:

Apply H_t to indexes in α_t

$$[\iota_0(x_0), \iota_1(x_1), \ldots, \iota_k(x_k), add_t] \rightsquigarrow$$
$$[\iota_0(x_0), \iota_1(x_1), \ldots, \iota_k(x_k), addR_t(x_0)] + \Sigma_{j=1}^k H_t(S_j)^{x_j} : true$$

Apply H_t to an index inside a region and advance pointer

$$[\iota_i(v), \iota_{i+1}(v'), addR_t(u), \checkmark(u)] \rightsquigarrow$$
$$[\iota_i(v), \iota_{i+1}(v'), addR_t(u')] + H_t(R_i)^u : v < u < v', u < u'$$

Terminate simulation of transition t

$$[\iota_0(x_0), \iota_1(x_1), \ldots, \iota_{k+1}(x_{k+1}), addR_t(u)] \overset{\lambda(t)}{\rightsquigarrow} [first(x_0), last(x_{k+1})] : true$$

Fig. 5. CMRS rules for add_t and $addR_t$ (all silent except the last one): $H_t(\pi)^x$ is the multiset that, for each $p \in P$, contains $H_t(\pi, p)$ occurrences of the term $p(x)$ for any $\pi \in R(\alpha_t)$.

for some positive integers $v_0 < v_1 < \ldots < v_{n+1}$.

Since an LCS transition (q_1, Op, q_2) operates simultaneously on all the queues, the corresponding CMRS rule (with the same label) has the following form:

$$[\, q_1 \,] + B_1 + \ldots + B_n \rightsquigarrow [\, q_2 \,] + B'_1 + \ldots + B'_n : C_1 \cup \ldots \cup C_n$$

where B_i, B'_i and C_i define the encoding of $Op(c_i)$ for $i : 1, \ldots, n$. The encoding of the operation is defined by atomic formulas defined on a distinct variables $x_1, y_1, \ldots, x_n, y_n$ as follows.
For $Op(c_i) = !a$,

$$B_i = [tail^i(x_i)] \quad B'_i = [a^i(x_i), tail^i(y_i)] \quad C_i = \{x_i < y_i\}$$

For $Op(c_i) = ?a$,

$$B_i = [head^i(x_i), a^i(y_i)] \quad B'_i = [head^i(y_i)] \quad C_i = \{x_i < y_i\}$$

Start copy :
$[first(f), last(l)] \rightsquigarrow [first_1(f), scan_t^0(f), last_1(l), copy_t^0(l')] : \ l < l'$

Copy a term :
For $i \in \{0, 1\}$
$[scan_t^i(x), p(x), copy_t^i(y)] \rightsquigarrow [scan_t^i(x), copy_t^i(y), p(y)] : \ true$

Move to the right :
For $i \in \{0, 1\}$
$[scan_t^i(x), last_1(l), copy_t^i(y)] \rightsquigarrow$
$\quad [scan_t^i(x'), last_1(l), copy_t^i(y')] : \ x < x' < l, y < y'$

Move to the right and reserve a fresh value in the copy tape :
$[scan_t^0(x), last_1(l), copy_t^0(y)] \rightsquigarrow$
$\quad [scan_t^1(x'), last_1(l), fresh_t(y'), copy_t^1(y'')] : \ x < x' < l, y < y' < y''$

Terminate copy (copy tape becomes current configuration) :
$[first_1(f), scan_t^1(x), last_1(l), copy_t^1(y)] \rightsquigarrow [first_1(l), last_1(y)] : \ true$

Subtraction :
For $j \in \{1, \ldots, k\}$
$[first_1(f), fresh_t(x), last_1(l)] + \Sigma_{i=1}^{k} F_t(S_i)^{x_i} \rightsquigarrow$
$\quad [i_0(f), i_1(x_1), .., i_j(x), .., i_k(x_k), i_{k+1}(l), new_t] :$
$\qquad f < x_1 < .. < x_j < x < x_{j+1} < x_k < l$

Fig. 6. Silent CMRS rules for the simulation of subtraction with selection of one fresh value.

For $Op(c_i) = empty?$,

$$B_i = \left[head^i(x_i) \mid tail^i(y_i)\right] \quad B_i' = \left[head^i(x_i') \mid tail(y_i')\right] \quad C_i = \{y_i < x_i' < y_i'\}$$

For $Op(c_i) = nop$,
$$B_i = B_i' = [\,] \quad C_i = true$$

Now let $\gamma_{acc} = (s_{acc}, \epsilon)$ be an accepting LCS configuration (i.e. all configurations with control state s_{acc} allow to accept words).

The accepting CMRS configuration is $[s_{acc}]$. Let us consider a LCS with one channel. Note that n channels can be encoded into one channel. Hence, considering a unique channel is not restrictive.

The following properties then hold. Given an LCS configuration $\gamma = (s, w)$, let $\gamma^\bullet$ be the corresponding CMRS encoding. Moreover, given $\gamma^\bullet$ containing $head^1(c)$, let $\mathsf{G}(\gamma^\bullet)$ be the set of CMRS configurations built from $\gamma^\bullet$ by adding some ground terms $a^1(c')$ where $a \in M$ and $c' < c$, i.e. by adding useless ground terms corresponding to lost messages.

It is easy to check that (1) if $\gamma_1^\bullet \xrightarrow{\rho} \eta$ with $\eta \in \mathsf{G}(\gamma_2^\bullet)$ in $\mathcal{S}$, then $\gamma_1 \xRightarrow{\sigma_\rho} \gamma_2$ in $\mathcal{F}$ where ρ is the CMRS rule corresponding to the LCS transition σ_ρ. Indeed, notice that in the CMRS implementation of the *dequeue* operation we move the *head* pointer to an arbitrary position within the queue and thus we perform a lossy step followed by a dequeue step. Similarly, the emptiness test is simulated by means of a lossy step in which all elements are removed from the queue (with

For $k = \alpha_t$, $i \in \{0, \ldots, k\}$, and any $p \in P$:

Remove token, apply G_t to indexes inside region and to $\alpha_t - indexes$

$[\iota_0(x_0), \iota_1(x_1), \ldots, \iota_i(x_i), \iota_{i+1}(x_{i+1}), \ldots, \iota_{k+1}(x_{k+1}), \overline{p}(u), trR_t(u)] \rightsquigarrow$
$[\iota_0(x_0), \iota_1(x_1), \ldots, \iota_i(x_i), \iota_{i+1}(x_{i+1}), \ldots, \iota_{k+1}(x_{k+1}), applyR_{i,p,t}(u)]$
$\quad + G_t(R_i, p, R_i)^u + \Sigma_{j=1}^k G_t(R_i, p, S_j)^{x_j} : \ x_i < u < x_{i+1}$

Apply G_t to other region $-$ indexes, move to the right For $i \neq j$, $\pi \in \{\overline{p}, \checkmark\}$:

$[\iota_j(v), applyR_{t,i,p}(u), \pi(u), \iota_{j+1}(v')] \rightsquigarrow$
$[\iota_j(v), applyR_{t,i,p}(u'), \pi(u), \iota_{j+1}(v')] + G_t(R_i, p, R_j)^u : v < u < v', u < u'$

Apply M_t to other index in the same region, move to the right

$[\iota_i(v), applyR_{t,i,p}(u), \pi(z), \iota_{i+1}(v')] \rightsquigarrow$
$[\iota_i(v), applyR_{t,i,p}(u'), \pi(z), \iota_{i+1}(v')] + M_t(R_i, p, R_j)^z : z \neq u, v < u < v', u < u'$

Terminate visit continue with next token

$[apply_{t,i,p}(u)] \rightsquigarrow [trR_t] : \ true$

Move trR_t pointer to the right

$[trR_t(u), \iota_{k+1}(l)] \ \rightsquigarrow \ [trR_t(u'), \iota_{k+1}(l),] : \ u < u' < l$

Terminate visit, move to addition step

$[trR_t(u)] \rightsquigarrow [add_t] : \ true$

Fig. 7. Silent CMRS rules for the simulation of transfer between distinct regions, or distinct indexes inside the same region.

the weak reduction of LCS the emptiness test is always executable and it has the effect of emptying the queue). Finally, the enqueue operation is simulated in an exact way. We can also easily see that for all $\eta \in G(\gamma_1^\bullet)$: (2) if $\eta \xrightarrow{\rho} \eta'$, then $\gamma_1^\bullet \xrightarrow{\rho} \gamma_2^\bullet$ with $\eta' \in G(\gamma_2^\bullet)$. Hence, if $\gamma_0^\bullet \xrightarrow{\rho_1} \eta_1 \ldots \xrightarrow{\rho_n} \eta_n$ with $[s_{acc}] \preceq_c \eta_n$, then we deduce from (2) that for all $i : 1 \leq i \leq n$, there exists $\gamma_i^\bullet$ such that $\eta_i \in G(\gamma_i^\bullet)$. Moreover, for all $i : 0 \leq i < n$, $\gamma_i^\bullet \xrightarrow{\rho_i} \eta_{i+1}'$ with $\eta_{i+1}' \in G(\gamma_{i+1}^\bullet)$. Since $[s_{acc}] \preceq_c \eta_n$ and $\eta_n \in G(\gamma_n^\bullet)$ we also have $[s_{acc}] \preceq_c \gamma_n^\bullet$. Following (1), we deduce $\gamma_0 \xRightarrow{\sigma_{\rho_1}} \gamma_1 \ldots \xRightarrow{\sigma_{\rho_n}} \gamma_n$ with $\gamma_{acc} \preceq_l \gamma_n$.

Vice versa, suppose that $\gamma_1 \xRightarrow{\sigma} \gamma_2$ in $\mathcal{F}$. Then, we have that there exists η such that $\gamma_1^\bullet \xrightarrow{\rho_\sigma} \eta^\bullet$ and $\gamma_2^\bullet \preceq_c \eta^\bullet$ where ρ_σ is the CMRS rule corresponding to σ. This is immediately verified for the enqueue operation and for the empty test (their simulation is exact, and thus returns a more precise representation of the queues). The same holds for the dequeue operation since we cannot forget elements to the right of the new position of the header.

Now let $\gamma_1 \xRightarrow{\rho_0} \gamma_2 \ldots \xRightarrow{\rho_n} \gamma_n$ with $\gamma_{acc} \preceq_l \gamma_n$. Then, we know that there exist $\eta_2', \ldots, \eta_n'$ such that $\gamma_i^\bullet \xrightarrow{\sigma_{\rho_{i-1}}} \eta_{i+1}^\bullet$ and $\gamma_{i+1}^\bullet \preceq_c \eta_{i+1}^\bullet$ for $i : 1, \ldots, n-1$. By the monotonicity of CMRS, we have that $\gamma_1^\bullet \xrightarrow{\sigma_{\rho_0} \ldots \sigma_{\rho_n}} \eta_n^\bullet$, and $[s_{acc}] \preceq_c \gamma_{acc}^\bullet \preceq_c \eta_n^\bullet$.
$\square$

Autonomous transition t_1 :

$$F_{t_1}{}^x + [\checkmark(x), last(l)] \overset{\lambda(t_1)}{\rightsquigarrow} H_{t_1}{}^x + [\checkmark(x), last(l)] : x < l$$

Synchronized transitions t_2 *and* t_3 :

$$F_{t_2}{}^x + F_{t_3}{}^y + [\checkmark(x), \checkmark(y), last(l)] \overset{\lambda(t_2, t_3)}{\rightsquigarrow}$$
$$H_{t_2}{}^x + H_{t_3}{}^y + [\checkmark(x), \checkmark(y), last(l)] : \ x \neq y < l$$

Replication rule t_4 *labelled with the marking* M :

$$F_{t_4}{}^x + [\checkmark(x), last(y)] \overset{\lambda(t_4)}{\rightsquigarrow} H_{t_4}{}^x + M^y + [\checkmark(x), \checkmark(y), last(z)] : z > y$$

Fig. 8. CMRS rules that simulate autonomous, synchronized, and replication transitions.

Proposition 6. $L_c(\Gamma_0) \subseteq L_c(LCS)$

Proof. Consider a Γ_0 $\mathcal{S}$ over the finite set of predicate symbols $\mathbb{P}$, an initial configuration γ_{init} and an accepting configuration γ_{acc}. W.l.o.g. let us assume that for each $p(v) \in \gamma_{init} + \gamma_{acc} : 0 \leq v \leq cmax$ (to satisfy this condition we can add a rule that is never fireable and in which there is a constant greater than all values in $\gamma_{init} + \gamma_{acc}$).

The proof follows three steps: first, we show how to encode a configuration as words (i.e. contents of LCS queues). Second, we show how a rule $L \rightsquigarrow R : \psi$ can be applied to the word representation of configurations, and finally we show how to simulate such an application using an LCS.

Γ_0 *configurations as words.* Γ_0 configurations consisting of terms with strictly increasing parameters can be naturally viewed as words defined over the corresponding predicate symbols. The execution of a Γ_0 on such a configuration, however, may lead to a new configuration with two terms with the same value.

As an example, consider the rule ρ defined as $[p(x), q(y)] \rightsquigarrow [p(x), r(z), q(y)] : x < z < y$ and a configuration $\gamma = [p(0), t(3), q(6)]$. We notice here that the application of ρ to γ may lead to different results depending on the valuation of z. One of the possible successors is $\gamma' = [p(0), r(3), t(3), q(6)]$. γ' obtained by applying the valuation $x \mapsto 0, z \mapsto 3, y \mapsto 6$. The question now is if we gain something in assigning to z the same value of a parameter in another term. The answer is no. Indeed, since in Γ_0 we cannot test for = in a rule, the effect of mapping z to 3 only restricts the set of rules that can be fired at γ', i.e., this choice can lead to dead ends.

This intuition is made formal in Prop. 3. This lemma tells us that all strings in $L_c(\mathcal{S})$ can be recognized by an execution that passes through configurations where all the terms with a value greater than $cmax$ are totally ordered on the values of their parameters w.r.t. < (i.e. they can be viewed as words). Notice that this reasoning can be applied only to terms with values greater than $cmax$. Indeed, for this kind of terms Prop. 2 tells us that if we fire a sequence of transitions and reach a configurations γ from γ_{init} then we can fire the same sequence of transitions from γ_{init} and reach a configuration with larger gap than in γ. Prop. 2 also implies that we do not have to retain gap between parameters

greater than $cmax$ since it is always possible to increase them. Terms with values smaller than $cmax$ must be treated in a special way.

More precisely, a configuration γ is encoded as a word $w_1 \cdot w_2$ where w_1 and w_2 are built as follows.

- Each ground term $p(c) \in \gamma$ with $0 \le c \le cmax$ is encoded as a (message) symbol (p, c). Thus, from γ we first extract the word

$$w_1 = w_1^0 \cdots w_1^{cmax}$$

  where w_1^i has many occurrences of (p, i) as those of $p(i)$ for any predicate $p \in \mathbb{P}$ and $0 \le i \le cmax$ (multiple occurrence of the same term produce different symbols, to disambiguate the encoding we assume a total order on symbols in $\mathbb{P}$).
- Each ground term $p(c) \in \gamma$ with $c > cmax$ is encoded as a symbol p. Thus, from γ we also extract the word

$$w_2 = p_1 \cdots p_k$$

  where $p_i(c_i) \in \gamma$, $c_i > cmax$ and $c_i < c_j$ for $1 \le i < j \le k$. Here we assume that there cannot be two terms with the same value for parameters greater than $cmax$.

Applying rewriting rules to words. W.l.o.g. we assume that each rule $L \rightsquigarrow R : \psi$ in $\mathcal{S}$ with set of variables V satisfies the following conditions.

- For each $x \in V$, either $(x = c) \in \psi$ and $0 \le c \le cmax$ or $(x > cmax) \in \psi$.
- Furthermore, we assume that for all pair of variables x, y in $L + R$ such that $x > cmax \in \psi$ and $y > cmax \in \psi$ we have that $x \bullet y \in \psi$ with $\bullet \in \{=, >, <\}$.

Given a Γ_0 rule ρ, we can compile ρ in a finite set of Γ_0 rules that satisfy the above mentioned conditions and that model the possible effects of applying ρ. The rules are obtained by completing the order in ρ with all possible missing relations between variables. By Prop. 3, we can safely introduce new equality constraints only when the resulting rule respecting the restrictions of Γ_0 (i.e. we do not need to introduce equality constraints involving more than two variables). As an example, the effect of the rule $[p(x), q(y)] \rightsquigarrow [r(z)] : x < y$ on a configuration in Γ_0 is modelled by the rules $[p(x), q(y)] \rightsquigarrow [p(z)] : x < y < z$, $[p(x), q(y)] \rightsquigarrow [p(z)] : x < z < y$, $[p(x), q(y)] \rightsquigarrow [p(z)] : z < x < y$, $[p(x), q(y)] \rightsquigarrow [p(x)] : x < y$, and $[p(x), q(y)] \rightsquigarrow [p(y)] : x < y$. Notice that in the last two rules we introduce an implicit equality between a variable in the rhs and a variable in the lhs.

Under these assumptions, a rule $\rho = L \rightsquigarrow R : \psi$ in $\mathcal{S}$ defined over the variables $V = \{x_1, \ldots, x_{m+r}\}$ can be represented by the word

$$w^\rho = w_1^\rho \cdot w_2^\rho$$

where w_1^ρ describes the effect of ρ on w_1, i.e. on ground terms with parameter smaller than $cmax$, and w_2^ρ describes the effect of ρ on w_2, i.e. on ground terms with parameter greater than $cmax$. More precisely,

$$w_2^\rho = \begin{bmatrix} \alpha_1 \\ \beta_1 \end{bmatrix} \cdots \begin{bmatrix} \alpha_r \\ \beta_r \end{bmatrix}$$

is the maximal sequence that satisfies the following conditions:

- For $1 \leq k \leq r$,
  - $\alpha_k = p \in \mathbb{P}$, if $p(x_{j_k}) \in L$ and $(x_{j_k} > cmax) \in \psi$,
  - $\beta_k = q \in \mathbb{P}$, if $q(x_{j_k}) \in R$ and $(x_{j_k} > cmax) \in \psi$,
  - α_k and β_k are equal to ϵ in all other cases.
- For $1 \leq k \leq r-1$, $x_{j_k} < x_{j_{k+1}}$ follows from ψ;

Pairs of the form $\begin{bmatrix} \epsilon \\ \epsilon \end{bmatrix}$ are not included in w_2^ρ. The word w_2^ρ specifies the order of terms in ρ and how a single term of a configuration (element in a word) is modified (using the pair $\begin{bmatrix} \alpha \\ \beta \end{bmatrix}$) by the rule. Notice that the syntactic restrictions of Γ_0 ensure that there cannot be elements $\begin{bmatrix} \alpha \\ \beta \end{bmatrix}$ with more than one predicate in α or β. Furthermore, if $\alpha = \epsilon$ then ρ adds a new occurrence of β, if $\beta = \epsilon$ then ρ removes an occurrence of α.

As an example, for $cmax = 2$, the rule ρ defined as $[p(x), q(y)] \rightsquigarrow [q(x), r(z)] : 2 < x < z < y$ is represented by the word

$$w_2^\rho = \begin{bmatrix} p \\ q \end{bmatrix} \cdot \begin{bmatrix} \epsilon \\ r \end{bmatrix} \cdot \begin{bmatrix} q \\ \epsilon \end{bmatrix}$$

The word $w_1^\rho = w_{1,0}^\rho \cdots w_{1,cmax}^\rho$ is such that for all $i : 0 \leq i \leq cmax, w_{1,i}^\rho = \varepsilon$ if there is no variable x such that $x = i \in \psi$, otherwise $w_{1,i}^\rho$ is a sequence of elements of the form $\begin{bmatrix} (p, i) \\ \epsilon \end{bmatrix}$ or $\begin{bmatrix} \epsilon \\ (p, i) \end{bmatrix}$ that describe the modifications that concerns ground terms with parameter $x = i$ as mentioned above.

Given a word $w_1 \cdot w_2$ associated to a configuration γ and a word $w_1^\rho \cdot w_2^\rho$ associated to a rule ρ, it should be clear now that the application of ρ to γ can be simulated by rewriting w_1 according to the ordered pairs in w_1^ρ and w_2 according to the ordered pairs in w_2^ρ. Clearly, the application of w_2^ρ to w_2 has some non-determinism, since we only have to ensure that in the resulting string w_2' the order in the two strings is preserved.

Going back to our example, the application of w_2^ρ to the word $w_2 = q \cdot p \cdot s \cdot q \cdot t$ produces the strings $q \cdot q \cdot r \cdot s \cdot t$ and $q \cdot q \cdot s \cdot r \cdot t$. (we recall that Prop. 3 tells us that we can safely ignore configurations in which r gets the same value as s).

Simulation in LCS. We are ready now to define the encoding of a Γ_0 $\mathcal{S}$ into an LCS $\mathcal{F}$. The LCS $\mathcal{F}$ has one channel c that contains the word encodings of configurations and one channel c' used as auxiliary memory. The control states of $\mathcal{F}$ are used for encoding different steps of simulation of a rule $\rho = L \rightsquigarrow R : \psi$ where we assume that all pair of variables x, y in ρ are in the relation $< \cup =$ induced by ψ. In particular, we will assume to have one distinct control state for each pair in w^ρ.

First, we simulate the effect on the terms with parameters less or equal than $cmax$. For a fixed i $1 \leq i \leq cmax$, the simulation consists in dequeuing

symbols of the form (p, i) from c, and by copying them into c' after applying the transformations defined in w_1^ρ. Notice that the information on the structure of w_1^ρ can be stored in the control states of $\mathcal{F}$. When there are no more symbols of the form (p, i) in c, we moves to the value $i + 1$.

Second, we simulate the effect of w_2^ρ for ground terms $p(c)$ with $c > cmax$. Suppose that w_2^ρ has r pairs $\begin{bmatrix} \alpha \\ \beta \end{bmatrix}$. Starting from the first pair in w_2^ρ, we define control states in which we either copy symbols from c to c' or, non deterministically, decide to apply the current pair $\begin{bmatrix} \alpha \\ \beta \end{bmatrix}$ to the head p of the queue in w_2^ρ:

- If $\alpha = p$ and $\beta = q$, then we remove p from c, add q to c', and move to the next pair in w_2^ρ;
- If $\alpha = p$ and $\beta = \epsilon$, then remove p from c, and move to the next pair in w_2^ρ;
- If $\alpha = \epsilon$ and $\beta = q$, we add q to c', and move to the next pair in w_2^ρ.

Note that since we non-deterministically choose the positions where modifications must be applied, the LCS $\mathcal{F}$ may get into a deadlock. Deadlocked computations do not influence the language $L_c(\mathcal{F})$.

Once the new word has been written into c' (and c is empty), we copy the content of c' into c and get ready for simulating the execution of another rule. In this last step we also recognize the symbols $\lambda(\rho)$ that labels ρ (all the other transitions used to simulate ρ are labelled by ε).

Finally, note that channels may lose messages. As a consequence, we encode *lossy* Γ_0 into LCS where ground terms may non-deterministically disappear during executions. However, r-language accepted by Γ_0 and lossy Γ_0 are the same (for coverability acceptance). $\qquad\square$

We show next that CMRS are strictly more expressive than LCS and Γ_0.

Theorem 5. $L_c(LCS) \subset L_c(CMRS)$.

Proof. We define a language L_{ent} which is accepted by a CMRS and that cannot be accepted by any LCS. Assume a finite alphabet Σ such that $\{\$, \#\} \not\subseteq \Sigma$. For each $w = a_1 \cdots a_k \in \Sigma^*$, we interpret w in the following as the multiset $[a_1, \ldots, a_k]$. Hence, we do not distinguish words in Σ^* from the multiset they represent, and vice versa. In particular, we will use the notation $a_1 \cdots a_k \leq a'_1 \cdots a'_l$ to denote that $[a_1, \ldots, a_k] \leq [a'_1, \ldots, a'_l]$. Define V to be the set of words of the form $w_1 \# w_2 \# \cdots \# w_n$ where $w_i \in \Sigma^*$ for each $i : 1 \leq i \leq n$. Consider $v = w_1 \# w_2 \# \cdots \# w_m \in V$ and $v' = w'_1 \# w'_2 \# \cdots \# w'_n \in V$. We write $v \sqsubseteq v'$ to denote that there is an injection $h: \{1, \ldots, m\} \mapsto \{1, \ldots, n\}$ such that

1. $1 \leq i < j \leq m$ implies $h(i) < h(j)$ (h is monotonic) and
2. $w_i \leq w'_{h(i)}$ ($\leq$ is multiset inclusion) for each $i : 1 \leq i \leq m$.

We now define the language $L_{ent} = \{v\$v' \mid v' \sqsubseteq v\} \subseteq (\Sigma \cup \{\#, \$\})^*$. As an example, given $\Sigma = \{a, b\}$, we have that $[a, b, b]\#[a, b, b]\#[a, a]\$[b, a]\#[a, a]$ is in L_{ent}, whereas $[a, b, b]\#[b, a, b]\#[a, a]\$[a, a]\#[a, b]$ is not in L_{ent}.

We now exhibit a CMRS $\mathcal{S}$ with $L_c(\mathcal{S}) = L_{ent}$. The set of predicate symbols which appear in $\mathcal{S}$ consists of (i) a predicate symbol a for each $a \in \Sigma$, and (ii) the symbols *guess*, *check*, $sep_\#$ and *ok*. The initial configuration γ_{init} is defined as $[guess(0)]$. Furthermore, we have the following rules:

(1) For each $a \in \Sigma$, we have a rule labelled with a and which is of the form

$$[guess(x)] \rightsquigarrow [guess(x) , a(x)] \quad : \quad true$$

Rules of this form are used to guess the letters in w_i in the first part of a word in L_{ent}. We keep track of the symbols inside w_i through their argument. These arguments are all the same by definition of the rule.

(2) A rule labelled with $\#$ of the form:

$$[guess(x)] \rightsquigarrow [sep_\#(x) , guess(y)] \quad : \quad \{x < y\}$$

This rule is used to switch from the guessing of the part w_i to the guessing of the next part w_{i+1}. $sep_\#(x)$ remembers the parameter on which the switch has been executed.

(3) A rule labelled with $\$$ of the form:

$$[guess(x) \,] \rightsquigarrow [check(y) , sep_\#(x)] \quad : \quad \{y = 0\}$$

This rule is used to switch from the guessing of the part $w_1 \# \ldots \# w_n$ to the selection of the second part of the word. The parameter of *check* is equal to the initial value of *guess*, i.e. to 0. This way, we can scan the word stored in the first phase from left-to-right, i.e., working on the argument order we define a monotonic injective mapping h.

(4) For each $a \in \Sigma$, we have a rule labelled with a which is of the form

$$[check(y) , a(y)] \rightsquigarrow [check(y)] \quad : \quad true$$

This rule is used to read a word (multiset) u_i contained in $w_{h(i)}$.

(5) A rule labelled with $\#$ of the form:

$$[check(x) , sep_\#(x) , sep_\#(y) \,] \rightsquigarrow [check(y) , sep_\#(y)] \quad : \quad \{x < y\}$$

This rule is used to pass from u_i to u_{i+1} for $i \geq 1$.

(6) A rule labelled with ϵ of the form:

$$[check(x) \,] \rightsquigarrow [ok(y)] \quad : \quad \{y = 0\}$$

This rule is used to non-deterministically terminate the checking phase. The accepting configuration γ_{acc} is defined as $[ok(0)]$.

Assuming that $\Sigma = \{a, b\}$, we now show that L_{ent} is not an LCS language. Suppose that $L_c(\mathcal{F}) = L_{ent}$ for some LCS $\mathcal{F} = (Q, \{c\}, M, \delta)$. We show that this leads to a contradiction. Let γ_{init} be the initial global state in $\mathcal{F}$ and γ_{acc} be the accepting global state. We use a binary encoding $enc : Q \cup M \mapsto \Sigma^*$ such that $enc(m) \not\leq enc(m')$ if $m \neq m'$. We will also use a special word $v_{init} \in \Sigma^*$ such

that $v_{init} \not\preceq enc(m)$ for each $m \in Q \cup M$. It is clear that such enc function and v_{init} exist. As an example, if $|Q \cup M| = n$ then we define enc as an injective map from $Q \cup M$ to multisets of $n+1$ elements with $i+1$ occurrences of a and $n-i$ occurrences of b for $0 \leq i \leq n$, and we use the multiset with $n+1$ occurrences of b for v_{init}. For instance, for $n=2$ we use $[a,a,a], [a,a,b], [a,b,b]$ for control states and messages and $[b,b,b]$ for v_{init}. We extend enc to global states such that if $\gamma = (q, m_1 m_2 \cdots m_n)$ then

$$ enc(\gamma) = enc(q)\#enc(m_1)\#enc(m_2)\# \cdots \#enc(m_n) $$

Observe that (i) $enc(\gamma) \in V$; (ii) for global states γ_1 and γ_2, it is the case that $\gamma_1 \preceq_l \gamma_2$ iff $enc(\gamma_1) \sqsubseteq enc(\gamma_2)$; and (iii) $v_{init} \not\sqsubseteq enc(\gamma)$ for each global state γ.

Since $L_{ent} = L_c(\mathcal{F})$ and $v\$v \in L_{ent}$ for each $v \in V$, it follows that for each $v \in V$, there is a global state γ such that $\gamma_{init} \xrightarrow{v} \gamma \xrightarrow{\$v} \gamma'$ with $\gamma_{acc} \preceq_l \gamma'$. We use $reach(v)$ to denote γ. We define two sequences $\gamma_0, \gamma_1, \gamma_2, \ldots$ of global states, and $v_0, v_1, v_2, \ldots$ of words in V such that $v_0 = v_{init}$, $\gamma_i = reach(v_i)$, and $v_{i+1} = enc(\gamma_i)$ for each $i \geq 0$. By Higman's theorem we know that there is a j such that $\gamma_i \preceq_l \gamma_j$ for some $i < j$. Let j be the smallest natural number satisfying this property. First, we show that $v_i \not\sqsubseteq v_j$. There are two cases: if $i = 0$ then $v_i \not\sqsubseteq v_j$ by (iii); if $i > 0$ then we know that $\gamma_{i-1} \not\preceq_l \gamma_{j-1}$ and hence, following (ii), $v_i = enc(\gamma_{i-1}) \not\sqsubseteq enc(\gamma_{j-1}) = v_j$. Since $\gamma_j = reach(v_j)$, we know that $\gamma_{init} \xrightarrow{v_j} \gamma_j$. By monotonicity, $\gamma_i \xrightarrow{\$v_i} \gamma_i'$, $\gamma_{acc} \preceq_l \gamma_i'$, $\gamma_i \preceq_l \gamma_j$ implies $\gamma_j \xrightarrow{\$v_i} \gamma_j'$ with $\gamma_{acc} \preceq_l \gamma_i' \preceq_l \gamma_j'$. We conclude that $\gamma_{init} \xrightarrow{v_j} \gamma_j \xrightarrow{\$v_i} \gamma_j'$ with $\gamma_{acc} \preceq_l \gamma_j'$. Hence, $v_j\$v_i \in L_c(\mathcal{F}) = L_{ent}$ which is a contradiction since $v_i \not\sqsubseteq v_j$. $\qquad \square$

Let us now consider r-languages. As mentioned at the beginning of the section, the expressive power of LCS remains the same as for coverability accepting conditions, However, this property does not hold anymore for Γ_0.

Proposition 7. $L_c(\Gamma_0) \subset L_r(\Gamma_0) = L_r(CMRS) = RE$.

Proof. It is well known that *perfect* FIFO channel systems with reachability accepting condition recognize the class RE. We prove that perfect channel systems accept the same languages as Γ_0 with reachability accepting condition. Given an LCS $\mathcal{F}$, let $\mathcal{S}$ be the Γ_0 used to encode an LCS in the proof of Theorem 4. In each step of a run σ in $\mathcal{S}$ the head and tail delimiters are moved to the right of their current positions. Thus, a "lost" ground term to left of the head delimiter, i.e. with parameter smaller than that of h_c, can never be removed in successive steps of σ. This implies that an accepting configuration in which all ground terms have parameters strictly greater than the parameter of the head delimiter characterize reachable configurations of a perfect FIFO channel system. $\qquad \square$

Hence, we have the following property.

Corollary 1. $L_r(LCS) \subset L_r(CMRS)$.

7 Petri Nets Extensions

Petri nets (PN), a well-known model of concurrent computation [21], can naturally be reformulated in a multiset rewriting system operating on nullary predicates only (i.e. predicates with no parameters). This class of rewriting rules corresponds to those in the fragment Γ_1 of CMRS defined in Section 3.2. It is easy to see that, if we associate a predicate symbol to each place of a net, configurations and rules of a Γ_1 model are just alternative representations of markings and transitions of a Petri net. As an immediate consequence of this connection, we have that $L_c(\Gamma_1) = L_c(PN)$ and $L_r(\Gamma_1) = L_r(PN)$. To formally compare Γ_1 with the other models, we use some known results on languages accepted by extensions of Petri nets. A *lossy Petri net with inhibitor arcs* (LN) is a Petri net in which it is possible to test if a place has no tokens and in which tokens may get lost before and after executing a transition. A *transfer net* [11] (TN) is a Petri net extended with transfer arcs.

Transfer nets (TN) are Petri nets extended with *transfer arcs*. A transfer arc is a pair $T \hookrightarrow q$ where T is a set of places of the net and q is a place. Given a set of places P, let us consider a transition $t = (I_t, O_t)$ with transfer $T \hookrightarrow q$. Given a marking M, t is fireable if $I_t \leq M$. Its firing leads to the new marking M' computed in three steps: we first compute $M_1 = M - I_t$, then we move all tokens in the places in T to the place q obtaining M_2; finally, we compute M' as $M_2 + O_t$.

Reset nets (RN) are Petri nets extended with *reset arcs*, i.e., with a transfer arc $T \hookrightarrow \bot$ where $\bot$ is a special place used only to reset places.

As an example, let $P = \{p, q, r, s\}$ and consider a transition t with $I = \{p, q\}$, $O = \{p, s\}$, and transfer arc $\{p, q\} \hookrightarrow r$. Now, consider the marking $M = [p, p, q, q, q]$. Then, the execution of t leads to the marking $M' = [p, s, r, r, r]$ (we first compute $M - I = [p, q, q]$, then execute the transfer obtaining $[r, r, r]$, and, finally, add O). If the transfer arc is instead $\{p, q\} \hookrightarrow \bot$, the execution of t leads to the marking $M' = [p, s]$ (marking do not refer to tokens in the special place $\bot$).

We first notice that the lossy version of RN(TN) (i.e. where tokens can be lost before and after applying the effect of transitions)define the same c-languages as RN(TN). We now prove that Lossy TN, Lossy RN and LN recognize the same class of c-languages.

- *Lossy TN as Lossy RN*: Let P and T be the set of places and transitions of a Lossy TN $\mathcal{N}$. We build a lossy RN with places P augmented by place n and s_t for each transition $t \in T$. The new places are used to distinguish normal transitions from simulations of the transfer of transition t. Consider now a transition t with label ℓ, $I_t = \{p_1, \ldots, p_m\}$, $O_t = \{q_1, \ldots, q_n\}$ and transfer $S \hookrightarrow q$ with $S = \{r_1, \ldots, r_k\}$. Transition t is simulated via the following set of transitions:
    - A transition t_0 labelled with ϵ such that $I_{t_1} = I \cup \{n\}$ and $O_{t_1} = \{s_t\}$. This transition checks if t is fireable and then activates the simulation of its transfer by adding a token to s_t.

- A set of transitions $t_1, \ldots, t_k$ labelled with ϵ such that $I_{t_i} = \{s_t, r_i\}$ and $O_{t_i} = \{s_t, q\}$ for $i : 1, \ldots, k$. Each such transition moves a single token from a places in S to q.
- A transition t' labelled with ℓ such that $I_{t'} = \{s_t\}$, $O_{t'} = O_t \cup \{n\}$, and with the reset arc $S \hookrightarrow \bot$. This transition non-deterministically terminates the simulation of the transfer.

- *Lossy RN as LN*: Given a Lossy RN $\mathcal{N}$, we can build a LPN+0 $\mathcal{N}'$ that accepts the same c-language simply by replacing each reset arc $T \hookrightarrow \bot$ of a transitions t with an inhibitor arc $Z_t = T$. Indeed, notice that the firing of a transition t with reset arc $T \hookrightarrow \bot$ in $\mathcal{N}$ at a marking M has the effect of forcing all places in T to be empty in the successor marking of M. Now, the corresponding transition t' in $\mathcal{N}'$ can be fired at M only if each place in T is empty in M. However, since $\mathcal{N}'$ is lossy this condition can always be verified (all tokens in places in T may get lost) and it has the same effect on M as t. Vice versa, if all tokens in places in T get lost, then we can fire t' and its firing as the same effect of t.
- *LN as Lossy TN*: Given a LN $\mathcal{N}$, we build a Lossy TN $\mathcal{N}'$ that accepts the same c-language simply by replacing each inhibitor arc $Z_t = T$ of a transitions t with a transition t' with a transfer arc $T \hookrightarrow p_t$ where p_t is a new place. We assume that tokens in p_t can never be re-used (i.e. p_t cannot occur in the preset of a transition in $\mathcal{N}'$). Indeed, notice that, since $\mathcal{N}$ is lossy, the inhibitor arcs $Z_t = T$ in $\mathcal{N}$ are enabled if we first lose all tokens in places in T. Thus, the inhibitor arcs have the same effect of a transfer to the new place p_t from which tokens can never be re-used. Vice versa, a transition t with transfer of all tokens of places in T to place p_t can be simulated by its corresponding transition with inhibitor arc t'. Indeed, in a lossy step all tokens in places in T may get lost thus enabling the inhibitor arc Z_t.
- *LN as LCS*: Given a LN $\mathcal{N}$ with places P and transitions in T, we build an LCS $\mathcal{F}$ that accepts the same c-language as follows. The LCS $\mathcal{F}$ has messages defined over the singleton set of symbols $\{t\}$. Furthermore, it uses a distinguished channel c_p to model each place $p \in P$. Thus, we use a queue c_p with k occurrences of t to simulate a place p with k tokens. Notice that we do not need to exploit the FIFO ordering of channels. Based on this idea, the simulation of a transition becomes straightforward. The consumption of a token from place p is simulated by a *dequeue* operation of message t executed on channel c_p, the production of a token in place p is simulated by an *enqueue* operation on channel c_p, and an inhibitor arc on place p is modelled by the *empty* test on channel c_p.

Thus, we have that $L_c(LN) = L_c(RN) = L_c(TN)$, $L_c(LN) \subseteq L_c(LCS)$, and, as for LCS, $L_r(LN) = L_c(LN)$.

Furthermore, in [14] the authors proved that $L_c(PN) \subset L_c(TN)$. From all these properties, we obtain the following result.

Theorem 6. $L_c(\Gamma_1) \subset L_c(\Gamma_0)$.

For r-languages, the classification changes as follows.

Theorem 7. $L_r(\Gamma_1) \not\sim L_r(LCS)$, $L_r(\Gamma_1) \not\sim L_r(LN)$, and $L_r(\Gamma_1) \subset L_r(\Gamma_0)$.

Proof. We first prove that $L_r(\Gamma_1) = L_r(PN) \not\subseteq L_c(LCS) = L_r(LCS)$, hence $L_r(\Gamma_1) \not\subseteq L_c(LN) = L_r(LN)$ since $L_c(LN) \subseteq L_c(LCS) = L_r(LCS)$. Consider the language $L = \{a^n b^n \mid n \geq 0\}$. It is easy to verify that there exists a Petri net $\mathcal{N}$ such that $L_r(\mathcal{N}) = L$. We now prove that $L \notin L_r(LCS)$. Per absurdum, suppose there exists an LCS $\mathcal{F}$ such that $L_c(\mathcal{F}) = L$. For any $k \geq 1$, let γ_k and γ'_k be two global states s.t. γ_{init} leads to γ_k by accepting the word a^k, γ_k leads to γ'_k by accepting the word b^k, and $\gamma_{acc} \preceq_l \gamma'_k$. Since $\preceq_l$ is a well-quasi ordering, there exists $i < j$ such that $\gamma_i \preceq_l \gamma_j$. By monotonicity of $\mathcal{F}$, we have γ_j leads to γ'' by accepting the word b^i and $\gamma_{acc} \preceq_l \gamma'_i \preceq_l \gamma''$. We conclude that $a^j b^i \in L_c(\mathcal{F})$ with $i < j$, which gives us a contradiction.

We now prove that $L_c(LN) \not\subseteq L_r(\Gamma_1)$, hence $L_c(LCS) \not\subseteq L_r(\Gamma_1)$. Let $\Sigma = \{a, b\}$ and let L_{par} be the language over the alphabet $\Sigma \cup \{\#\}$ that contains all the words $w_1 \# \ldots \# w_n$ with $n \geq 0$ such that $w_i \in \Sigma^*$ and there is no prefix of w_i that contains more occurrences of symbol b than those of symbol a, for $i : 1 \leq i \leq n$. Notice that the number of occurrences of symbols a and b in w_i may be different. The language can be accepted by a LN defined as follows. When we accept the symbol a we add one token in a special place p_a. To accept the symbol b, we remove one token from p_a. To pass from w_i to w_{i+1}, we accept symbol $\#$ whenever p_a is empty (in LN the empty test is just a reset). We now show that L_{par} cannot be recognized by a Petri net. Suppose that there exists a Petri net $\mathcal{N}$ such that $L_r(\mathcal{N}) = L_{par}$. Starting from $\mathcal{N}$, we build a net $\mathcal{N}_1$ by adding a new place d that keeps track of the difference between the number of occurrences of symbols a and b in the prefix of the word that is being processed in $\mathcal{N}$. Furthermore, we add the condition that d is empty to the accepting marking of $\mathcal{N}$. It is easy to verify that $\mathcal{N}_1$ accepts the language L_{bal} consisting of words of the form $w = w_1 \# \cdots \# w_n$ where w_i belongs the the language of *balanced parentheses* on the alphabet Σ for $i : 1 \leq i \leq n$. We exploit now [15, Lemma 9.8] that states that L_{bal} cannot be recognized by a Petri net with reachability accepting condition, which gives us a contradiction.

Finally, the property $L_r(\Gamma_1) = L_r(PN) \subset L_r(\Gamma_0)$ follows from [15, Lemma 9.8] and Prop. 7, Indeed, we have that $L_{bal} \in L_r(\Gamma_0) = RE$ and $L_{bal} \notin L_r(\Gamma_1)$.
□

Finally, we observe that we can use an argument similar to that used in the proof of Theorem 7 to show that $L_r(PN) \not\sim L_c(CMRS)$.

7.1 Affine well-structured nets

Affine well-structured nets (aWSNs) [18] are a generalization of Petri nets with black tokens and *whole-place* operations like reset and transfer arcs [11]. They can also be viewed as a subclass of data nets in which a configuration s is such that $s(d)(p) > 0$ only for a specific data d chosen a priori from D. Furthermore, all transitions have arity 1 and we can remove from F_t, H_t and G_t all the components in regions different from S_1, i.e., F_t and H_t are vectors in $\mathbb{N}^P$, and

G_t is a matrix in $\mathbb{N}^P \times \mathbb{N}^P$. $aWSN$ configurations, called markings, are vectors in $\mathbb{N}^P$, i.e., finite multisets with symbols in P. In the rest of the paper we use $[a_1, \ldots, a_n]$ to indicate a multiset with elements $a_1, \ldots, a_n$. Furthermore, for a multiset M, we use $M(a)$ to denote the number of occurrences of a in M. An aWSN-transition t is enabled at marking M if $F_t \leq M$ where $\leq$ denotes multiset inclusion. The firing of t at M produces a new marking $M' = (M - F_t)G_t + H_t$ where $-$ and $+$ denote resp. multiset difference and union. aWSN are well-structured with respect to inclusion of markings (vectors in $\mathbb{N}^P$).

Example 4. The projection of F_t, H_t and G_t in Fig. 1 on S_1 (i.e. restricted to the single data d_1) gives us the aWSN-transition t with $\alpha_t = 1$ defined as

$$F_t = (1\ 0)\ \ H_t = (0\ 1)\ \ G_t = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

This transition removes a token from p and resets the number of tokens in q to 1, i.e. for $M = (m_1, m_2)$ with $m_1 \geq 1$, it yields $M' = (m_1 - 1, 1)$.

We compare now aWSNs and LCSs.

Theorem 8. $L_c(aWSN) \subset L_c(LCS)$.

Proof. (1) We first prove the inclusion $L_c(aWSN) \subseteq L_c(LCS)$. Assume an aWSN $W = (P, T, F, G, H)$ with $P = \{p_1, \ldots, p_n\}$. We build an LCS $\mathcal{F} = (Q, C, N, \delta)$ such that $L_c(W) = L_c(\mathcal{F})$. The set of channels is defined as $C = P \cup P'$ where P' (auxiliary channels) contains a primed copy of each element in P. The set of messages N contains the symbol $\bullet$ (a representation of a black token).

Assume that $q_0 \in Q$ is the initial state of $\mathcal{F}$. Then, a marking $M = (m_1, \ldots, m_n)$ is encoded as an LCS configuration $enc(M)$ with state q_0 and in which channel $p_i \in P$ contains a word $\bullet^{m_i}$ containing m_i occurrences of symbol $\bullet$ for $i \in \overline{n}$.

For each transition t with label ℓ, we need to simulate the three steps (subtraction, multiplication, and addition) that correspond to F_t, G_t and H_t. Subtraction and addition can be simulated in a straightforward way by removing/adding the necessary number of tokens from/to each channel. The multiplication step is simulated as follows. For each $i \in \overline{n}$, we first make a copy of the content of channel p_i in the auxiliary channel p'_i. Each copy is defined by repeatedly moving a symbol from p_i to p'_i and terminates when p_i becomes empty. After the copy is terminated for all channels, we start the multiplication step. For each $i \in \overline{n}$, we remove a symbol from p'_i and add as many symbol to channel p_j as specified by $G_t(p_i, p_j)$ for $j \in \overline{n}$. The analysis terminates when the channels $p'_1, \ldots, p'_n$ are all empty. The following properties then hold:

i) We first notice that $M \leq M'$ iff $enc(M) \preceq_l enc(M')$;

ii) Furthermore, if $M_0 \stackrel{w}{\Longrightarrow} M_1$ in W, then $enc(M_0) \stackrel{w}{\Longrightarrow} enc(M_1)$ in $\mathcal{F}$;

iii) Finally, since $\bullet$ symbols may get lost in $\mathcal{F}$, if $enc(M_0) \stackrel{w}{\Longrightarrow} enc(M_1)$ then there exists M_2 such that $M_0 \stackrel{w}{\Longrightarrow} M_2$ and $M_1 \leq M_2$.

If the accepting marking is $M_f = (m_1, \ldots, m_k)$ then the accepting LCS configuration contains the control state q_0, the channel $p_i \in P$ contains m_i symbols $\bullet$, and the channels $p' \in P'$ are empty. Since we consider languages with coverability acceptance, $L_c(W) = L_c(\mathcal{F})$ immediately follows from properties *(i),(ii)*, *(iii)* and Lemma 1.

(2) We prove now that $L_c(LCS) \not\subseteq L_c(aWSN)$. For this purpose, we exhibit a language in $L_c(LCS)$ and prove that it cannot be recognized by any aWSN.

Fix a finite alphabet $\Sigma = \{a, b, \sharp\}$ and let $\mathcal{L} = \{w\sharp w' \mid w \in \{a, b\}^*$ and $w' \preceq_w w\}$. It is easy to define a LCS that accepts the language L: we first put w in a lossy channel and then remove one-by-one all of its messages. Thus, we have that $\mathcal{L} \in L_c(LCS)$. We now prove that there is no aWSN that accepts $\mathcal{L}$. Suppose it is not the case and there exists a aWSN N, with (say) n places, that recognizes $\mathcal{L}$ with initial marking M_{init} and accepting marking M_f.

For each $w \in \{a, b\}^*$, there is a marking M_w such that $M_{init} \overset{w\sharp}{\Longrightarrow} M_w \overset{w}{\Longrightarrow} M$ and $M_f \leq M$ (otherwise $w\sharp w$ would not be in $L_c(N)$). Consider the sequences $w_0, w_1, w_2, \ldots$ and $M_{w_0}, M_{w_1}, M_{w_2}, \ldots$ of words and markings defined as follows:

- $w_0 := b^n$;
- If $M_{w_i} = (m_1, \ldots, m_n)$ then $w_{i+1} := a^{m_1} b\, a^{m_2} b \cdots b\, a^{m_n}$, for $i = 0, 2, \ldots$

We observe that *(a)* $w_0 \not\preceq_w w_i$ for all $i > 0$, since w_0 contains n occurrences of b, while w_i contains only $n - 1$ occurrences of b; and *(b)* for any $i < j$, $M_{w_i} \leq M_{w_j}$ iff $w_{i+1} \preceq_w w_{j+1}$. By Dickson's lemma [9], there are $i < j$ such that $M_{w_i} \leq M_{w_j}$. Without loss of generality, we can assume that j is the smallest natural number satisfying this property. Remark that we have that $w_i \not\preceq_w w_j$. Indeed, $w_0 \not\preceq_w w_j$ for any $j > 0$ by *(a)*, and in the case of $i > 0$ we have by *(b)* that $w_i \not\preceq_w w_j$ since $M_{w_{i-1}} \not\leq M_{w_{j-1}}$. Since $M_{w_i} \leq M_{w_j}$, by monotonicity of aWSNs, we have that $M_{w_i} \overset{w_i}{\Longrightarrow} M$ with $M_f \leq M$ implies that $M_{w_j} \overset{w_i}{\Longrightarrow} M'$ with $M_f \leq M \leq M'$. Hence, $M_{init} \overset{w_j \sharp w_i}{\Longrightarrow} M'$ and $w_j \sharp w_i \in L_c(N) = \mathcal{L}$, which is a contradiction. $\square$
$\square$

It is interesting to notice that aWSNs can also be simulated by reset nets by using an encoding similar to the one based on LCSs. The emptiness test on a channel is replaced by a reset on the corresponding place. From this observation and from the results in [3], we have the following classification

Proposition 8. $L_c(Petri\ nets) \subset L_c(transfer/reset\ nets) = L_c(aWSN) \subset L_c(LCS)$.

This result shows that c-language recognized by reset/transfer nets are strictly included in those recognized by LCSs.

8 (Integral) Relational Automata

In this section we compare the class of languages accepted by the fragment Γ_2 of CMRS defined in Section 3.2 with those accepted by *relational automata* [8].

An *(integral) relational automaton* (RA) operates on a finite set X of positive integer variables, and is of the form (Q, δ) where Q and δ are finite sets of control states and transitions respectively. A transition is a triple (q_1, op, q_2) where $q_1, q_2 \in Q$ and op is of one of the following three operations: (i) *reading*: $read(x)$ reads a new value of variable x (i.e. assigns a non-deterministically chosen value to x), (ii) *assignment*: $x := y$ assigns the value of variable y to x; (ii) *testing*: $x < y$, $x = y$, $x < c$, $x = c$, and $x > c$ are guards which compare the values of variables x, y and the natural constant c. Assume a RA $\mathcal{A} = (Q, \delta)$. A *valuation* v is a mapping form X to $\mathbb{N}$. A *configuration* is of the form (q, v), where $q \in Q$ and v is a valuation. We define γ_{init} to be (q_{init}, v_{init}) where $q_{init} \in Q$ and $v_{init}(x) = 0$ for all $x \in X$. For a transition $\rho \in \delta$ of the form (q_1, op, q_2), we let $\gamma_1 \xrightarrow{\rho} \gamma_2$ if and only if $\gamma_1 = (q_1, v_1)$, $\gamma_2 = (q_2, v_2)$, and one of the following holds: $op = read(x)$ and $v_2(y) = v_1(y)$ for each $y \in X - \{x\}$; $op = (y := x)$, $v_2(z) = v_1(z)$ for each $z \in X - \{y\}$, and $v_2(y) = v_1(x)$; $op = (x < y)$, $v_2 = v_1$, and $v_1(x) < v_1(y)$. Other testing operations are defined in a similar manner. In [8] Cerans has shown that RA equipped with the *sparser-than* order of tuples of natural numbers are well-structured. In the case of RA with that order, the coverability accepting condition is equivalent to the control state acceptance, i.e., a word is accepted if it is recognized by an execution ending in a particular control state $q_{acc} \in Q$.

As stated in the following propositions, RA and Γ_2 define the same class of c- and r-languages.

Proposition 9. $L_c(\Gamma_2) = L_c(RA)$.

Proof. Given an RA $\mathcal{A} = (Q, \delta)$ over the set of variables X, we can build the Γ_2 $\mathcal{S}$ defined below. The set of predicate symbols in $\mathcal{S}$ consists of the following: (i) for each $q \in Q$, there is a predicate symbol q in $\mathcal{S}$; and (ii) for each variable x in X, there is a predicate symbol q_x in $\mathcal{S}$. Transitions in δ are encoded via the following CMRS rules (with the same labels)

$$
\begin{aligned}
(q_1, read(x), q_2) &\Rightarrow & [q_1, p_x(z)] &\rightsquigarrow [q_2, p_x(w)] : true \\
(q_1, x := y, q_2) &\Rightarrow & [q_1, p_x(z), p_y(w)] &\rightsquigarrow [q_2, p_x(w), p_y(w)] : true \\
(q_1, x < y, q_2) &\Rightarrow & [q_1, p_x(z), p_y(w)] &\rightsquigarrow [q_2, p_x(z), p_y(w)] : \{z < w\}
\end{aligned}
$$

For $X = \{x_1, \ldots, x_n\}$, the initial configuration is $\gamma_{init} = [q_0, p_{x_1}(0), \ldots, p_{x_n}(0)]$. The accepting configuration γ_{acc} is the multiset $[q_{acc}]$.

For the other inclusion, by using Prop. 2, we assume w.l.o.g. that there is no gap order formula $x <_c y$ with $c > 0$ in $\mathcal{S}$ and that $\gamma_{acc} = [ok]$. To justify the second assumption, notice that we can always introduce new predicate symbols ko and ok and a new rule that can be executed only on a configuration γ with $\gamma_{acc} \preceq_c \gamma$ and that replace ko with ok. All other rules are modified to be enabled only at configurations containing ko. Finally, we also observe that we can assume that all configurations of $\mathcal{S}$ have the same size (the size of the initial configuration of the Γ_2 model). Thus, we associate a variable of X to each ground term of the initial CMRS configuration and compose the predicate symbols in a CMRS configuration to form a single control state. CRMS rules can then be simulated

in several steps by operations on variables and updates of control states. To each control state containing ok, we add a transition labeled with ϵ to the accepting control state q_{acc}. $\qquad\square$

We now prove that $L_c(\Gamma_2)$ is the class of regular languages. For this purpose, we first need some preliminary definitions. Given a configuration γ with

$$index(\gamma) = D_0 \ldots D_{cmax} d_0 B_0 \ldots d_n B_n$$

we define

$$index'(\gamma) = D_0 \ldots D_{cmax} \sharp B_0 \ldots B_n.$$

Let us now consider an initial $cmax$-bounded configuration γ_{init} and a final $cmax$-bounded configuration γ_{acc}. The symbolic graph $\mathcal{G}_\mathcal{S}$ associated to a Γ_2 specification $\mathcal{S}$ is an automaton $(V, \to_{\mathcal{G}_\mathcal{S}}, c_0, F)$ where

- $V = \{index'(\gamma) \mid |\gamma| \leq |\gamma_{init}|\}$;
- $\to_{\mathcal{G}_\mathcal{S}} \subseteq V \times \mathcal{S} \times V$ such that $\forall c_1, c_2 \in V : (c_1, \rho, c_2) \in \to_{\mathcal{G}_\mathcal{S}}$ iff there exist two configurations γ_1 and γ_2 such that $index'(\gamma_1) = c_1, index'(\gamma_2) = c_2$ and $\gamma_1 \xrightarrow{\rho} \gamma_2$;
- $c_0 = index'(\gamma_{init})$;
- $F = \{index'(\gamma) \in V \mid \gamma_{acc} \preceq_c \gamma\}$.

We easily see that $\mathcal{G}_\mathcal{S}$ is a finite automata since the number of predicate symbols that appears in states is bounded by the size of γ_{init}.
In the following, we use $c \xrightarrow{\rho}_{\mathcal{G}_\mathcal{S}} c'$ to denote that $(c, \rho, c') \in \to_{\mathcal{G}_\mathcal{S}}$.
Moreover, given a sequence of rules $w = \rho_1 \ldots \rho_l$, $c \xrightarrow{w}_{\mathcal{G}_\mathcal{S}} c'$ denotes that there exists $c_1, \ldots, c_{l-1}$ in V such that $c \xrightarrow{\rho_1}_{\mathcal{G}_\mathcal{S}} c_1 \ldots \xrightarrow{\rho_l}_{\mathcal{G}_\mathcal{S}} c_{l-1} \xrightarrow{\rho_l}_{\mathcal{G}_\mathcal{S}} c'$ and $w = \rho_1 \cdots \rho_{l+1}$.
The next lemma states the main property of $\mathcal{G}_\mathcal{S}$: all the executions of $\mathcal{G}_\mathcal{S}$ corresponds to an execution in $\mathcal{S}$ starting from γ_{init}.

Lemma 4. *If* $index'(\gamma_{init}) \xrightarrow{w}_{\mathcal{G}_\mathcal{S}} c$, *then* $\gamma_{init} \xrightarrow{w} \gamma$ *such that* $index'(\gamma) = c$.

Proof. The proof is by induction on the number of transitions to reach c.
$n = 0$: Immediate.
$n > 0$: Let

$$index'(\gamma_{init}) \xrightarrow{w}_{\mathcal{G}_\mathcal{S}} c \xrightarrow{\rho}_{\mathcal{G}_\mathcal{S}} c'$$

and suppose that $c = D_0 \ldots D_{cmax} \sharp B_0 \ldots B_n$ and $c' = D'_0 \ldots D'_{cmax} \sharp B'_0 \ldots B'_{n'}$.
By ind. hypothesis, we have $\gamma_{init} \xrightarrow{w} \gamma$ with $index'(\gamma) = c$.
Since $c \xrightarrow{\rho}_{\mathcal{G}_\mathcal{S}} c'$, we know from definition of symbolic graph that there exist two configurations γ_1, γ_2 such that $index'(\gamma_1) = c$, $index'(\gamma_2) = c'$ and $\gamma_1 \xrightarrow{\rho} \gamma_2$.
Since

$$index'(\gamma) = index'(\gamma_1) = D_0 \ldots D_{cmax} \sharp B_0 \ldots B_n$$

we have

$$index(\gamma) = D_0 \ldots D_{cmax} d_0 B_0 \ldots d_n B_n$$

and
$$index(\gamma_1) = D_0 \ldots D_{cmax} d'_0 B_0 \ldots d'_n B_n$$

Moreover, suppose that
$$index(\gamma_2) = D'_0 \ldots D'_{cmax} b_0 B'_0 \ldots b_{n'} B_{n'}$$

Following Prop. 2, we have $\gamma_{init} \xrightarrow{w} \gamma'$ such that
$$index(\gamma') = D_0 \ldots D_{cmax} d''_0 B_0 \ldots d''_n B_n$$

and for any $i : 0 \leq i \leq n \ d''_i \geq d_i + d'_i$.

Since CMRS are monotonic (so do Γ_0) $\gamma_1 \xrightarrow{\rho} \gamma_2$ and $\gamma_1 \preceq_c \gamma'$, we have that $\gamma' \xrightarrow{\rho} \gamma''$ with $\gamma_2 \preceq_c \gamma''$. Since the number of ground terms in configurations is bounded, $\gamma_2 \preceq_c \gamma''$ implies $\gamma_2 \prec \gamma''$. Thus, we have that $index'(\gamma'') = index'(\gamma_2) = c'$. $\square$

Theorem 9. $L_c(\Gamma_2) = Regular\ Languages.$

Proof. We first show how to encode a finite automata in Γ_2. The encoding of a finite automaton is direct: each state corresponds to a predicate and CMRS rules mimic the transition relation. Acceptance of words is simulated as follows: for any final state c we have a rule $\{c(x)\} \rightsquigarrow \{p_{acc}(x)\} : true$ and the final configuration is $\{p_{acc}(0)\}$. Finally, the initial configuration is $\{c_0(0)\}$ where c_0 is the initial state of the automaton.

We now show that all the c-languages accepted by a Γ_2 are regular. Consider a Γ_2 $\mathcal{S}$ with an initial $cmax$-bounded configuration and a final $cmax$-bounded final configuration. Note that considering $cmax$-bounded configurations is not restrictive because we can always add a rule $\{p(x)\} \rightsquigarrow \{p(x)\} : (x = c)$, i.e. a rule that has no effect, where c is greater than the maximal constant that appears in the initial and final configurations. From Lemma 4 we have that a word accepted by the symbolic graph $\mathcal{G}_\mathcal{S}$ corresponds to sequence of rules corresponding to a word accepted by $\mathcal{S}$ (following definition of $\mathcal{G}_\mathcal{S}$, $\gamma_{acc} \preceq_c \gamma$ iff $index'(\gamma) \in F$). Moreover, from the definition of $\mathcal{G}_\mathcal{S}$ we have

$$\gamma_{init} \xrightarrow{\rho_1} \gamma_1 \xrightarrow{\rho_2} \ldots \xrightarrow{\rho_{l-1}} \gamma_l$$

implies

$$index'(\gamma_{init}) \xrightarrow{\rho_1}_{\mathcal{G}_\mathcal{S}} index'(\gamma_1) \xrightarrow{\rho_2}_{\mathcal{G}_\mathcal{S}} \ldots \xrightarrow{\rho_{l-1}}_{\mathcal{G}_\mathcal{S}} index'(\gamma_l)$$

Furthermore, from definition of accepting states F, $\gamma_{acc} \preceq_c \gamma_l$ if and only if $index'(\gamma_l) \in F$. Hence, if we replace symbols ρ in $\mathcal{G}_\mathcal{S}$ by $\lambda(\rho)$ we conclude that a words w is accepted by $\mathcal{S}$ if and only if w is accepted by $\mathcal{G}_\mathcal{S}$. $\square$

We are ready now to compare Γ_2 (hence RA) with the other models studied in this paper. For this purpose, we first observe that Petri nets can accept regular languages (finite automata can be encoded as Petri nets). Furthermore, it is straightforward to build a Petri net that accepts a non-regular language like $L = \{a^n \# b^m \mid n \geq m\}$. As a consequence of this observation and of Theorem 9, we have the following result.

Corollary 2. $L_c(\Gamma_2) \subset L_c(\Gamma_1)$.

Let us now consider the reachability accepting condition. We first notice that $L_c(\Gamma_2) = L_r(\Gamma_2) = L_c(RA) = L_r(RA)$. Indeed, in both cases of Γ_2 and RA we can encode the reachability acceptance into the coverability acceptance by adding transitions (labelled with ϵ) that can be fired only from the accepting configuration and leads to a configuration with control state q_{acc} in the case of RA and a configuration containing a special accepting predicate symbol in the case of Γ_2. Thus, we have the following property.

Theorem 10. $L_r(\Gamma_2) \subset L_r(\Gamma_1)$.

9 Conclusions

In this paper we have compared wsts by using languages with coverability acceptance as a measure of their expressinevess. From our results we obtain the following classification:

$$
\begin{pmatrix} L_c(FA) \\ = \\ L_c(RA) \end{pmatrix} \subset L_c(PN) \subset \begin{pmatrix} L_c(aWSN) \\ = \\ L_c(TN) \\ = \\ L_c(RN) \\ = \\ L_c(LN) \end{pmatrix} \subset L_c(LCS) \subset \begin{pmatrix} L_c(CMRS) \\ = \\ L_c(data\ nets) \end{pmatrix}
$$

Furthermore, since CMRS and Petri data nets (data nets without whole-place operations) recognize the same class of c-languages (coverability in CMRS can be reduced to coverability in Petri data nets [19]) we have that data nets, Petri data nets, and transfer data nets (another subclass of data nets with restrictions on the type of transfers) all define the same class of c-languages as CMRS.

References

1. P. A. Abdulla, K. Čerāns, B. Jonsson, and T. Yih-Kuen. General decidability theorems for infinite-state systems. LICS 1996: 313–321.
2. P. A. Abdulla and G. Delzanno: On the Coverability Problem for Constrained Multiset Rewriting Systems. AVIS 2006 (ETAPS 2006 workshop). Available in the authors' home page.
3. P. A. Abdulla, G. Delzanno, and L. Van Begin: Comparing the Expressive Power of Well-Structured Transition Systems. CSL 2007: 99-114
4. P. A. Abdulla and B. Jonsson: Verifying Programs with Unreliable Channels. Inf. Comput. 127(2): 91-101 (1996)
5. P. A. Abdulla and B. Jonsson. Undecidable verification problems for programs with unreliable channels. Inf. Comput. 130(1): 71-90 (1996)
6. N. Bertrand and Ph. Schnoebelen: A short visit to the STS hierarchy. ENTCS 154(3): 59-69 (2006)

7. G. Cécé, A. Finkel, and S. P. Iyer. Unreliable channels are easier to verify than perfect channels. Inf. Comput. 124(1): 20-31 (1996)

8. K. Čerāns. Deciding properties of integral relational automata. ICALP 1994: 35–46.

9. L. E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with r distinct prime factors. Amer. Journal Math. 35: 413422.

10. J. Esparza, A. Finkel, and R. Mayr: On the Verification of Broadcast Protocols. LICS 1999: 352-359

11. C. Dufourd, A. Finkel, and Ph. Schnoebelen: Reset Nets Between Decidability and Undecidability. ICALP 1998: 103-115

12. A. Finkel, G. Geeraerts, J.-F. Raskin, and L. Van Begin. On the ω-language expressive power of extended petri nets. TCS 356(3):374-386, 2006.

13. A. Finkel and Ph. Schnoebelen: Well-structured transition systems everywhere! TCS 256(1-2): 63-92 (2001)

14. G. Geeraerts, J.-F. Raskin, and L. Van Begin. Well-structured languages. Acta Informatica 44(3-4):249-288 (2007).

15. M. Hack. Petri net languages. Technical report TR-159, MIT, 1976.

16. T. A. Henzinger, R. Majumdar, and J.-F. Raskin. A classification of symbolic transition systems. ACM Trans. Comput. Log. 6(1): 1-32 (2005)

17. F. Rosa Velardo and D. de Frutos-Escrig: Name Creation vs. Replication in Petri Net Systems. ICATPN 2007: 402-422

18. A. Finkel, P. McKenzie, and C. Picaronny: A well-structured framework for analysing petri net extensions. Inf. Comput. 195(1-2): 1-29 (2004)

19. R. Lazic, T. C. Newcomb, J. Ouaknine, A. W. Roscoe, and J. Worrell: Nets with Tokens Which Carry Data. ICATPN 2007: 301-320

20. R. Mayr. Undecidable problems in unreliable computations. LATIN 2000: 377–386.

21. C. A. Petri. *Kommunikation mit Automaten*. PhD Thesis. Univ. of Bonn, 1962.

A   Appendix: Technical Lemmas and Propositions

A.1   Proof of Proposition 2

The proof is by induction on the size k of the sequence of rules $\rho_1 \ldots \rho_k$.

Base case: $(k = 0)$ Since we assume that all the constants that appear in γ_{init} are lesser or equal than $cmax$, we have that $index(\gamma_{init})$ is of the form $D_0 \ldots D_{cmax}$. Hence, the lemma trivially holds.

Induction step: $(k > 0)$ Suppose that we have $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_{k-1}} \gamma_1 \xrightarrow{\rho_k} \gamma$ with

$$index(\gamma_1) = E_0 \ldots E_{cmax} e_0 F_0 \ldots e_m F_m$$

Suppose that γ is built from γ_1 by applying the instance $\rho = L_1 \rightsquigarrow R_1$ of $\rho_k = L \rightsquigarrow R : \psi$. This means that there exists a multi-set of ground terms η such that $\gamma_1 = L_1 + \eta$ and $\gamma = R_1 + \eta$.

Under this hypothesis, the multisets in $index(\gamma_1)$ satisfy the following conditions:

- For any $i : 0 \leq i \leq cmax$, we have that $E_i = G_i + E_i^L$ where
  - E_i^L is the maximal (possibly empty) multiset of predicate symbols with parameter equal to i that occur in L_1;

- G_i is the maximal multiset of predicates with parameter i that are not consumed by ρ (i.e. they also occur in $index(\gamma)$).
  - For any $i : 0 \leq i \leq m$ and given $v_i = cmax + \Sigma_{0 \leq j \leq i} e_j$, we have that $F_i = H_i + F_i^L$ where
    - F_i^L is the maximal (possibly empty) multiset of predicate symbols with parameter equal to v_i that occur in L_1;
    - H_i is the maximal multiset of predicates with parameter v_i that are not consumed by ρ (i.e. they also occur in $index(\gamma)$).

Let us now suppose that instead of removing L_1 from γ_1, we add R_1 to γ_1. The resulting configuration $\gamma_2 = \gamma_1 + R_1$ has index

$$index(\gamma_2) = E_0' \dots E_{cmax}' \sigma_0 c_0 F_0' \dots \sigma_m c_m F_m' \sigma_{m+1}$$

with

$$\forall 0 \leq i \leq m+1 : \sigma_i = c_0^i K_0^i \dots c_{n_i}^i K_{n_i}^i$$

where (assuming $\forall 0 \leq k \leq m : e_k' = c_k + \Sigma_{0 \leq j \leq n_k} c_j^k$)

- For any $i : 0 \leq i \leq cmax$, we have that

$$E_i' = G_i + E_i^L + E_i^R$$

  where E_i^R is the maximal (possibly empty) multiset of predicate symbols with parameter equal to i that occur in R_1;
- for any $i : 0 \leq i \leq m+1$, for any $j : 0 \leq j \leq n_i$, K_j^i is the maximal (possibly empty) multiset of predicate symbols that occur in R_1 with parameter equal to
- Furthermore, for any $i : 0 \leq i \leq m$:

$$F_i' = G_i + F_i^L + F_i^R$$

  where F_i^R is the maximal (possibly empty) multiset of predicate symbols with parameter equal to $cmax + \Sigma_{0 \leq j \leq i} e_i'$ that occur in R_1.

Intuitively, σ_i represent the structure added to $index(\gamma_1)$ by R_1 for what concerned all predicate symbols with a parameter v not directly represented in $index(\gamma_1)$, i.e., such that

$$cmax + \Sigma_{0 \leq j \leq i-1} e_j < v < cmax + \Sigma_{0 \leq j \leq i} e_j.$$

The sequence

$$\varsigma = (E_0^L + E_0^R) \dots (E_{cmax}^L + E_{cmax}^R) \sigma_0 c_0 (F_0^L + F_0^R) \dots \sigma_m c_m (F_m^L + F_m^R) \sigma_{m+1}$$

can be transformed into $index(L_1 + R_1)$ by removing all empty multisets and summing up constants in order to correctly maintain gaps between non-empty multisets of predicates.

To simplify the presentation, let us assume that ς coincides with $index(L_1 + R_1)$. The following can be easily extended to the general case. We now observe that ρ corresponds to an instance of a specialization ρ' of ρ_k in which the variables in ρ_k are totally ordered w.r.t. $< \cup =$. In other words, from ς we can reconstruct the constraint ψ' of ρ' as follows:

- To each non-empty multiset M in ς we associate a distinct variable x_M, each predicate in M takes x_M as parameter in ρ';
- For each non empty multiset $M = E_i^L + E_i^R$ with $0 \le i \le cmax$, we associate the condition $x_M = i$ in ψ';
- If M and M' are two consecutive multisets in ς (and M occurs before M') with a constant c between them then ρ' contains the gap-order constraint $x_M <_c x_{M'}$.

Since the condition ψ' of ρ' corresponds to one of the possible linearizations of the condition of ρ_k, every instance of ρ' is also an instance of ρ_k. Furthermore, ψ' in ρ' represents the minimal gap-order constraints extracted from ς which is compatible with ψ (i.e. $\psi \wedge \psi' \equiv \psi'$). This implies that any other instance $L_2 \rightsquigarrow R_2$ of ρ' can be represented by a sequence

$$\varsigma_1 = (E_0^L + E_0^R) \ldots (E_{cmax}^L + E_{cmax}^R)\sigma'_0 f_0 (F_0^L + F_0^R) \ldots \sigma'_m f_m (F_m^L + F_m^R)\sigma'_{m+1}$$

where the following conditions are satisfied:

- for any $i : 0 \le i \le m$ we have $f_i \ge c_i$,
- for any $i : 0 \le i \le m+1, \sigma'_i = f_0^i K_0^i \ldots f_{n_i}^i K_{n_i}^i$ such that $f_j^i \ge c_j^i$ for all $j : 0 \le j \le n_i$.

Fixed a given instance $L_2 \rightsquigarrow R_2$ of ρ' with associated sequence ς_1, we define the new sequence

$$\varsigma_2 = S_0 \ldots S_{cmax}\sigma'_0 f_0 T_0 \ldots \sigma'_m f_m T_m \sigma_{m+1}$$

with $S_i = E_i^L + E_i^R + G_i$ for $i : 0 \le i \le cmax$ and $T_i = F_i^L + F_i^R + H_i$ for $i : 0 \le i \le m$.

Following the definition of $index$ and $\prec$, this sequence corresponds to the index of a configuration γ'_2 such that $\gamma + R_1 = \gamma_2 \prec \gamma'_2$. Now let us define the values

$$e_i'' = f_i + \Sigma_{0 \le j \le n_i} f_j^i$$

for $i : 0 \le i \le m$. Furthermore, let us define the sequence

$$\varsigma_3 = (E_0^L + G_0) \ldots (E_{cmax}^L + G_0)e_0''(F_0^L + H_0) \ldots e_m''(F_m^L + H_m)$$

Again, following definition of $index$ and $\prec$, there exists γ'_1 such that $\varsigma_3 = index(\gamma'_1)$ and $\gamma_1 \prec \gamma'_1$.

Now we note that γ'_2 corresponds to $\gamma'_1 + R_2$. This implies that the instance $L_2 \rightsquigarrow R_2$ of ρ can be applied at γ'_1.

If we now define $\gamma = \gamma_2 - L_1$ and $\gamma'' = \gamma'_2 - L_2$, then we have $\gamma \prec \gamma''$. Indeed, $index(\gamma)$ and $index(\gamma'')$ are obtained by removing predicate symbols in multiset occurring in the same position in $index(\gamma_2)$ and $index(\gamma'_2)$, respectively. Furthermore, we have that $\gamma'_1 \xrightarrow{\rho'} \gamma''$. Finally, note that for any sequence ς_1 there exists an instance ρ'' of ρ' (the specialization of ρ_k we consider). Hence, there exists an instance ρ'' such that $\gamma'' = \gamma'$.

By applying the inductive hypothesis, we have that there exists an execution $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_{k-1}} \gamma'_1$ that subsumes $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_{k-1}} \gamma_1$ such that $\gamma'_1 \xrightarrow{\rho_k} \gamma'$. We conclude that there exists an execution $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma'$ that subsumes $\gamma_{init} \xrightarrow{\rho_1 \ldots \rho_k} \gamma$. $\quad\square$

A.2 Proof of Lemma 2

If the length of the execution is one the thesis trivially holds because γ_{init} is $cmax$-bounded. Now suppose that

$$\gamma_{init} \xrightarrow{\rho_1} \gamma_1 \ldots \xrightarrow{\rho_k} \gamma_k \xrightarrow{\rho} \gamma$$

and γ is not linear. Suppose that $L_1 \rightsquigarrow R_1$ is the instance of ρ applied to γ_k to obtain γ. As in the proof of Prop. 2, we have that

$$index(\gamma_k) = E_0 \ldots E_{cmax} e_0 (F_0^L + G_0) \ldots e_m (F_m^L + G_m)$$

where F_i^L are the predicate symbols of terms in L_1 with parameter $cmax + \Sigma_{0 \leq j \leq i} e_j$ for $i : 0 \leq i \leq m$.
Now let us consider $\gamma_k + R_1$. Then we have that

$$index(\gamma_k + R_1) = E_0' \ldots E_{cmax}' \sigma_0 c_0 T_0 \ldots \sigma_m c_m T_m \sigma_{m+1}$$

where

- $E_i' = E_i + E_i^R$ where E_i^R are the predicate symbols of terms in R_1 with parameter i for $i : 0 \leq i \leq cmax$;
- $T_i = F_i^L + F_i^R + G_i$ where F_i^R are the predicate symbols of terms in R_1 with parameter $cmax + \Sigma_{0 \leq j \leq i} e_j$ for $i : 0 \leq i \leq m$; and
- σ_i is a sub-sequence
$$c_0^i K_0^i \ldots c_{n_i}^i K_{n_i}^i$$
that represents terms with new values added by R_1 for $i : 0 \leq i \leq m + 1$.

Since γ is not linear, there is a multiset $F_r^R + G_r$ or a multiset K_j^i that contains at least two predicates p and q with the same parameter say v. Let us suppose such a multiset $F_r^R + G_r$. The case where a multi-set K_j^i contains at least two predicates is treated is a similar way. Since by hypothesis $F_r^L + G_r$ contains at most one symbol and by the syntactic restriction of Γ_0, we have that (at least) one between $p(v)$ and $q(v)$ is produced by a valuation to a variable in ρ which is not involved in $=$ constraints.

Following Prop. 2, for any γ_k' with $\gamma_k \prec \gamma_k'$ we know that there exists a linear execution from γ_{init} to γ_k' with the same rules $\rho_1, \ldots, \rho_k$ and passing through $\gamma_1', \ldots, \gamma_{k-1}'$ such that $\gamma_i \prec \gamma_i'$ for $i : 1 \leq i \leq k$. This implies that we can choose γ_k' and instance $L_2 \rightsquigarrow R_2$ of ρ_k such that

$$index(\gamma_k' + R_2) = E_0' \ldots E_{cmax}' \sigma_0 c_0 T_0 \ldots \sigma_r c_r T_r \sigma_{r+1}' c_{r+1}' T_{r+1} \ldots \sigma_m' c_m' T_m \sigma_{m'+1}$$

where for any $j > r$ c_j' and the constants in σ_j' are strictly greater than c_j and the values in σ_j, respectively (i.e. we "shift to the right" all values greater than v).
Now notice that in a Γ_0 rule it is not possible to impose the equalities over more than two parameters. Furthermore, when imposing equality of two paramters of ground terms, one ground term is removed by the rule and the second one is

added to configurations by the rule. Hence, there is no constraints that impose that the parameter of $p(v)$ and $q(v)$ must be equal. W.l.o.g. we assume that tere is no constraint that impose that the parameter v of $p(v)$ must be equal to another parameter. This means that ρ remains applicable to γ'_k whenever the evaluation for the argument of predicate p is the value $v' = v + 1$.
With this new instance $L_3 \rightsquigarrow R_3$ of ρ we have that

$$index(\gamma'_k + R_3) = E'_0 \ldots E'_{cmax}\sigma_0 c_0 T_0 \ldots \sigma_r c_r T_r 1[p]\sigma''_{r+1} c'_{r+1} T_{r+1} \ldots \sigma'_m c'_m T_m \sigma_{m'+1}$$

where σ''_{r+1} is obtained from σ'_{r+1} by decrementing by 1 the first constant that appears.

We conclude by noticing that from $\gamma'_k + R_3$ we can compute γ' by removing L_3. This operation maintains the same structure of the index of $\gamma'_k + R_3$ for what concerns predicate $[p]$, i.e. assuming that $v = cmax + \sum_{j=1..i} d_i$ (remember d_i are constants of $index(\gamma)$) we have

$$index(\gamma') = D_0 \ldots D_{cmax} d_0 B_0 \ldots d_i B_i 1[p]d'_{i+1} B_{i+1} \ldots d'_n B_n$$

such that $\forall i + 1 \leq j \leq n : d'_j \geq d_j$. $\qquad\qquad\square$

A.3 Proof of Lemma 3

We prove by induction on the number of transitions that $\gamma_1 \overset{\rho_1 \ldots \rho_k}{\longrightarrow} \gamma_3$ implies there exists γ_4 such that $\gamma_2 \overset{\rho_1 \ldots \rho_k}{\longrightarrow} \gamma_4$ with either $\gamma_3 \prec \gamma_4$ or γ_4 is a linearisation of γ_3.

Base case:($k = 1$) Suppose that the lemma does not hold. Let $L_1 \rightsquigarrow R_1$ be the instance of ρ_1 that allows to build γ_3 from γ_1, i.e. $\gamma_3 = \gamma_1 - L_1 + R_1$. Suppose that $\gamma_1 = \gamma'_1 + [p(v)], \gamma_2 = \gamma'_2 + [p(v')]$ and $\gamma'_1 \prec \gamma'_2$. In other words, the predicate p is "isolated" in the index of γ_2. We consider two cases: either $L_1 \leq \gamma'_1$ or not, i.e. the instance $L_1 \rightsquigarrow R_1$ does not remove $p(v)$ or it does.

In the case of $L_1 \leq \gamma'_1$, let

$$index(L_1 + R_1) = E_1 \ldots E_{cmax} e_0 H_0 \ldots e_r H_r.$$

To any sequence,

$$\varsigma = E_1 \ldots E_{cmax} e'_0 H_0 \ldots e'_r H_r$$

such that $\forall 0 \leq i \leq r : e'_i \geq e'_r$ correspond another instance $L_2 \rightsquigarrow R_2$ of ρ_1. Furthermore, since γ_2 is a linearization of γ_1, all the constants that appear in $index(\gamma_2)$ are greater than the corresponding ones in $index(\gamma_1)$. Hence, there exists an instance $L_3 \rightsquigarrow R_3$ of ρ_1 that has the same effect on the structure of $index(\gamma_2)$ than the instance $L_1 \rightsquigarrow R_1$ on $index(\gamma_1)$, i.e. predicates are removed from and added to the same multi-sets and the same sequences of multi-sets (interleaved with constants) are added at the same point into $index(\gamma_2)$. Hence, $\gamma_2 \overset{\rho_1}{\longrightarrow} \gamma_4 = \gamma_2 - L_3 + R_3$ and γ_4 is a linearisation of γ_3.

In the second case, i.e. $L_1 \not\leq \gamma'_1$ and $p(v)$ is removed from γ_1 when applying $L_1 \rightsquigarrow R_1$, let

$$index(L_1 + R_1) = E_1 \ldots E_{cmax} e_0 H_0 \ldots e_i H_i + [p]e_{i+1} H_{i+1} \ldots e_r H_r$$

such that $v = cmax + \sum_{j=0..i} e_i$. Following the syntactic restriction of Γ_0, either (i) ρ_1 imposes no equality constraint between the parameter of $p(v)$ and the parameter of another ground term $q(v')$, or (ii) ρ_1 imposes such an equality constraint on the parameters of $p(v)$ and $q(v')$ which is added by ρ_1. In case (i), the sequence

$$\varsigma = E_1 \ldots E_{cmax} e_0 H_0 \ldots e_i H_i 1[p] e_{i+1} H_{i+1} \ldots e_r H_r$$

corresponds to another instance $L_2 \rightsquigarrow R_2$ of ρ_1, i.e. $index(L_2 + R_2) = \varsigma$, since the gap-orders between predicates defined by $L_1 + R_1$ are not violated. Furthermore, if we increase the gap orders defined by ς we still obtain a sequence that corresponds to an instance of ρ_1. Hence, any sequence

$$\varsigma' = E_1 \ldots E_{cmax} e'_0 H_0 \ldots e'_i H_i e[p] e'_{i+1} H_{i+1} \ldots e'_r H_r$$

such that $e \geq 1$ and $e'_i \geq e_i$ for any $i : 0 \leq i \leq r$ corresponds to an instance of ρ_1. Furthermore, since γ_2 is a linearisation of γ_1, all the constants that appear in $index(\gamma_2)$ are greater than the corresponding ones in $index(\gamma_1)$. Hence, there exists a instance $L_3 \rightsquigarrow R_3$ of ρ_1 that has the same effect on the structure of $index(\gamma_2)$ than $L_1 \rightsquigarrow R_1$ on $index(\gamma_1)$, except that $L_3 \rightsquigarrow R_3$ removes the multi-set $[p]$ that corresponds to the ground term $p(v)$. We conclude that $\gamma_2 \xrightarrow{\rho_1} \gamma_4 = \gamma_2 - L_3 + R_3$ and $\gamma_3 \prec \gamma_4$.

In case (ii), the sequence

$$\varsigma = E_1 \ldots E_{cmax} e_0 H_0 \ldots e_i H_i 1[p, q] e_{i+1} H_{i+1} \ldots e_r H_r$$

corresponds to another instance $L_2 \rightsquigarrow R_2$ of ρ_1, i.e. $index(L_2 + R_2) = \varsigma$, since the gap-orders (and equality between parameters of p and q) defined by $L_1 + R_1$ are not violated. Again, any sequence

$$\varsigma' = E_1 \ldots E_{cmax} e'_0 H_0 \ldots e'_i H_i e[p, q] e'_{i+1} H_{i+1} \ldots e'_r H_r$$

such that $e \geq 1$ and $e'_i \geq e_i$ for any $i : 0 \leq i \leq r$ corresponds to an instance of ρ_1. Furthermore, since γ_2 is a linearisation of γ_1, all the constants that appear in $index(\gamma_2)$ are greater than the corresponding ones in $index(\gamma_1)$. Hence, there exists a instance $L_3 \rightsquigarrow R_3$ of ρ_1 that has the same effect on the structure of $index(\gamma_2)$ than $L_1 \rightsquigarrow R_1$ on $index(\gamma_1)$, except that $L_3 \rightsquigarrow R_3$ replaces the multi-set $[p]$ that corresponds to the ground terms $p(v')$ by the multi-set $[q]$. We conclude that $\gamma_2 \xrightarrow{\rho_1} \gamma_4 = \gamma_2 - L_3 + R_3$ and γ_4 is a linearisation of γ_3.

Induction Step:$(k > 1)$ By induction Hypothesis, we have that $\gamma_1 \xrightarrow{\rho_1 \ldots \rho_{k-1}} \gamma'_3$ implies there exists γ'_4 such that $\gamma_2 \xrightarrow{\rho_1 \ldots \rho_{k-1}} \gamma'_4$ with either $\gamma'_3 \prec \gamma'_4$ or γ'_4 is a linearisation of γ'_3.

In the case where γ'_4 is a linearisation of γ'_3, we apply the same resoning than in the base base.

In the case where $\gamma'_3 \prec \gamma'_4$, let $\gamma'_3 \xrightarrow{\rho_k} \gamma_3$ and $L_1 \rightsquigarrow R_1$ be the instance of ρ_k used to build γ_3 from γ'_3, i.e. $\gamma_3 = \gamma'_3 - L_1 + R_3$. Consider

$$index(L_1 + R_1) = E_1 \ldots E_{cmax} e_0 H_0 \ldots e_r E_r.$$

To any sequence
$$\varsigma = E_1 \ldots E_{cmax} e'_0 H_0 \ldots e'_r E_r$$
with $\forall 0 \leq i \leq r : e_i \leq e'_i$ corresponds to another instance of ρ_k where the gap-orders between parameters of ground terms increase. Hence, there exists an instance $L_2 \rightsquigarrow R_2$ of ρ_k such that $\gamma_4 = \gamma'_4 - L_2 + R_2$ and $\gamma_3 \prec \gamma_4$. $\qquad \square$
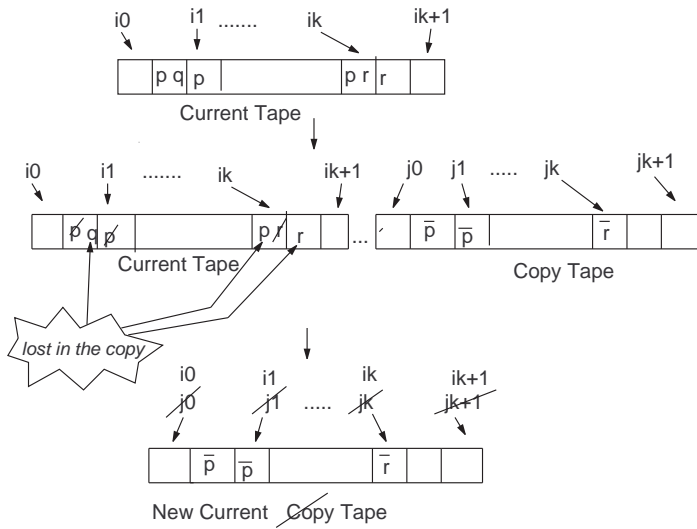
Fig. 9. Simulation of: (a) subtraction and selection of indexes; (b) creation of a new configuration.
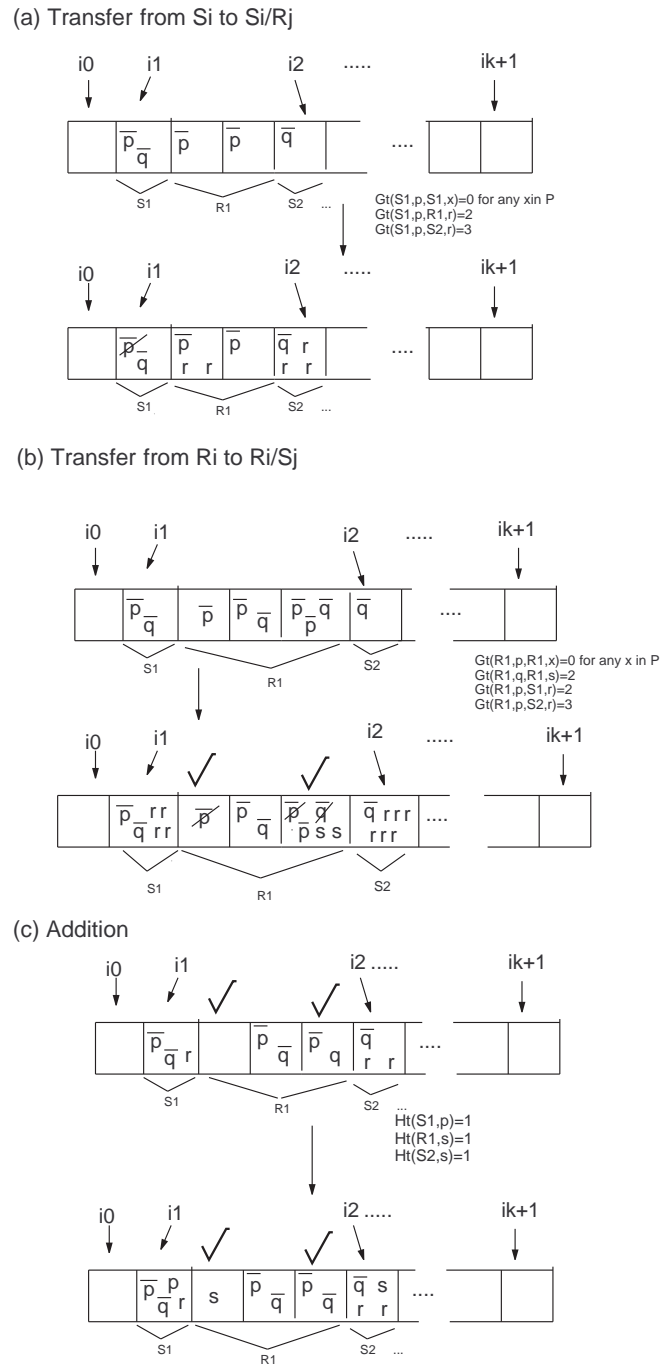
Fig. 10. Simulation of: (a) transfer from S_i to R_j/S_j; (b) transfer from R_i to R_i/S_j; (c) addition.