



HAL
open science

RDF: A Reconfigurable Dataflow Model of Computation

Pascal Fradet, Alain Girault, Ruby Krishnaswamy, Xavier Nicollin, Arash Shafiei

► **To cite this version:**

Pascal Fradet, Alain Girault, Ruby Krishnaswamy, Xavier Nicollin, Arash Shafiei. RDF: A Reconfigurable Dataflow Model of Computation. ACM Transactions on Embedded Computing Systems (TECS), 2022, 10.1145/3544972 . hal-03940615

HAL Id: hal-03940615

<https://inria.hal.science/hal-03940615>

Submitted on 16 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

RDF: A Reconfigurable Dataflow Model of Computation

PASCAL FRADET, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, France

ALAIN GIRAULT, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, France

RUBY KRISHNASWAMY, Orange Labs, France

XAVIER NICOLLIN, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, France

ARASH SHAFIEI, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, Orange, France

Dataflow Models of Computation (MoCs) are widely used in embedded systems, including multimedia processing, digital signal processing, telecommunications, and automatic control. In a dataflow MoC, an application is specified as a graph of actors connected by FIFO channels. One of the first and most popular dataflow MoCs, Synchronous Dataflow (SDF), provides static analyses to guarantee boundedness and liveness, which are key properties for embedded systems. However, SDF and most of its variants lacks the capability to express the dynamism needed by modern streaming applications. In particular, the applications mentioned above have a strong need for reconfigurability to accommodate changes in the input data, the control objectives, or the environment.

We address this need by proposing a new MoC called Reconfigurable Dataflow (RDF). RDF extends SDF with transformation rules that specify how and when the topology and actors of the graph may be reconfigured. Starting from an initial RDF graph and a set of transformation rules, an arbitrary number of new RDF graphs can be generated at runtime. A key feature of RDF is that it can be statically analyzed to guarantee that all possible graphs generated at runtime will be consistent and live. We introduce the RDF MoC, describe its associated static analyses, and present its implementation and some experimental results.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; • **Theory of computation** → **Parallel computing models**.

Additional Key Words and Phrases: models of computation, synchronous dataflow, reconfigurable systems, graph rewriting, static analyses, boundedness, liveness

ACM Reference Format:

Pascal Fradet, Alain Girault, Ruby Krishnaswamy, Xavier Nicollin, and Arash Shafiei. 2022. RDF: A Reconfigurable Dataflow Model of Computation. 1, 1 (June 2022), 30 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Dataflow Models of Computation (MoCs) are convenient for multimedia processing and digital signal processing since they model the application as a network of processing units, which is very natural for applications in these domains [18]. One of the first and most popular dataflow MoCs is Synchronous Dataflow (SDF) [20]. In a nutshell, an SDF graph

Authors' addresses: Pascal Fradet, Univ. Grenoble Alpes, and Inria, and CNRS, and Grenoble INP, and LIG, France, Pascal.Fradet@inria.fr; Alain Girault, Univ. Grenoble Alpes, and Inria, and CNRS, and Grenoble INP, and LIG, France, Alain.Girault@inria.fr; Ruby Krishnaswamy, Orange Labs, France, Ruby.Krishnaswamy@orange.com; Xavier Nicollin, Univ. Grenoble Alpes, and Inria, and CNRS, and Grenoble INP, and LIG, France, Xavier.Nicollin@inria.fr; Arash Shafiei, Univ. Grenoble Alpes, and Inria, and CNRS, and Grenoble INP, and LIG, and Orange, France, Arash.Shafiei@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

consists of so-called actors connected by FIFO channels. When it is executed (or fired in SDF terminology), an SDF actor consumes a fixed number of data (referred as tokens) on each of its input edges, performs some computation and produces a fixed number of tokens on each of its output edges. These numbers of consumed and produced tokens are fixed integers, which allows static analyses to check boundedness and liveness of SDF graphs.

Being able to check statically these properties is a strong advantage of SDF, but it comes at the price of forbidding dynamic changes of the graph. For this reason, several extensions of SDF have been explored, such as the parametric production and consumption rates (e.g., PSDF [4], BPDF [3], PiSDF [10]), allowing limited changes of the topology using scenarios (e.g., SADF [16]) or the possibility to dynamically enable and disable the edges of the graph (BPDF [3]). The common point of these variants is to remain statically analyzable [7], a crucial feature for embedded systems. Other MoCs have gone further along the road towards dynamicity (e.g., BDF [9] or DDF [19]), but properties such as boundedness or liveness become undecidable.

One aspect of dataflow MoCs that has not been explored is the *dynamic changes to the graph topology*. For example, this would be very useful for many kinds of embedded systems, including telecommunication applications (to allocate more pipelines when the number of IP packets to be handled increases or for software-defined radio), embedded video processing (to add filters as the light conditions change), automatic control (to change the control law depending on stability criteria).

We propose in this paper a variant of SDF called *Reconfigurable Dataflow (RDF)*. RDF allows dynamic changes to the graph topology thanks to *transformation rules* (expressed as graph rewrite rules) and to a *controller* that applies these rules depending on runtime conditions. In RDF, the number of graphs that can be produced using transformation rules is potentially *unbounded*. This contrasts with SADF where the number of scenarios is fixed and, in practice, rather small. We show that RDF remains statically analyzable and we describe conditions on transformations to ensure connectivity, boundedness, and liveness of RDF graphs. Finally, RDF programs can be deployed on different kinds of target architectures, from single-cores to multi-cores, possibly heterogeneous with DSPs and FPGAs.

The paper is organized as follows. We start by recalling the basic notions of SDF in Sec. 2 before presenting the RDF MoC in Sec. 3. Sec. 4 describes the conditions on transformations and static analyses ensuring that RDF reconfigurations preserve connectivity, consistency, and liveness. We present in Sec. 5 the main features of the implementation of RDF and provide some experimental results on a multi-core desktop in Sec. 6. Finally, Sec. 7 presents related work and Sec. 8 concludes. The appendix gathers the proofs of the theorems stated in Sec. 4.

This article extends and revises the work presented in [12]. Since then, RDF has been equipped with variable arity actors, an implementation has been developed, and experiments have been conducted. Sections 3.3, 5, and 6 are entirely novel, and other sections have been extended or rewritten. Explanations and examples have been added throughout. Additional details can be found in a PhD thesis [26].

2 SYNCHRONOUS DATAFLOW

An SDF graph [20] is a directed graph, where vertices – called *actors* – are functional units. Actors are connected by *edges*, which represent FIFO communication channels. The atomic execution of a given actor – referred to as actor *firing* – consumes data tokens from all its incoming edges (its *inputs*) and produces data tokens to all its outgoing edges (its *outputs*). The number of tokens consumed (resp. produced) on a given edge at each firing is called the *consumption* (resp. *production*) *rate*. An actor can fire only when *all* its input edges contain enough tokens (*i.e.*, at least the number specified by the consumption rate of the corresponding edge). In SDF, all rates are non-null integers known at compile time.

Formally, an SDF graph is defined by a 4-tuple $G = (V, E, \rho, \iota)$ where:

- V is a finite set of actors; among those, we distinguish *source* actors that have no incoming edges, and *sink* actors that have no outgoing edges;
- E is a finite set of directed edges: $E \subseteq V \times V$;
- $\rho : E \rightarrow \mathbb{N}^* \times \mathbb{N}^*$ is a function that returns for each edge a pair (x, y) , where x is the production rate of its origin actor (producer) and y is the consumption rate of its destination actor (consumer)¹;
- $\iota : E \rightarrow \mathbb{N}$ is a function that returns for each edge the number of its initial tokens (possibly 0).

When necessary, we will use V_G instead of V to refer to the set of vertices of graph G (and similarly for the other constituents).

Fig. 1 shows a simple SDF graph G_1 with 5 actors. The edge between A and B has a production rate of 2 and a consumption rate of 3.




Fig. 1. The SDF graph G_1 .

Each edge carries zero or more tokens at any moment. The *state* of a dataflow graph is the vector of the number of tokens present on each edge. The *initial state* of a graph is the vector of the number of initial tokens on its edges. For instance, the initial state of G_1 is the vector $[0; 0; 0; 0]$.

The *minimal iteration* of an SDF graph is a smallest set of firings of its actors such that (1) all actors fire at least once, and (2) the graph is returned to its initial state. For instance, the minimal iteration of G_1 is $(I^3, A^3, B^2, C^2, D^4)$, where X^i means that X is fired i times. We note $sol_G(X)$ the number of firings of X in the iteration of the graph G , or $sol(X)$ when no ambiguity can arise. The *basic repetition vector* \vec{Z} indicates the number of firings of actors per minimal iteration. For G_1 , it is $\vec{Z}_{G_1} = [3, 3, 2, 2, 4]$ assuming the actor ordering $[I, A, B, C, D]$.

An SDF graph is said to be *consistent* if it admits a repetition vector. The repetition vector is obtained by solving a *system of balance equations*. Each balance equation of that system corresponds to an edge of the graph. The edge $X \xrightarrow{p, q} Y$ is associated with the balance equation $sol(X) \cdot p = sol(Y) \cdot q$, which states that all produced tokens during an iteration must be consumed within the same iteration. The graph is consistent if and only if its system of balance equations admits a non-null solution [20], which is easy to check statically. An important advantage is that a consistent graph can be executed infinitely with *bounded* memory: all produced tokens are eventually consumed.

The next step is to determine a static order of the actors' firings, a *schedule*, in which the firings in the repetition vector can be executed without deadlock. It is obtained by an abstract computation where an actor is fired only when it has enough input tokens. Such a schedule ensures that the graph returns to its initial state and that each actor is eventually fired. A consistent SDF graph is said to be *live* if it admits a schedule [20].

Among all admissible schedules, we distinguish *single appearance schedules* (SAS)² where, once factorized³, each actor appears exactly once. For instance, G_1 admits exactly one SAS: $\{I^3; A^3; B^2; C^2; D^4\}$.

An acyclic SDF graph always admits an SAS, while a cyclic SDF graph admits an SAS if and only if each cycle includes at least one *saturated edge*, that is, an edge (X, Y) containing enough initial tokens to fire Y at least $sol(Y)$ times. In the context of this paper, we only consider SAS, but RDF can also operate with general schedules.

¹ \mathbb{N}^* is the set of non-null natural integers.

²Also called *flat* SASs in [1].

³Any sequence $X; \dots; X$ of n consecutive firings of X is replaced by X^n .

An SAS can be executed on a single-core chip or on a multi-core chip. On a single-core, it suffices to fire the actors sequentially as specified in the SAS. On a multi-core, each actor is first allocated to a core, and then on each core an ordering is chosen among all the actors allocated to it. In this paper, we consider *As Soon As Possible (ASAP)* scheduling, where each actor X is embedded in a private thread `act_X` consisting of the *periodic execution loop* presented in Fig. 2.

```

thread act_X {
  while (true) {
    consume_input_tokens();
    fire();
    produce_output_tokens();
  }
}

```

Fig. 2. Periodic execution loop for actor X .

The `consume_input_tokens` instruction blocks when (at least) one of the input buffers of X does not contain enough tokens, while the `produce_output_tokens` instruction blocks when (at least) one of the output buffers of X is full. On each core, one such thread `act_X` is started for each actor X allocated to it.

This multi-threaded ASAP execution guarantees that the graph can be executed in bounded memory and without deadlock, provided that each buffer has at least the minimal size required for liveness (which is easy to compute statically [22]). The dataflow semantics guarantees *functional determinism* whatever the order in which the actors are fired [18]. Moreover, provided enough cores and sufficiently large buffers, ASAP scheduling permits to achieve the maximal throughput.

3 RDF: A RECONFIGURABLE DATAFLOW MOC

The RDF MoC extends SDF with *transformation rules*, *actor types*, and *explicit ports*. A transformation rule describes how the current dataflow graph is modified. Actors and communication links can be moved, removed, and/or added. Adding new actors motivates the introduction of actor types. A type can be seen as a class of actors having the same functionality. Types allow transformation rules to add new actors in the graph as new type instances. For instance, a video application may require the dynamic introduction of several noise filter actors at different places in the graph. This may be done by introducing new actors in the graph, instances of the noise filter type. Transformation rules and type instances allow the number of actors and possible RDF graphs to be *unbounded*. RDF also introduces explicit actor ports to allow transformation rules to select specific edges more easily. For instance, ports allow two outgoing edges of the same actor and bearing the same producing rate to be discriminated.

An RDF application is specified as a pair (G, C) where:

- G is a dataflow graph, basically an SDF graph where each actor is equipped with a type;
- C is a *reconfiguration controller*, which consists of a set of transformation rules that specifies *how* an RDF graph may be transformed, and of a reconfiguration program using conditions to specify *when* the transformation rules should be applied.

An RDF application starts by executing its initial graph, until a condition is true and some transformation rules are applied, resulting in a new graph that is executed, and so on and so forth. The transformation rules allow a potentially infinite number of graphs to be produced dynamically from the initial graph.

3.1 RDF graph

RDF graphs extend SDF graphs with a set of *actor types* T and a notion of *ports*. Formally, an RDF graph is defined as a tuple $G = (T, V, E, \iota)$ where

- $T \subseteq Id_T \times (\mathbb{N} : k_i) \times (\mathbb{N} : k_o) \times ([1, k_i] \rightarrow \mathbb{N}^*) \times ([1, k_o] \rightarrow \mathbb{N}^*)$ is a finite set of types consisting of a unique identifier, a number of input ports, a number of output ports, and two functions returning the rate associated with each input and output port respectively. Accordingly, a type $t = (i, k_i, k_o, f_i, f_o)$ is composed of
 - an identifier i (a capital letter in this article);
 - two integers k_i and k_o denoting its numbers of input and output ports respectively;
 - two functions f_i and f_o returning the rate associated with their input and output port argument respectively.
 The auxiliary functions *idof*, *nbin*, *nbout*, *finr*, and *foutr* return respectively the identifier, number of input ports, number of output ports, input rate function, and output rate function of their type argument. For instance, $finr(T)(nbin(T))$ returns the rate of the last input port of type T ;
- $V \subseteq T \times \mathbb{N}^*$ is a finite set of actors, each one consisting of a type ($\tau \in T$) and an index ($i \in \mathbb{N}^*$). In the following, we use capital letters for type identifiers and we denote an actor of type X and index i by X_i . The functions *typeof* and *indof* return the type and index of their actor argument. Among actors, we distinguish *source* actors that have no incoming ports, and *sink* actors that have no outgoing ports;
- $E \subseteq (V \times \mathbb{N}^*) \times (V \times \mathbb{N}^*)$ is a finite set of directed edges. The edge $((a, i), (b, j))$ connects the i th output port of actor a to the j th input port of actor b . We will also denote an edge between a and b in the graph G by $a \xrightarrow[G]{} b$ and the fact that actors a and b are connected in G (i.e., $a \xrightarrow[G]{} b$ or $b \xrightarrow[G]{} a$) by $a \xleftrightarrow[G]{} b$;
- $\iota : E \rightarrow \mathbb{N}$ is a function that returns, for each edge, the number of its initial tokens (possibly 0).

We consider only well-formed graphs, that is, graphs properly connected and typed. In RDF, we check that initial graphs are well-formed and that transformations preserve well-formedness. Formally,

DEFINITION 1 (WELL-FORMEDNESS). *An RDF graph is well formed if it is (weakly) connected, all its actors are fully linked, and all its edges are valid.*

An RDF graph G is (weakly) connected if there exists an undirected path between any two actors a and a' , which we write $a \xleftrightarrow[G]^* a'$. Formally,

DEFINITION 2 (GRAPH CONNECTIVITY). *An RDF graph $G = (T, V, E, \iota)$ is weakly connected if and only if $\forall (a, a') \in V \times V$, we have $a \xleftrightarrow[G]^* a'$.*

Actors are *fully linked* if all their ports belong to a unique edge. Formally,

DEFINITION 3 (ACTORS FULLY LINKED). *An RDF graph $G = (T, V, E, \iota)$ has its actors fully linked if*

$$\forall b \in V, \forall 1 \leq i \leq nbin(typeof(b)), \forall 1 \leq o \leq nbout(typeof(b)), \\ \exists!(a, o') \in V \times \mathbb{N}^*, ((a, o'), (b, i)) \in E \wedge \exists!(c, i') \in V \times \mathbb{N}^*, ((b, o), (c, i')) \in E$$


Edges are *valid* if they connect actors only through ports permitted by the actors' type. Formally,

DEFINITION 4 (EDGE VALIDITY). *An RDF graph $G = (T, V, E, \iota)$ has valid edges if*

$$\forall ((a, o), (b, i)) \in E, 1 \leq o \leq nbout(typeof(a)) \wedge 1 \leq i \leq nbin(typeof(b))$$

To facilitate the reading, RDF graphs are often represented as in SDF with implicit ports and explicit rates. The graph of Fig. 1, can be seen as an RDF graph where S_1 , A_1 , B_1 , C_1 , and D_1 are actors with index 1 and types S , A , B , C , and D respectively. It has the same repetition vector and schedules as its SDF counterpart.

Another representation closer to the formal definition, with explicit ports, is given in Fig. 3.



with $foutr(I)(1) = 1$, $finr(A)(1) = 1$, $foutr(A)(1) = 2$, $finr(B)(1) = 3$, $foutr(B)(1) = 1$, \dots

Fig. 3. The RDF graph corresponding to G_1 with explicit ports.

3.2 Reconfiguration Controller

The RDF controller specifies when and how the dataflow graph is modified. The basic operations are *transformation rules* which are specified as graph rewrite rules. The *reconfiguration program* combines these transformation and specifies the conditions for their application. We present these two components in turn.

3.2.1 Transformation rules. An RDF transformation rule is a graph rewrite rule of the form

$$tr : lhs \Rightarrow rhs,$$

which selects a sub-graph matching the pattern lhs , and replaces it by the graph specified by rhs . We use the set-theoretic approach of [25] to graph rewriting: the terms lhs and rhs are seen as non empty sets of edges possibly with pattern variables matching either actor types, actor indices, ports, or rates.

Pattern variables require us to introduce the set \mathcal{V}_t of type variables, the set \mathcal{V}_i of actor index variables, the set \mathcal{V}_p of port variables, and the set \mathcal{V}_r of rate variables. With these sets, we define:

- the set of pattern nodes as $\tilde{V} \subseteq (T \cup \mathcal{V}_t) \times (\mathbb{N}^* \cup \mathcal{V}_i)$; for instance, A_1 , B_y , α_x all belong to \tilde{V} ;
- the set of pattern edges as $\tilde{E} \subseteq (\tilde{V} \times (\mathbb{N}^* \cup \mathcal{V}_p)) \times (\tilde{V} \times (\mathbb{N}^* \cup \mathcal{V}_r))$.



As it is standard in programming languages, pattern matching amounts to finding a variable substitution identifying the pattern with a sub-term. In RDF, a pattern lhs matches a sub-graph of G if there is a substitution σ mapping types (resp. indices, ports) variables to actual types (resp. indices, ports) such that the set of edges $\sigma(lhs)$ belongs to G : i.e., $\sigma(lhs) \subseteq G$. The rule removes the matched sub-graph and replaces it by rhs after substituting its variables by their matches, i.e., $\sigma(rhs)$.

In all examples, we note α , β , \dots the pattern variables matching types, x , y , \dots the pattern variables matching indices, p_1 , p_2 , \dots the pattern variables matching ports, and r_1 , r_2 , \dots the pattern variables matching rates. For instance, A_x matches any actor of type A and α_x matches any actor.

For the same reasons as we represent graphs with rates instead of explicit ports, we use patterns matching rates instead of ports. In the case of ambiguity, we may use explicit port index (such as $\textcircled{1}$ in Fig. 3) or port variables p_i in transformation rules.

As an example, consider the transformation rule tr_1 depicted in Fig. 4. In this figure, the terms α_x and β_z match any actor of any type, whereas the term B_y matches any actor of type B .


When applied to the graph of Fig. 1 (seen as an RDF graph), the lhs of the rule tr_1 matches the sub-graph


 Fig. 4. The transformation rule tr_1 .

 Fig. 5. The resulting graph $G_2 = tr_1(G_1)$.

and yields the substitution

$$\sigma = \{\alpha \mapsto A, x \mapsto 1, r_1 \mapsto 2, y \mapsto 1, r_2 \mapsto 1, \beta \mapsto C, z \mapsto 1\}.$$

The rule tr_1 replaces the actor B_1 by a new actor of type E . When a transformation introduces a new actor, its index is chosen so that the actor's name is fresh. Since no E actor occurs in G_1 , the variable t in tr_1 can be instantiated with index 1. As a result, tr_1 transforms the RDF graph G_1 of Fig. 1 into the RDF graph G_2 of Fig. 5.


 Fig. 5. The resulting graph $G_2 = tr_1(G_1)$.

The numbers of incoming and outgoing ports must be consistent with types. Actors occurring in the *lhs* and *rhs* must have the same number of edges in both parts (Cond. (C1)). Actors occurring only in the *lhs* or *rhs* must be fully linked: they must have explicit types and all their ports connected (Conds (C2) and (C3)). The following conditions must be respected by each transformation rule:

- (C1) An actor occurring in the *lhs* and in the *rhs* is *preserved* by the rule. It must have exactly the same edges and ports connected in the *rhs* and in the *lhs*. For an actor with an unknown type (*i.e.*, denoted by a pattern variable), that condition ensures that since it was fully linked before the transformation, it remains so afterwards.

For instance in tr_1 , actors α_x and β_z are preserved and it is easy to check that they have the same edges and ports connected in both sides.

- (C2) An actor occurring in the *lhs* but not in the *rhs* is *suppressed* by the rule. To be valid, all incoming and outgoing edges of that actor should appear in the *lhs*. Otherwise, suppressing an actor would create dangling edges. To verify this point, we request the type of removed actors to appear explicitly in the rule. Indeed, when the type is known, the numbers of incoming and outgoing edges are also known and the rule can be checked statically.

For instance in tr_1 , actor B_y is suppressed. In the *lhs*, it has exactly one incoming and one outgoing edge. It must be checked that type B has exactly one incoming and one outgoing edge.

- (C3) An actor occurring in the *rhs* but not in the *lhs* is *created* by the rule. Necessarily it must be of the form X_y with X an explicit type from T (*i.e.*, not a type variable) and y an index variable from \mathcal{V}_i . To guarantee statically that an actor is never created with unconnected ports, we check that it is fully linked in the *rhs*.

For instance in tr_1 , actor E_t is created, with the explicit type E . In the *rhs*, it has exactly one incoming and one outgoing edge. It must be checked that actors of type E have exactly one incoming and one outgoing edges.

These conditions are easy to check *syntactically* on each transformation rule. Additional constraints are required to guarantee that transformations preserve connectivity, consistency, and liveness. They are presented in Sec. 4.

RDF transformations can be formalized by representing a dataflow graph G as set of edges, and the rule $tr : lhs \Rightarrow rhs$ applied to G as the set rewrite rule

$$\underbrace{X \cup \sigma(lhs)}_G \Rightarrow \underbrace{X \cup \sigma(rhs)}_{G' = tr(G)}. \quad (1)$$

The graph G is the set of edges $X \cup \sigma(lhs)$ where σ is the substitution returned by the pattern matching. The resulting graph $G' = tr(G)$ is G where the sub-graph $\sigma(lhs)$ has been replaced by $\sigma(rhs)$. The context X (i.e., the graph or set of edges “surrounding” the matched part) remains unchanged.

A final remark is that initial tokens raise semantic issues. For instance, if the rhs of a transformation rule contains initial tokens, we would need a way to specify the origin or values of these tokens. To keep things simple, we allow the initial RDF graph to have edges with initial tokens but impose that transformations do not manipulate them. In other words, an edge with initial tokens cannot be matched nor created.

3.2.2 Reconfiguration programs. RDF transformation rules may be composed freely. The controller describes how to compose transformations into reconfiguration programs and when to apply these programs. A controller is specified by a *sequence* of pairs “(condition : reconfiguration program)” separated by semicolons:

$$[cond_1 : P_1; \dots; cond_n : P_n].$$

If a condition $cond_i$ is satisfied, then the controller stops the execution of the RDF graph *at the end of the iteration*, applies the transformations specified by P_i , and finally resumes the execution. Only one pair $(cond_i, P_i)$ is selected. If the conditions are not mutually exclusive, the first true condition in the sequence is chosen. Typically, the conditions depend on dynamic non-functional properties (e.g., buffer size, throughput, quality of the input signal, etc.). The language for describing these non-functional properties is not part of the MoC nor is it in the scope of this paper.

The simplest option for specifying reconfiguration programs is to consider them made of a single transformation. This is the language we used to perform our experiments (see Sec. 6). Many other, more expressive, options are possible. We describe in the following one such option. A reconfiguration program can be a combination of *transformation rules* with the following syntax:

$$\begin{array}{ll} P & ::= \quad tr \quad \textit{Transformation rule} \\ & | \quad P_1 \triangleright P_2 : P_3 \quad \textit{Choice} \\ & | \quad P^* \quad \textit{Iteration} \end{array}$$

The application of a transformation rule on a given RDF graph G is said to be *successful* if it has *matched* a sub-graph of G . By extension, an application of a program is considered successful if at least one of the transformation rules it tries to apply has been successful. The choice construction $P_1 \triangleright P_2 : P_3$ tries to apply P_1 ; if P_1 was successful, then P_2 is applied next, otherwise P_3 is applied. The iteration P^* applies P as long as it is successful. We can also write $P_1; P_2$ for the program $P_1 \triangleright P_2 : P_2$, which try to apply P_1 then P_2 in sequence regardless of the success or not of P_1 .

To ensure that a controller always preserves connectivity, consistency, and liveness of the dataflow graphs it transforms, it is sufficient to verify that the initial graph satisfies these properties *and* that each individual transformation rule preserves them. This will be the topic of Sec. 4.

This expressive language raises another issue, however: an iteration P^* may loop infinitely. To guarantee the termination of such iterations, a solution could be to enforce that P decreases some measure (e.g., the number of actors of type T in the graph).

3.3 Variable arity actors

An important application of RDF is to permit dataflow programs whose parallelism level can vary dynamically when needed by the environment (for instance according to some performance measures). Consider the dataflow graph G_3 of Fig. 6 that applies a filter F_1 on a flow of image macroblocks.



Fig. 6. The RDF graph G_3 with a single computing line.

When the resolution of the images in the video flow increases, it might be needed to increase the computational power and change the graph G_3 into the new graph G_4 of Fig. 7.




Fig. 7. The RDF graph G_4 with two computing lines.

In Fig. 7, the split actor S_1 reads two image blocks and distributes them towards the two filters F_1 and F_2 , while the join actor J_1 reads the two resulting blocks and passes them to actor O_1 . Provided enough hardware computing resources, the actors F_1 and F_2 can be fired in parallel and the throughput is thus improved compared to the initial RDF program.




Fig. 8. The RDF graph G_5 with three computing lines. The split actor S'_1 differs from the split actor S_1 from Fig. 7.

Should a third computation line be needed, one would have to introduce new split and join actors so as to distribute and read the three blocks, as shown in Fig. 8. The split actor S'_1 now reads three image blocks and distributes them to its outputs, so its type differs from that of S_1 . Similarly, the type of the join actor J'_1 differs from that of J_1 . The graphs G_4 and G_5 illustrate the complexity of modifying the graph topology to increase (and reduce) dynamically the number of computing lines:

- It requires an arbitrary number of split actor types like that of S_1 in Fig. 7 and S'_1 in Fig. 8 to distribute an arbitrary number of tokens read from its single input to all its outputs (and similarly for the join actor types like that of J_1 and J'_1).
- It requires an arbitrary number of transformation rules, because the rule used to increase the number of computing lines from 1 to 2 differs from the rule used to increase the number of computing lines from 2 to 3 (and similarly for the rules decreasing the number of computing lines).

To solve these issues, RDF provides a specific type named V for *variable arity* actors, shown in Fig. 9. To the best of our knowledge, RDF is the first dataflow MoC to offer such a variable arity actor.




Fig. 9. The variable arity actor type V (most general form).

When an actor of type V has i incoming edges and o outgoing edges, the consumption rate on each of its incoming edges is o and the production rate on each of its outgoing edges is i . At each firing, such an actor consumes and produces $i \cdot o$ tokens. It does not perform any computation; it just reads tokens from its input ports and distributes them evenly on its output ports.

The variable arity actor V_k is associated with two unique parameters: i_k representing both the input rates and the number of output ports, and o_k representing both the output rates and the number of input ports. Solving the system of balance equations is performed *symbolically* with those parameters, as in parametric variants of SDF [7]; the resulting iteration is also parametric. In this way, consistency is checked for all possible values of parameters.

Whenever a transformation adds or removes one or more edges of a variable arity actor V_k , the following modifications must occur:

- the number of ports $nbin(V_k)$ and $nbout(V_k)$ are updated;
- the value of the two parametric rates i_k and o_k are updated such that $i_k = nbout(V_k)$ and $o_k = nbin(V_k)$;
- ports are implemented as *lists* and adding a new edge involves adding a new port at the end of the list, while removing the edge of the ℓ th port makes the $\ell + 1$ th port become the ℓ th and so on; finally, the functions fnr and $foutr$ are updated such that:
 - $\forall 1 \leq \ell \leq i_k, fnr(V_k)(\ell) = o_k$,
 - $\forall 1 \leq \ell \leq o_k, foutr(V_k)(\ell) = i_k$.

To allow this updating, the functions $nbin$, $nbout$, fnr , and $foutr$ must now take as their first argument an *actor* instead of a *type*. Indeed, before introducing variable arity actors, all the actors of a given type T had exactly the same number of input and output ports. This is not the case anymore with variable arity actors.

Variable arity actors entail additional conditions on transformation rules. Indeed, suppressing a variable arity actor would require to select all its edges, the number of which cannot be statically known. Creating a new variable arity actor is also difficult since it involves introducing new parameters that play a role in the solutions of connected actors. We therefore enforce the following new condition on rules:

- (C4) Variable arity actors cannot be suppressed nor be created by transformation rules. All variable arity actors must appear in the initial graph.

Furthermore, a variable arity actor V_k is fully linked if (i) the value of its parameter i_k (resp. o_k) is equal to the number of its outgoing (resp. incoming) edges, and (ii) it has at least one incoming and one outgoing edge. Condition (i) is enforced by construction, as explained above. To enforce Condition (ii) for all possible RDF graphs generated dynamically, we add the following new condition:

- (C5) If a transformation rule removes an incoming (resp. outgoing) edge from a variable arity actor, then this actor must occur in the *rhs* with still at least one incoming (resp. outgoing) edge.

The rule t_{dec} in Fig. 11 removes a line of treatment and an outgoing and incoming edge of two variable arity actors (S_1 and J_1) while respecting that constraint.

Two special cases of this generic type V are particularly useful: the *split* actor type S and the *join* actor type J , depicted in Fig. 10 (and already seen in Fig. 11). The split actor type S has a single input with rate q , which is also the number of its outputs whose rates are 1. In other words, S is a special case of V where $p = 1$. The join actor type J has p inputs with rates 1 and a single output whose rate is p . In other words, J is a special case of V where $q = 1$. For an actor S_k (resp. J_k), we note its corresponding parameter s_k (resp. j_k). We only use splits and joins as variable arity actors in our examples and experiments.




Fig. 10. Variable arity actor types: (a) split S and (b) join J .

Other form of variable arity actors might be considered. Consider for example an actor *Sum* whose functionality consists in summing a collection of integer inputs. Such an actor could be represented as a variable arity actor returning the sum of its inputs on its single output whereas its number of inputs could be changed freely by transformations. Actually, any actor the functionality of which is defined on a list of inputs and/or outputs could be implemented as this form of variable arity actor. This generalization is presented in details in [26].


3.4 A complete RDF application

Fig. 11 shows an RDF application, with its initial graph G_6 , two transformation rules tr_{inc} and tr_{dec} , and a controller using two conditions for reconfigurations. This RDF application uses two variable arity actors, namely the split actor S_1 and the join actor J_1 . The reconfiguration controller applies the transformation rule tr_{inc} as soon as the throughput of the last actor O_1 drops below a threshold value equal to 20, and the transformation rule tr_{dec} as soon as the number of tokens present in the buffer from I_1 to S_1 drops below a threshold value equal to 10 and there are few data to process.

Applying tr_{inc} twice to G_6 yields the graph G_7 shown in Fig. 12, with $j_1 = 3$ and $s_1 = 3$, to be compared with the previous graph G_5 from Fig. 8. Rule tr_{inc} creates a new output port for actor S_1 (and similarly for J_1). This is only allowed for variable arity actors of course. As explained above, the ports are implemented as lists and the functions nb_{in} , nb_{out} , $finr$, and $foutr$ must be updated each time an input or output port is created or suppressed by a transformation rule. The parametric iteration for the initial dataflow graph G_6 of Fig. 11 is $(I_1^{s_1} S_1 F_1 J_1 O_1^1)$. When the transformation tr_{inc} is applied, the values of the parameters are updated and are propagated to all the actors with a number of firings per iteration depending on them (here, I_1 and O_1).

The RDF application of Fig. 11 implements an adaptive video processing application performing edge detection on a video stream with variable image quality. It will be used as a case study in Sec. 6.

Note that the following initial graph G'_6



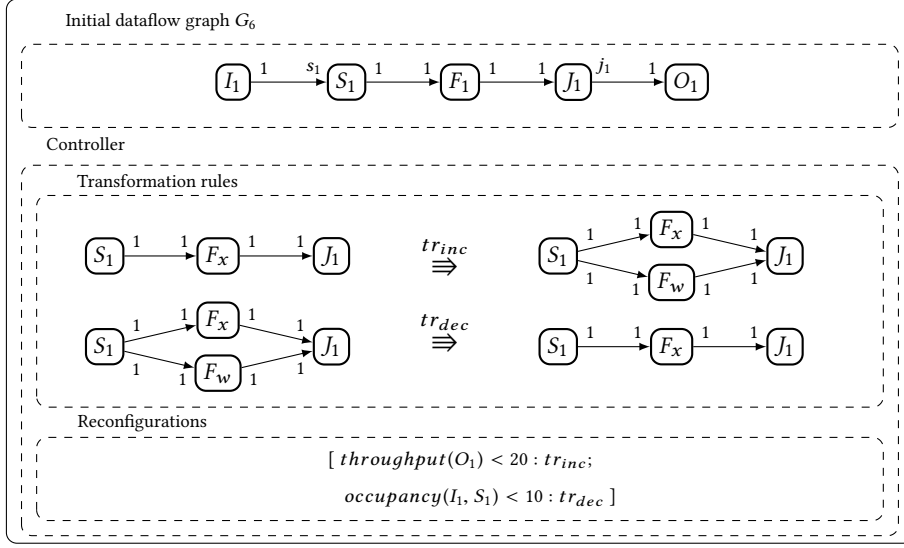




Fig. 11. An adaptive video processing application.

Fig. 12. The RDF graph $G_7 = tr_{inc}(tr_{inc}(G_6))$; we have $j_1 = 3$ and $s_1 = 3$.

would be inconsistent in the RDF MoC, because the edge between I_1 and O_1 requires that the parameters j_1 and s_1 be equal. This condition is indeed satisfied in the initial graph, but it could perfectly be invalidated by some transformation. Indeed, in contrast with the SDF MoC, the consistency of an RDF graph (initial or not) must be guaranteed *whatever* the transformation rules (see Sec. 4.2). For instance, adding an edge between S_1 and a new sink actor of type H , with the rule



would yield the following graph $G'_7 = tr_{sink}(G'_6)$:



which is inconsistent because the value of s_1 is now 2 while that of j_1 is still 1, so the iteration $(I_1^{s_1} S_1 F_1 H_1 J_1 O_1^{j_1})$ will result in an accumulation of tokens in the $I_1 \rightarrow O_1$ buffer. Recall that a rule like tr_{sink} is possible only if S_1 is a variable arity actor.

A simple extension would be to allow consistency checking (see Sec. 4.2) to produce such additional constraints and to statically check that all the transformations respect these constraints. In the following, we do not consider this extension and ensure consistency by checking that transformations do not change the solutions of existing actors.

4 RDF STATIC ANALYSES

The ability to guarantee statically consistency and liveness is of paramount importance for embedded systems. For this reason, improving the expressivity and dynamicity of SDF should not come at the price of losing these static analyses. This is the main technical issue of RDF. We present in this section how well-formedness, consistency, and liveness can be analyzed and guaranteed for RDF applications.

Of course, we want to guarantee that these three properties hold for all possible RDF graphs a given RDF application can generate at run-time. Our key contribution here is to show that it is sufficient:

- to check these three properties on the initial graph. SDF static analyses can be reused for that matter;
- to check that each individual transformation rule preserves these properties, that is to say, assuming that the considered property holds on the (unknown) source graph, it still holds on the transformed graph.

An RDF application is said to be valid if all its transformation rules satisfy these checks. Therefore, a valid RDF application transforms, produces, and runs only well-formed, consistent, and live graphs. We present in turn the conditions that a transformation rule must satisfy to preserve these three properties.

4.1 Well-formedness

We show that a well-formed graph (see Def. 1) remains so after a transformation respecting the previous conditions **(Ci)**. We have to show that all actors remain fully linked, all edges remain valid, and the graph remains weakly-connected.

All actors in the transformed graph remain fully linked since:

- Cond. **(C1)** ensures that all the other actors of the *rhs* keep the same ports connected, so they remain fully linked.
- Cond. **(C3)** ensures that newly introduced actors are fully linked.
- By construction, all the ports of a variable arity actor remain connected after a transformation (see Sec. 3.3).
- Finally, all the actors not present in the transformation rule remain untouched in the graph.

All edges in the *rhs* (new and remaining) can be checked to connect valid ports. Cond. **(C2)** ensures that removing an actor cannot create dangling edges. Therefore, all edges occurring in the graph remain valid.

SDF graphs are always connected, that is, there always exist an undirected path between every pair of vertices. In contrast, an RDF transformation rule removing edges could easily transform a connected graph into several disconnected ones. Theorem 1 states that, in order to guarantee that connectivity is preserved by the transformation rule $tr : lhs \Rightarrow rhs$, it is sufficient to ensure that *rhs* is a connected (pattern) graph $(x \xrightarrow[rhs]{*} y)$ states that there is an undirected path between x and y in *rhs*). Note that, in contrast to *rhs*, *lhs* may be disconnected, and therefore match disconnected subgraphs.

THEOREM 1. *Let G be a weakly connected graph and $tr : lhs \Rightarrow rhs$ be a transformation rule such that*

$$\forall x \neq y \in rhs, x \xrightarrow[rhs]{+} y, \quad (\mathbf{C}^{conn})$$

then $tr(G)$ is a weakly connected graph.

The proof of Theorem 1, as well as the proofs of Theorems 2 and 3, can be found in the appendix. Well-formedness follows from the preservation of complete linkage, validity, and connectivity.

COROLLARY 1. *Let G be a well-formed graph and $tr : lhs \Rightarrow rhs$ be a transformation rule satisfying the syntactic constraints of Sec.3.2.1 and (C^{conn}) , then $tr(G)$ is a well-formed graph.*

Clearly, the transformation tr_1 in Fig. 4 on page 7 preserves connectivity, but the rule tr_2 shown in Fig. 13 is invalid because its rhs is not a connected graph.




Fig. 13. The invalid transformation rule tr_2 .

Applying this transformation to G_1 of Fig. 1 would produce two disconnected graphs.

4.2 Consistency

The graph resulting from a transformation rule must remain consistent, meaning that its system of balance equations should have *non-null solutions*. Our condition for consistency, stated in Theorem 2, enforces a *stronger* property: all actors remaining in the transformed graph must keep their *original solution*.

For each transformation rule $tr : lhs \Rightarrow rhs$, we check that both pattern graphs lhs and rhs are consistent and we compute the (possibly symbolic) solutions of their actors. Actors occurring both in the lhs and rhs should have exactly the same solution. New actors (*i.e.*, occurring only in the rhs) only need to have a non-null solution, which is ensured by the fact that rhs must be consistent.

THEOREM 2. *Let G be a consistent graph and let $tr : lhs \Rightarrow rhs$ be a transformation rule such that lhs and rhs are consistent and*

$$\forall \alpha_x \in lhs \cap rhs, sol_{lhs}(\alpha_x) = sol_{rhs}(\alpha_x), \quad (C^{sol})$$

then $tr(G)$ is consistent.

Note that $sol_{pat}(\alpha_x)$ denotes the *minimal* symbolic solution (see [13]) of α_x in the system of equations corresponding to the pattern graph pat . If pat is a pure SDF graph, then this solution is an integer. In contrast, if pat has pattern variables matching rates, then the solution can also be computed and is, in general, symbolic. If pat has variable arity actors, then it also contains actors with parametric solutions. It is quite simple to deal with symbolic systems of equations and to define their minimal symbolic solutions [13].

Example: The transformation rule tr_1 of Fig. 4 (p. 7) preserves consistency. Indeed, both the lhs and rhs are consistent pattern graphs, and their common actors have the same symbolic solutions. Moreover, the solutions of the lhs actors are:

$$sol_{lhs}(\alpha_x) \quad sol_{lhs}(B_y) = \frac{r_1 sol_{lhs}(\alpha_x)}{3} \quad sol_{lhs}(\beta_z) = \frac{r_1 sol_{lhs}(\alpha_x)}{3 r_2}$$

while those the rhs actors are:

$$sol_{rhs}(\alpha_x) \quad sol_{rhs}(E_t) = \frac{r_1 sol_{rhs}(\alpha_x)}{3} \quad sol_{rhs}(\beta_z) = \frac{r_1 sol_{rhs}(\alpha_x)}{3 r_2}$$

The actors common to the *lhs* and *rhs* (α_x and β_z) keep their solutions, while the fresh actor E_t has a non-null solution. Besides, since $sol_{rhs}(E_t) = sol_{lhs}(B_y)$, we know that it is an integer solution.

Applied to the graph G_1 from Fig. 1, it yields the consistent graph $G_2 = tr_1(G_1)$ shown in Fig. 5. The actors I_1 , A_1 , C_1 , and D_1 keep their solutions (3, 3, 2, and 4, respectively), while the solution of the new actor E_1 is 2.

On the contrary, the transformation rule tr_3 of Fig. 14 is invalid. The reason is that, even though *rhs* is consistent, the actor B_y with solution $\frac{r_1 sol(\alpha_x)}{3}$ is replaced by actor F_u with solution $\frac{r_1 sol(\alpha_x)}{5}$, so we cannot be sure that this new solution is an integer.




Fig. 14. The invalid transformation rule tr_3 .

For instance, the graph $tr_3(G_1)$ is consistent but some of the solutions change: $sol_{G_1}(I_1) = sol_{G_1}(A_1) = 3$ but $sol_{tr_3(G_1)}(I_1) = sol_{tr_3(G_1)}(A_1) = 5$. Rules such as tr_3 can produce inconsistent graphs. For instance, when applied to the graph G_8 of Fig. 15a, tr_3 would produce the inconsistent graph G_9 of Fig. 15b.



Fig. 15. (a) Consistent graph G_8 . (b) Inconsistent graph $G_9 = tr_3(G_8)$.

The transformed graph G_9 is inconsistent since, even if the *rhs* of tr_3 is consistent, its system of balance equations

$$5 sol_{G_9}(E_1) = sol_{G_9}(F_1), 3 sol_{G_9}(E_1) = sol_{G_9}(H_1), sol_{G_9}(F_1) = sol_{G_9}(H_1)$$

does not have a non null solution. The reason is that the edge (E_1, H_1) enforces a constraint on the system of balance equations that does not appear in the transformation rule alone.

Note that we could have chosen a weaker condition for Theorem 2, namely

$$\forall \alpha_x \in lhs \cap rhs, \exists k, sol_{lhs}(\alpha_x) = k sol_{rhs}(\alpha_x). \quad (2)$$

This would allow a transformation to weaken some constraints (e.g., by removing edges) so that the minimal solutions of the *rhs* are possibly smaller than the solutions of *lhs*. In that case, consistency would be still preserved, the solutions of all actors could remain the same although they would not be minimal anymore.

4.3 Liveness

A consistent graph is live if it can be scheduled. We present here conditions on transformation rules so that they preserve liveness for graphs with single appearance schedules (SAS). The general case (i.e., a schedule exists, but is not an SAS) can also be dealt with, but it is more involved and would require more space to present. Recall also that, as stated in Sec. 3.2.1, transformation rules do not match nor create edges with initial tokens.

For each transformation rule $tr : lhs \Rightarrow rhs$, it suffices to check that rhs is live (*i.e.*, acyclic) and that tr does not add a path between common actors of lhs and rhs that did not exist before. These two conditions ensure that tr cannot introduce new cycles in the graph.

THEOREM 3. *Let G be a live graph with an SAS and $tr : lhs \Rightarrow rhs$ be a transformation rule such that rhs is live and*

$$\forall x, y \in lhs \cap rhs, x \xrightarrow{+}_{rhs} y \Rightarrow x \xrightarrow{+}_{lhs} y, \quad (\mathbf{C}^{live})$$

then $tr(G)$ is live and admits an SAS.

The transformation rule tr_1 of Fig. 4 (page 7) preserves liveness. Indeed, its rhs is live (it admits the schedule $[\alpha_x^{3r_2}; E_t^{r_1 r_2}; \beta_z^{r_1}]$) and it does not introduce new paths between actors occurring both in the lhs and rhs (namely between α_x and β_z).

On the other hand, the transformation tr_4 in Fig. 16 is invalid.




Fig. 16. The transformation rule tr_4 . All rates are 1.

Indeed, actor Y_y is connected to Z_z in the rhs but not in the lhs . If the only schedule in the initial graph is one where Z_z needs to be fired before Y_y , then rule tr_4 would produce a deadlocked (*i.e.*, non live) graph. Such a case is shown in Fig. 17.




Fig. 17. (a) Live graph G_{10} . (b) Deadlock graph $G_{11} = tr_4(G_{10})$. All rates are 1.

The rule tr_4 transforms the live graph G_{10} of Fig. 17a into the deadlocked graph G_{11} of Fig. 17b, where the new (and blocking) directed cycle created by tr_4 is highlighted in red.

5 IMPLEMENTATION

We have implemented a prototype of RDF to perform experiments, to evaluate the reconfiguration costs, and to explore its practicability. In this section, we present the main characteristics of our prototype. In particular, we describe (i) how an RDF application is executed in normal mode, *i.e.*, between reconfigurations; (ii) the steps needed to perform a reconfiguration; (iii) the kinds of conditions the controller may use; (iv) how the pattern matching of a transformation is implemented efficiently, and finally (v) how to deal with the placement of actors on a multi-core architecture when actors can be added or removed dynamically.

5.1 Standard execution

The initial graph is implemented by creating each actor as an instance of its type, and by allocating a circular buffer for each of its output ports. For an actor A and an output port with rate p , the size of the allocated buffer is $p \cdot \text{sol}(A) \cdot \text{sizeof}(\text{token})$. This is enough to achieve maximal throughput [22], but smaller bounds are known for special classes of graphs [6, 8]. The types of tokens as well as the values of the initial tokens, which are not part of the MoC, must be specified in the application. The types of the tokens allow to compute their size. The initial tokens are pushed in the circular buffer they belong to. Then, the communication links (the edges of the RDF graph) are created by providing to each input port a reference to the buffer it reads from.

Depending on the architecture (*e.g.*, single or multiple servers), actors can communicate (*i.e.*, read and write from and to buffers) through shared memory or message passing. Our prototype runs on a single multi-core processor and uses only buffers in shared memory to communicate. Finally, one thread is created for each actor as well as for the controller.

Actors are executed according to an as soon as possible (ASAP) policy, meaning that each actor fires as soon as it has enough tokens on all its incoming edges. When it does so, it extracts from each of its input buffers a number of tokens equal to the input rate of this buffer, it processes them according to its functionality, and finally it writes output tokens into its output buffers (see Fig. 2). Note that, at this step, it may have to wait to have enough available space in its output buffers.

Provided enough resources, all actors can run in parallel independently of each other. Synchronization is ensured by communication buffers. Using the ASAP schedule and properly sized buffers guarantees that the maximal throughput is obtained.

5.2 Reconfigurations

Reconfigurations cannot be performed at any moment. Indeed, transforming the dataflow graph in the middle of an iteration, or when all actors are not in the same iteration, would raise many semantic issues. Therefore, a reconfiguration should only occur when the RDF graph is in a coherent state, that is, after an iteration has completed and the graph has returned to its initial state (meaning implicitly that all actors have completed the same iteration).

Our prototype uses reconfiguration programs (see Sec. 3.2) made of a single transformation rule. The controller (which runs inside its own thread) is thus of the form

$$[cond_1 : lhs_1 \Rightarrow rhs_1; \dots; cond_n : lhs_n \Rightarrow rhs_n].$$

It periodically watches whether one of its reconfiguration condition $cond_i$ is satisfied. In our prototype and experiments, the reconfiguration conditions are mainly performance metrics (throughput, latency, buffer occupancy) measured at runtime. Many other criteria (*e.g.*, arity of split actors, internal variables of an actor, absolute time or delays) could be considered as well. Whenever a condition is true, the controller retrieves the rule corresponding to this condition and checks whether the lhs matches the current graph. If so, the transformation can be applied and the reconfiguration starts. Whenever several conditions are true, only the first matching rule in the controller list is selected.

As explained above, before applying a transformation, the graph must return to its initial state, and all actors must have completed the same iteration. To implement this, all actors keep track of their iteration number and of their number of firings within the current iteration. Since actors run in separate threads, at a given time they may not necessarily belong to the same iteration. Here are the main steps of a reconfiguration:

- When the controller has decided to apply a transformation rule, it prompts all actors for their iteration number; it then computes the maximum iteration number received and asks the actors to proceed until the end of this maximum iteration.
- On their side, all the actors stop at the end of their current iteration when they are prompted for their iteration number, *i.e.*, they finish all their firings but do not start a new iteration; when they receive from the controller the maximum iteration number, they either do nothing (if that number was indeed their own last iteration number) or they resume firing until they reach the end of the maximum iteration (otherwise). Then, they all send an acknowledgment to the controller.
- When the controller has received the acknowledgment from all the actors, the graph is its initial state and the transformation can be applied. The subgraph matching the *lhs* is replaced by the instantiated *rhs*. For each actor suppressed by the rule, we terminate its thread and we deallocate its buffers. For each actor created by the rule, we create a fresh name, we start a new thread, we allocate the new buffers, and we connect them.
- Finally, the controller asks all actors to resume their execution (or to start in the case of created actors). The computation proceeds as before, each actor firing as soon as its incoming edges have enough tokens.


Other, more local, options for reconfiguration could be implemented as well. For instance, stopping only the matched actors while ensuring that the matched edges are empty seems feasible and more efficient.

5.3 Pattern matching

Graph pattern matching is, in general, a NP-complete problem [15] that involves costly graph traversals with potentially many backtrackings. However, in our context of dynamic graph reconfiguration, this operation should be performed as fast as possible, ideally with a time complexity linear in the size of the *lhs* of the rule under consideration and without any backtracking.

To achieve this, we enforce that a *lhs* has at least one fully named actor *i.e.*, neither an actor variable such as α_x nor an actor type with an index variable such as B_x which can serve as a root. Further, we enforce that all edges of the *lhs* can be traversed starting from such roots by following explicit ports. Informally, each edge of the *lhs* must be reachable from a root actor via a *non-ambiguous* (undirected) path. That path is non unambiguous, because either the outgoing or incoming port of each edge is a fully identified one, *i.e.*, not a port variable.

Consider, for instance, the following pattern which may appear as the *lhs* of a transformation rule:



This pattern has two explicitly named actors that can be directly selected in the graph: A_5 and B_3 . From these roots, α_x , β_y , and γ_z and the four edges of the *lhs* can be selected unambiguously. Indeed, the edges (α_x, A_5) and (A_5, β_y) can be selected from A_5 by following the input port 1 and output port 2 of A_5 respectively. The edge (γ_z, B_3) can be selected (and actor γ_z determined) by following the input port 3 of B_3 , and then the edge (β_z, γ_z) from the input port 1 of γ_z .

If the actor A_5 was not named (*e.g.*, A_x), then the pattern matching would have to consider all typed A actors of the graph. Similarly, if the output port of A_5 was a variable, then the pattern matching would have to consider all possible ports. In both cases, this may involve failure and backtracking.

Our constraints can be relaxed. For instance, if the actor type A had a single input port, then the pattern would not need to make it explicit. Similarly, if the second output port of type A was the only one to have the rate 4, then the

pattern could use that rate instead to know with which edge to explore. The key idea here is that the pattern matching should always be able to proceed without performing choices.

Our approach guarantees that pattern matching can be achieved by traversing the pattern without backtracking, *i.e.*, with a time complexity *linear* in the size of the *lhs*. In practice, we have not noticed a loss of expressive power while observing these constraints. It is always possible to make patterns precise enough by introducing dummy named actors to act as pointers on the graph and as roots in patterns.

5.4 Placement strategy

Actors should be placed on the architecture so as to maximize parallelism and minimize communication costs (which may be quite high on distributed memory architectures). With more actors than resources, one should also take care of load balancing. A simple heuristic is to place a newly created actor X on the core that minimizes the load and the communication cost. These can be expressed in terms of the execution and production costs of this actor X during an iteration, the values of which are given by its type and by its solution $sol_G(X)$.

Our experiments were conducted on a multi-core, single socket server, where communication costs remain small. Preliminary experimental results showed that the heuristics described above was not significantly better than Linux's Completely Fair Scheduler (CFS), even with transformations drastically changing the initial graph. We then relied on CFS, by running each actor in a separate thread with equal priority. Whenever a transformation rule creates new actors, the CFS appears to place these new threads to optimize load balancing.

6 EXPERIMENTAL RESULTS

We experimented our prototype on an Intel Core i7-8700 CPU @ 3.20 GHz with 12 cores running Linux. We present an experimental evaluation of reconfiguration costs, show typical RDF transformations that change the throughput, and describe a small case study. All experiments consider graphs with a single source and sink actors.

Instead of measuring the throughput (number of iterations per time unit) of a graph G , denoted $\mathcal{T}(G)$, we consider a closely related measure, namely the number of tokens produced by the graph per time unit, denoted $\mathcal{H}(G)$. For a graph with a single sink actor O_1 , this measure is equal to the number of tokens consumed by O_1 per time unit.

Formally, the maximal throughput of an SDF graph G is determined by the actor whose execution takes the most time during an iteration. For an acyclic graph with a set of actors V , it is equal to

$$\mathcal{T}(G) = \frac{1}{\max_{v \in V} sol_G(v) t(v)}. \quad (3)$$

For a graph G with a single sink actor consuming n tokens per iteration, the number of tokens per time unit is

$$\mathcal{H}(G) = n \mathcal{T}(G). \quad (4)$$

The measure $\mathcal{H}(G)$ is more relevant than the throughput since the graph and the computation performed during an iteration may change dynamically. For instance, consider the initial graph G_6 and the transformation tr_{inc} of Fig. 11. This transformation does not change the throughput of G_6 since a fully parallel ASAP iteration takes the same execution time. In contrast, the number of tokens produced (and consumed) doubles. In other words, $\mathcal{T}(tr_{inc}(G_6)) = \mathcal{T}(G_6)$ while $\mathcal{H}(tr_{inc}(G_6)) = 2 \mathcal{H}(G_6)$.

In the following, we often use the term throughput to refer to the number of tokens produced.

6.1 Reconfiguration costs

An important point to evaluate is the cost of applying a transformation and the global reconfiguration cost. Indeed, RDF would lose part of its interest for streaming applications if reconfiguring takes too long. This cost can be decomposed in two parts:

- the cost of the transformation itself, *i.e.*, matching the *lhs* and replacing it by the *rhs*, possibly creating/suppressing actors and communication links;
- plus the cost of halting the graph’s execution and restarting it until it reaches again its steady state and maximal throughput.

In order to measure the transformation costs, we considered the two following dual transformation rules:

$$I \rightarrow O \xrightarrow{tr_5} I \rightarrow A_{x_1} \rightarrow \dots \rightarrow A_{x_n} \rightarrow O$$

$$I \rightarrow A_{x_1} \rightarrow \dots \rightarrow A_{x_n} \rightarrow O \xrightarrow{tr_6} I \rightarrow O$$

for various values of n . Experiments show that the matching and transformation costs are *linear* in the size of the rule. This was expected since there is no backtracking while matching and (de)allocating actors and buffers is main part of the costs. The transformation costs range from around 1 *ms* for $n = 10$ to 4 *ms* for $n = 40$.

To evaluate the total reconfiguration costs we used initial graphs of the form

$$I \rightarrow A_{x_1} \rightarrow \dots \rightarrow A_{x_n} \rightarrow O$$

and with the “identity” transformation $I \rightarrow A_1 \xrightarrow{tr_{id}} I \rightarrow A_1$. The difference of duration between the execution of the graph with the transformation rule tr_{id} and the duration without allows us to evaluate the cost of pausing and restarting the graph. For our setting, it is around 100 *ms*.

In conclusion, the cost of a reconfiguration remains low enough to be used in a video streaming application. This will be also shown in the case study of Sec. 6.3. If it was higher, a solution to make the reconfiguration seamless would be to introduce output buffers to continue the streaming and prevent glitches during reconfigurations.

6.2 Synthetic transformations

The great originality of RDF is to allow the dynamic reconfiguration of a dataflow graph of an application in order to increase its throughput when it is required. For instance, changing the desired resolution of a video stream could lead to increment or decrement the parallel levels of computation in a video streaming application.

To exemplify this, consider the initial graph G_6 of Fig. 11. The execution times of V_1 and O_1 are 10 *ms* each, S_1 and J_1 take 2 *ms* each, and C_1 takes 50 *ms*. Consider also the two transformation rules tr_{inc} and tr_{dec} , also from Fig. 11. To increase the throughput, we apply tr_{inc} at times 5 *s* and then 10 *s*. Then, to decrease the throughput, we apply tr_{dec} at times 20 *s* and then 25 *s*.

The evolution of G_6 ’s throughput is shown in Fig. 18. As it was expected, after applying tr_{inc} twice, the throughput increases from 20 to nearly 40 and then to 60 tokens/second. Then, by applying tr_{dec} twice, the throughput returns to its initial value. If the throughput is not exactly multiplied by 2 and then by 3, it is due to the small overhead of the split and join actors, which have to distribute and to gather tokens. Provided sufficient resources, these two rules are sufficient to adapt the application to any throughput demand. Other dynamic MoC like SADF would have to plan for a fixed number of configurations beforehand.

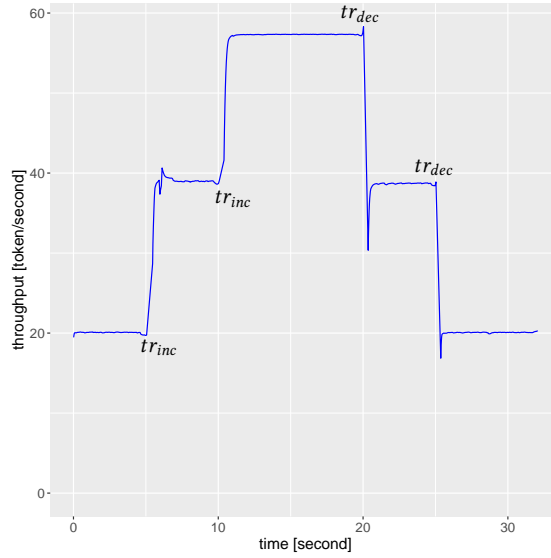


Fig. 18. Evolution of G_6 's throughput when applying successively tr_{inc} and tr_{dec} .

We use this kind of transformations in our case study (a Canny edge detection application) presented next.

6.3 Case study

We have used RDF to implement a Canny edge detection, an application that captures a video stream from a source, decodes it, detects the edges in the decoded images, and finally displays the images reduced to their edges with a constant frame rate. The quality of the input video stream can vary dynamically. When the quality increases, the processing (decoding & edge detection) takes more time and a static application may fail to maintain the same frame rate. We show that, by using transformation rules triggered by conditions on throughput, our RDF application can accommodate dynamic changes in the quality of the input video stream. The full RDF program is shown in Fig. 11.

We have implemented this example only partially. Our application reads the video stream from a file and we simulate the increase of processing time by using a dummy actor D_1 , the execution time of which increases artificially. The RDF modified dataflow graph G_c , including the dummy actor, therefore contains 6 actors (see Fig. 19): V_1 captures and decodes the input video stream; the decoded images are sent to the split actor S_1 that dispatches s_1 images to its s_1 output ports; C_1 receives an image and extracts its edges using a canny edge detector; D_1 is the dummy actor with a dynamically increasing execution time; J_1 is a join actor; and finally O_1 displays the output image.





Fig. 19. The initial RDF graph G_c of our Canny edge detection application.


The transformation in Fig. 20 adds one line of treatment and increases the arity of S_1 and J_1 , and increases accordingly the value of their parameters s_1 and j_1 . The controller applies this transformation whenever the throughput goes below 20 *fps*.

Fig. 20. The transformation rule tr_{canny} .

In our experiment, the execution time of actor D_1 increases from 30 ms to 60 ms at the 100th iteration, and then to 120 ms at the 200th iteration. Without any transformation (e.g., using SDF), the throughput decreases from 25 fps to 17 fps and then to 8 fps .

In Fig. 21, we see the throughput measured at the sink actor O_1 . At the beginning, the actor V_1 , whose execution time is 42 ms , is the most costly actor in the iteration and determines the throughput: $\mathcal{H}(G) = \frac{1}{0.042} = 24\text{ fps}$. When the execution time of actor D_1 increases to 60 ms , it becomes the most costly actor; the throughput thus drops to $\mathcal{H}(G) = \frac{1}{0.06} = 16.6\text{ fps}$. At iteration 110, the throughput goes below the threshold so the transformation tr_{canny} is applied. Actor V_1 becomes again the most costly actor since its solution after reconfiguration is 2 and its cost becomes 84 per iteration. Therefore, $\mathcal{H}(G) = 2 \times \frac{1}{2 \times 0.042} = 24\text{ fps}$.

At iteration 200, the execution time of actors D_1 and D_2 increases to 120 ms . Again the throughput decreases to $\mathcal{H}(G) = 2 \times \frac{1}{0.12} = 16.6\text{ fps}$. The transformation tr_{canny} is applied for the second time and the throughput returns to $\mathcal{H}(G) = 3 \times \frac{1}{3 \times 0.042} = 24\text{ fps}$.

Fig. 21. Evolution of the throughput of our Canny edge detection application when the transformation tr_{canny} is applied twice.

The reconfiguration costs are small enough so that the two reconfigurations are hardly visible when viewing the output video stream. This case study shows that RDF can be used to design easily adaptive image processing that can maintain the throughput at a desired value even with dynamic resolution changes.

We have seen in this section an example where RDF rules allow to maintain the throughput whereas the quality of the input video stream increases dynamically. It does so by adding more parallelism thanks to the variable arity actors.

Other options could have been chosen as well. For instance, changing the canny edge detector by coarser but more efficient versions could be considered to maintain the throughput on single-core systems. Let us mention that there are many other kinds of embedded systems requiring dynamic changes for which RDF can be useful. As one example, in software-defined radio [17], RDF could specify the switches between a GMSK and an QPSK demodulator depending on whether the system is receiving GSM or UMTS packets.

7 RELATED WORK

Many different dataflow MoCs have been proposed in the few last decades. More recently, several *parametric* dataflow MoCs have been presented as an interesting trade-off between expressiveness and analyzability. Among the existing parametric MoCs let us cite PSDF [5], VRDF [28], SPDF [14], BPDF [3], PiSDF [11], and PFSM-SADF [27]. They all offer a controlled form of dynamism under the form of parameters (*e.g.*, parametric rates) along with run-time parameter configuration.

Among those, BPDF [3] can model dynamic topology changes by adding Boolean conditions to FIFO channels. When a condition switches to false (resp. true) the channel is *disabled* (resp. *enabled*). Boundedness and liveness remain statically analyzable, and static or quasi-static schedules can be produced [2].

A different approach is taken by SADF [16] and its parametric version PFSM-SADF [27]. Here reconfigurability is modeled as a set of predefined configurations (called scenarios), coupled with a non-deterministic finite-state machine that specifies the transitions between these scenarios. The number of available topologies is statically fixed and specified in the source model. Analyzing a (PFSM-)SADF model relies on the standard analyses of SDF applied to each scenario.

Both BPDF and SADF only allow a fixed, and usually small, number of graph topologies. Imagine a video application that may need to apply a collection of n filters depending on the characteristics of the input video stream. Since there are 2^n combinations of such filters, describing so many graph configurations would be cumbersome in such MoCs.

To the best of our knowledge, RDF is the only dataflow MoC allowing both the dynamic reconfiguration (in the general sense) of the graph topology and static analyses for boundedness and liveness. It allows the engineer to generate an unbounded number of consistent and live graphs that do not have to be planned nor declared in advance.

Reconfigurability using rewriting rules has also been studied for Petri nets (see [23] for an overview). In the general case, reconfigurable Petri nets do not preserve properties such as liveness, boundedness, or reversibility. In [21], a restricted class of transformations (called INRS) is proposed that preserves these properties. It has been applied to design Petri net controllers for the supervision of reconfigurable manufacturing systems. Model checking of reconfigurable Petri nets has been considered by converting the net and the set of rewriting rules into a Maude specification [24]. This approach allows the absence of deadlocks to be verified.

8 CONCLUSION

We addressed the question of dynamic reconfigurations of SDF graphs. To this aim, we introduced the RDF MoC consisting in a dataflow graph (an SDF graph with typed actors) and a controller (a sequence of reconfiguration programs triggered by conditions). The transformation rules determine *how* the RDF graph is reconfigured and the conditions specify *when* these reconfigurations take place. A key feature and advantage of RDF is that it retains static analyses to guarantee boundedness and liveness properties of all possible graphs produced by the dynamic reconfigurations. To allow richer reconfiguration of the graph topology, we introduced a notion of *variable arity* actors, which does not seem to exist in other dataflow MoCs found in the literature. Finally, we outlined the main characteristics

of our RDF prototype implementation and presented some experiments. They show that the dynamic reconfiguration cost of an RDF graph is, in practice, very small.

Several extensions of RDF come to mind. First, RDF rates could be generalized to accept parameters. Such rates are different from the parametric rates of variable arity actors that denote their number of edges (incoming edges for the join actor and outgoing edges for the split actor). In parametric MoCs, a rate parameter may be changed, usually between iterations, to an arbitrary value. We expect this generalization to be relatively straightforward. Indeed, in such MoCs, static analyses become parametric but remain similar to those of SDF. Second, it is likely that some constraints we enforced to guarantee boundedness and liveness could be relaxed. A clear candidate is the condition prohibiting transformations to manipulate edges with initial tokens. Finally, the pattern and transformation language might be enriched. Consider the problem of duplicating a line of actors between a split and a join whereas this line may change overtime (by adding/removing actors). This would require to have as many transformation rules as there are versions of the line. Allowing to match and duplicate (unambiguously) some part of the graph without enumerating all its actors would be useful here.

REFERENCES

- [1] Shuvra S. Bhattacharyya, Edward A. Lee, and Praveen K. Murthy. 1996. *Software Synthesis from Dataflow Graphs*. Kluwer Academic Publishers, Norwell, MA, USA.
- [2] V. Bebelis, P. Fradet, and A. Girault. 2014. A Framework to Schedule Parametric Dataflow Applications on Many-Core Platforms. In *International Conference on Languages, Compilers and Tools for Embedded Systems, LCTES'14*. ACM, Edinburgh, UK.
- [3] Vagelis Bebelis, Pascal Fradet, Alain Girault, and Bruno Lavigueur. 2013. BPDF: A Statically Analyzable Dataflow Model with Integer and Boolean Parameters. In *International Conference on Embedded Software, EMSOFT'13*. 1–10.
- [4] Bishnupriya Bhattacharyya and Shuvra S. Bhattacharyya. 2001. Parameterized Dataflow Modeling for DSP Systems. *IEEE Trans. on Signal Processing (TSP)* 49, 10 (2001), 2408–2421.
- [5] Bishnupriya Bhattacharyya and Shuvra S. Bhattacharyya. 2001. Parameterized Dataflow Modeling for DSP Systems. *Trans. Sig. Proc.* 49, 10 (2001), 2408–2421.
- [6] A. Bouakaz, P. Fradet, and A. Girault. 2016. Symbolic Buffer Sizing for Throughput-Optimal Scheduling of Dataflow Graphs. In *Real-Time and Embedded Technology and Applications Symposium, RTAS'16*. Vienna, Austria, 199–208.
- [7] A. Bouakaz, P. Fradet, and A. Girault. 2017. A Survey of Parametric Dataflow Models of Computation. *ACM Trans. on Design Automation of Electronic Systems (TODAES)* 22, 2, Article 38 (March 2017).
- [8] Adnan Bouakaz, Pascal Fradet, and Alain Girault. 2017. Symbolic Analyses of Dataflow Graphs. *ACM Trans. on Design Automation of Electronic Systems (TODAES)* 22, 2 (2017), 39.
- [9] J.T. Buck and E.A. Lee. 1993. Scheduling Dynamic Data-Flow Graphs with Bounded Memory Using the Token Flow Model. In *International Conference on Acoustics, Speech, and Signal Processing, ICASSP'93*, Vol. I. IEEE, Minneapolis (MN), USA, 429–432.
- [10] K. Desnos, M. Pelcat, J.-F. Nezan, S.S. Bhattacharyya, and S. Aridhi. 2013. PiMM: Parameterized and Interfaced Dataflow Meta-Model for MPSoCs runtime reconfiguration. In *International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation, SAMOS'13*. IEEE, Samos Island, Greece, 41–48.
- [11] Karol Desnos, Maxime Pelcat, Jean-François Nezan, Shuvra S Bhattacharyya, and Slaheddine Aridhi. 2013. PiMM: Parameterized and Interfaced dataflow Meta-Model for MPSoCs runtime reconfiguration. In *13th Int. Conf. on Embedded Computer Systems: Architecture, Modeling and Simulation*. 41 – 48.
- [12] Pascal Fradet, Alain Girault, Ruby Krishnaswamy, Xavier Nicollin, and Arash Shafiei. 2019. RDF: Reconfigurable Dataflow. In *DATE 2019 - Design, Automation & Test in Europe Conference & Exhibition*. Florence, Italy.
- [13] Pascal Fradet, Alain Girault, and Peter Poplavko. 2011. *SPDF: A Schedulable Parametric Data-Flow MoC (Extended Version)*. Research Report RR-7828. INRIA. 24 pages. <https://hal.inria.fr/hal-00666284>
- [14] Pascal Fradet, Alain Girault, and Peter Poplavko. 2012. SPDF: A Schedulable Parametric Data-flow MoC. In *Proceedings of the Conference on Design, Automation and Test in Europe*. 769–774.
- [15] M. R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.
- [16] Marc Geilen. 2010. Synchronous Dataflow Scenarios. *ACM Trans. on Embedded Computing Systems (TECS)* 10, 2 (2010), 16.
- [17] F. Jondral. 2005. Software-Defined Radio – Basics and Evolution to Cognitive Radio. *EURASIP Journal on Wireless Communications and Networking* 3 (2005), 275–283.
- [18] Gilles Kahn. 1974. The Semantics of a Simple Language for Parallel Programming. *Information Processing* 74 (1974), 471–475.

- [19] E.A. Lee, S. Neuendorffer, and G. Zhou. 2014. *System Design, Modeling, and Simulation using Ptolemy II*. Ptolemy.org.
- [20] Edward A. Lee and David G. Messerschmitt. 1987. Synchronous Data Flow. *Proc. IEEE* 75, 9 (1987), 1235–1245.
- [21] Jun Li, Xianzhong Dai, Zhengda Meng, and Libo Xu. 2008. Improved Net Rewriting Systems-extended Petri Nets Supporting Dynamic Changes. *Journal of Circuits, Systems, and Computers* 17, 6 (2008), 1027–1052.
- [22] O. Moreira, T. Basten, M. Geilen, and S. Stuijk. 2010. Buffer Sizing for Rate-Optimal Single-Rate Data-Flow Scheduling Revisited. *IEEE Trans. on Computer* 59, 2 (2010), 188–201.
- [23] Julia Padberg and Laid Kahloul. 2018. Overview of Reconfigurable Petri Nets. In *Graph Transformation, Specifications, and Nets - In Memory of Hartmut Ehrig*. 201–222.
- [24] Julia Padberg and Alexander Schulz. 2016. Model Checking Reconfigurable Petri Nets with Maude. In *Graph Transformation - 9th International Conference, ICGT*. 54–70.
- [25] Jean-Claude Raoult and Frédéric Voisin. 1994. Set-theoretic graph rewriting. In *Graph Transformations in Computer Science*. Springer, 312–325.
- [26] Arash Shafiei. 2021. *RDF: A Reconfigurable Dataflow Model of Computation*. Ph. D. Dissertation. Université Grenoble Alpes.
- [27] Mladen Skelin, Marc Geilen, Francky Catthoor, and Sverre Hendseth. 2015. Parametrized dataflow scenarios. In *Proceedings of the 12th International Conference on Embedded Software*. 95–104.
- [28] Maarten H. Wiggers, Marco J. G. Bekooij, and Gerard J. M. Smit. 2008. Buffer Capacity Computation for Throughput Constrained Streaming Applications with Data-Dependent Inter-Task Communication. In *Proc. of the 2008 IEEE Real-Time and Embedded Technology and Applications Symposium*. 183–194.

A APPENDIX

We first recall the following facts and notations:

- In this work a graph is seen as a set of (directed) edges. We say that an actor α_x belongs to graph G (and write $\alpha_x \in G$) if there is an edge in G having α_x as initial or terminal vertex.
- A transformation rule $tr : lhs \Rightarrow rhs$ is a set rewriting where both lhs and rhs are non empty. Applying it to a graph G consists in finding a substitution σ such that $G = X \cup \sigma(lhs)$. The graph is then rewritten into $tr(G) = X \cup \sigma(rhs)$. Here, the set of edges of G is partitioned into the two disjoint sets of edges X and $\sigma(lhs)$. X and $\sigma(lhs)$ may have some actors in common because any node connected to one or more edges of X and to one of more edges of $\sigma(lhs)$ appears both in X and in $\sigma(lhs)$.
- We write $\alpha_x \xrightarrow[G]{+} \beta_y$ for a directed edge from α_x to β_y belonging to graph G (set of edges) and use the corresponding transitive closure $\alpha_x \xrightarrow[G]{+} \beta_y$ (resp. reflexive transitive closure $\alpha_x \xrightarrow[G]{*} \beta_y$) to denote paths in G . We write $\alpha_x \xleftrightarrow[G]{+} \beta_y$ to denote that there is an edge from α_x to β_y or from β_y to α_x in graph G . We use the corresponding transitive closure $\alpha_x \xleftrightarrow[G]{+} \beta_y$ (resp. reflexive transitive closure $\alpha_x \xleftrightarrow[G]{*} \beta_y$) to denote an undirected path between α_x and β_y in G .
- For any actor α_x of a consistent graph G , $sol_G(\alpha_x)$ denotes the number of firings of α_x in the minimal iteration of G .

A.1 Proof of Theorem 1 (connectivity)

THEOREM 1. *Let G be a weakly connected graph and $tr : lhs \Rightarrow rhs$ be a transformation rule such that*

$$\forall x \neq y \in rhs, x \xleftrightarrow[rhs]{+} y, \quad (\mathbf{C}^{conn})$$

then $tr(G)$ is a weakly connected graph.

PROOF. Let x and y be two distinct actors in $tr(G)$; we must prove that $x \xleftrightarrow[tr(G)]{+} y$. As said above, we consider tr as the set rewriting $G = X \cup \sigma(lhs) \Rightarrow X \cup \sigma(rhs) = tr(G)$. Note that Cond. (\mathbf{C}^{conn}) implies that for all x, y in $\sigma(rhs)$, we have $x \xleftrightarrow[\sigma(rhs)]{+} y$.

We distinguish the following exclusive cases: x and y are in $\sigma(rhs)$; neither x nor y are in $\sigma(rhs)$; x is in $\sigma(rhs)$ whereas y is not; the last case ($y \in \sigma(rhs)$ and $x \notin \sigma(rhs)$) is identical to the previous one.

(Case 1): $x \in \sigma(rhs)$ and $y \in \sigma(rhs)$.

We have already stated $x \xleftrightarrow[\sigma(rhs)]{+} y$ for any two distinct actors x and y of rhs . Since $\sigma(rhs) \subseteq tr(G)$, we therefore conclude that $x \xleftrightarrow[tr(G)]{+} y$.

(Case 2): $x \notin \sigma(rhs)$ and $y \notin \sigma(rhs)$.

Recall that an actor belonging to lhs but not to rhs is removed from the graph. Since neither x nor y are removed by tr (by assumption they are in $tr(G)$), necessarily none of them belong to $\sigma(lhs)$. It follows that they both belong to X , and therefore to G . Since G is weakly connected, we thus have $x \xleftrightarrow[G]{+} y$.

Since both x and y belong to X , this undirected path between x and y in G must start and finish with an edge in X , meaning that it consists of an alternation of subpaths in X and subpaths in $\sigma(lhs)$, starting and ending with a subpath in X . Formally, this path must have one of the two following forms:

- $x \xrightarrow{X^+} x_1 \xrightarrow{\sigma(lhs)^+} x_2 \xrightarrow{X^+} x_3 \xrightarrow{\sigma(lhs)^+} \dots \xrightarrow{\sigma(lhs)^+} x_n \xrightarrow{X^+} y$ with $n \geq 1$;
- or $x \xrightarrow{X^+} y$.

In the second case, trivially $x \xrightarrow{tr(G)^+} y$ since $X \subseteq tr(G)$. We therefore focus on the first case.

For each $1 \leq i \leq n$, x_i belongs to two edges, one in X and one in $\sigma(lhs)$, hence x_i belongs to X and cannot be suppressed by tr (thanks to Cond. (C2)). It follows that, for each $1 \leq i \leq n$, x_i belongs to $\sigma(rhs)$.

Now, by Cond. (C^{conn}), $\sigma(rhs)$ is weakly connected, hence we have $x_1 \xrightarrow{\sigma(rhs)^+} x_n$. Furthermore, edges in X being untouched by tr , we thus have $x \xrightarrow{X^+} x_1 \xrightarrow{\sigma(rhs)^+} x_n \xrightarrow{X^+} y$. Since $tr(G) = X \cup \sigma(rhs)$, we therefore have $x \xrightarrow{tr(G)^+} y$.

(Case 3): $x \in \sigma(rhs)$ and $y \notin \sigma(rhs)$.

As in (Case 2), y belongs to X hence to G and does not belong to $\sigma(lhs)$. However, x does not necessarily belong to $\sigma(lhs)$. We consider both cases in turn.

(Case 3.1): $x \in \sigma(lhs)$.

Since y belongs to the weakly connected graph G , we have $x \xrightarrow{G^+} y$. Similarly to (Case 2), this path consists of an alternation of subpaths in X and subpaths in $\sigma(lhs)$, starting with a subpath in $\sigma(lhs)$ and ending with a subpath in X :

- $x \xrightarrow{\sigma(lhs)^+} x_1 \xrightarrow{X^+} x_2 \xrightarrow{\sigma(lhs)^+} \dots \xrightarrow{\sigma(lhs)^+} x_n \xrightarrow{X^+} y$ with $n \geq 1$

On the one hand, since x_n belongs to X and to $\sigma(lhs)$, it also belongs to $\sigma(rhs)$ (same reasoning as in (Case 2)). Furthermore, by hypothesis x also belongs to $\sigma(rhs)$. Therefore, by Cond. (C^{conn}), $x \xrightarrow{\sigma(rhs)^+} x_n$, hence $x \xrightarrow{tr(G)^+} x_n$.

On the other hand, edges in X such as $x_n \xrightarrow{X^+} y$ being untouched by tr , we have $x_n \xrightarrow{tr(G)^+} y$. Putting both facts together, we conclude that $x \xrightarrow{tr(G)^+} y$.

(Case 3.2): $x \notin \sigma(lhs)$.

In that case, x is a fresh actor created by tr .

Since, by definition $lhs \neq \emptyset$ (see Section 3.2.1), there is an actor $z \in \sigma(lhs)$ such that $z \xrightarrow{G^+} y$ (because G is weakly connected). This path is an alternation of subpaths in X and subpaths in $\sigma(lhs)$, starting with a subpath in $\sigma(lhs)$ and ending with a subpath in X :

- $z \xrightarrow{\sigma(lhs)^+} x_1 \xrightarrow{X^+} x_2 \xrightarrow{\sigma(lhs)^+} \dots \xrightarrow{\sigma(lhs)^+} x_n \xrightarrow{X^+} y$ with $n \geq 1$

Necessarily x_n is not touched by tr (because it belongs to X and $\sigma(lhs)$ in G), hence x_n belongs both to X and to $\sigma(rhs)$ in $tr(G)$. This implies two things. On the one hand we have $x_n \xrightarrow{X^+} y$, implying $x_n \xrightarrow{tr(G)^+} y$. On the other hand, $\sigma(rhs)$ being connected, thanks to Cond. (C^{conn}), we thus have $x \xrightarrow{\sigma(rhs)^+} x_n$, implying $x \xrightarrow{tr(G)^+} x_n$. By transitivity, we conclude that $x \xrightarrow{tr(G)^+} y$. \square

A.2 Proof of Theorem 2 (consistency)

THEOREM 2. *Let G be a consistent graph and let $tr : lhs \Rightarrow rhs$ be a transformation rule such that lhs and rhs are consistent and*

$$\forall \alpha_x \in lhs \cap rhs, sol_{lhs}(\alpha_x) = sol_{rhs}(\alpha_x), \quad (\text{C}^{sol})$$

then $tr(G)$ is consistent.

Recall that $sol_{pat}(\alpha_x)$ denotes the *minimal* solution of actor α_x in the system of equations corresponding to the pattern graph pat . If pat is a pure SDF graph, then this minimal solution is an integer. In contrast, if pat is a pattern (with possibly symbolic rates), then this solution can also be computed but it is, in general, *symbolic*. If a graph (or pattern) has actors of variable arity, it also contains actors with parametric solutions. The reader may refer to [13] for a definition of the minimal symbolic solutions of parametric systems of balance equations.

PROOF. First, consider a graph G (a set of edges) that can be partitioned into two disjoint subsets of edges (two subgraphs) G_1 and G_2 , *i.e.*, $G = G_1 \cup G_2$ with $G_1 \cap G_2 = \emptyset$. The system of balance equations of G is the union of the two systems of balance equations of G_1 and G_2 . If G is consistent (*i.e.*, its system of balance equation has a non-null solution), then clearly G_1 and G_2 are also consistent. Moreover, for any actor α_x of G , $sol_G(\alpha_x)$ is also a solution of α_x for the systems of balance equations of both G_1 and G_2 . This solution may be not minimal for the system of balance equations of G_1 because G may enforce additional constraints (and similarly for G_2), but we have:

$$\exists!k_1 \in \mathbb{N}^*, \forall \alpha_x \in G_1, sol_G(\alpha_x) = k_1 sol_{G_1}(\alpha_x), \quad (5)$$

and similarly for G_2 . Recall that, since $sol_G(\alpha_x)$ is the number of firings in the *minimal* iteration (and similarly for G_1 and G_2), both k_1 and k_2 are *unique*.

Dually, if G_1 and G_2 are consistent and if there exist two non-null integers k_1 and k_2 such that, for any actor α_x belonging to both G_1 and G_2 , we have $k_1 sol_{G_1}(\alpha_x) = k_2 sol_{G_2}(\alpha_x)$, then G is also consistent. Indeed, $k_1 sol_{G_1}(\alpha_x)$ is also a solution for α_x in the system of balance equations of G . Furthermore, the minimal (and necessarily coprime) pair of integers k_1 and k_2 gives the minimal solutions for G . Lemma 1 formalizes this fact.

LEMMA 1. *Let G be an SDF graph partitioned into G_1 and G_2 . We have:*

$$G \text{ is consistent} \iff \begin{cases} G_1 \text{ is consistent} \\ \wedge G_2 \text{ is consistent} \\ \wedge \exists!(k_1, k_2) \in \mathbb{N}^* \times \mathbb{N}^*, \forall \alpha_x \in G_1 \cap G_2, k_1 sol_{G_1}(\alpha_x) = k_2 sol_{G_2}(\alpha_x) = sol_G(\alpha_x) \end{cases} \quad (6)$$

Now, let G be a consistent RDF graph, and let tr be a transformation rule of the form

$$\underbrace{X \cup \sigma(lhs)}_G \Rightarrow \underbrace{X \cup \sigma(rhs)}_{tr(G)}$$

and satisfying the conditions of Theorem 2, namely lhs is consistent, rhs is consistent, and Cond. (C^{sol}) holds. Thanks to Lemma 1, to show that $tr(G)$ is consistent, it suffices to show that:

- (a) X is consistent;
- (b) $\sigma(rhs)$ is consistent;
- (c) and $\exists!(k_1, k_2) \in \mathbb{N}^* \times \mathbb{N}^*$ such that $\forall \alpha_x \in X \cap \sigma(rhs)$, we have $k_1 sol_X(\alpha_x) = k_2 sol_{\sigma(rhs)}(\alpha_x) = sol_{tr(G)}(\alpha_x)$.

Since G is consistent, by Lemma 1 we know that X is consistent too; hence item (a). Since rhs is consistent and since the only effect of the substitution σ on the system of balance equations is to make some rate variables concrete, we know that $\sigma(rhs)$ is consistent too; hence item (b).

To prove item (c), let x in $X \cap \sigma(rhs)$, since α_x belongs to X and to $\sigma(rhs)$, it is neither a created nor a suppressed actor, but instead a preserved actor that belongs to $\sigma(lhs)$, hence to $X \cap \sigma(lhs)$. Then, since G is consistent, by Lemma 1

we know that there exist two unique non-null integers k_1 and k_2 such that, for any actor α_x in $X \cap \sigma(lhs)$, we have:

$$k_1 \text{sol}_X(\alpha_x) = k_2 \text{sol}_{\sigma(lhs)}(\alpha_x) = \text{sol}_G(\alpha_x). \quad (7)$$

Moreover, thanks to Cond. (\mathbf{C}^{sol}), we know that for any α_x in $lhs \cap rhs$, we have $\text{sol}_{lhs}(\alpha_x) = \text{sol}_{rhs}(\alpha_x)$, in the symbolic domain. Applying the substitution σ does not change this equality, it only “shifts” it to the concrete domain. In other words, for any α_x in $\sigma(lhs) \cap \sigma(rhs)$, we have $\text{sol}_{\sigma(lhs)}(\alpha_x) = \text{sol}_{\sigma(rhs)}(\alpha_x)$.

Putting these two facts together, it follows that the two non null integers k_1 and k_2 from Eq. (7) are such that, for any actor α_x in $X \cap \sigma(rhs)$, we have:

$$k_1 \text{sol}_X(\alpha_x) = k_2 \text{sol}_{\sigma(rhs)}(\alpha_x) = \text{sol}_{tr(G)}(\alpha_x), \quad (8)$$

hence item (c). This concludes the proof. \square

The proof holds for variable arity actors too. The condition and reasoning deal with symbolic solutions that can accommodate parameters of actors in X . Indeed, Lemma 1 can be extended to RDF graphs by replacing the two non-null integers k_1 and k_2 by two symbolic expressions (still denoted k_1 and k_2) over the set of parametric rates \mathcal{R} used in the variable arity actors. More precisely, each symbolic expression k_i is actually a product of a constant factor in \mathbb{N}^* and several (possibly zero) rate variables from \mathcal{R} . Such expressions therefore admit a canonical form as a pair from $\mathbb{N}^* \times 2^{\mathcal{R}}$, and equality in Lemma 1 is the syntactic equality of the symbolic expressions. Besides, k_1 and k_2 are such that (i) their respective constant factors are coprime, and (ii) they have no common rate variable.

A.3 Proof of Theorem 3 (liveness)

THEOREM 3. *Let G be a live graph with an SAS and $tr : lhs \Rightarrow rhs$ be a transformation rule such that rhs is live and*

$$\forall x, y \in lhs \cap rhs, x \xrightarrow{rhs}^+ y \Rightarrow x \xrightarrow{lhs}^+ y, \quad (\mathbf{C}^{live})$$

then $tr(G)$ is live and admits an SAS.

PROOF. It is well known that any consistent acyclic SDF graph has a single appearance schedule [1]. We therefore focus on cycles and first prove the following lemma, which states that a transformation respecting Cond. (\mathbf{C}^{live}) cannot create new cycles.

LEMMA 2. *Let $tr : lhs \Rightarrow rhs$ a transformation rule satisfying Cond. (\mathbf{C}^{live}), then*

$$\forall G \text{ and } \forall x \in G, x \xrightarrow{tr(G)}^+ x \Rightarrow x \xrightarrow{G}^+ x \quad (9)$$

PROOF. Consider the rule $tr : G = X \cup \sigma(lhs) \Rightarrow X \cup \sigma(rhs) = tr(G)$. There are two exclusive cases depending on whether or not x belongs to X :

Case $x \in X$. The path $x \xrightarrow[tr(G)]{+} x$ is made of alternating subpaths from X and $\sigma(rhs)$. It can take one of the following forms depending on whether the path starts and terminates with a subpath in X or in $\sigma(rhs)$:

$$\begin{array}{c}
x \xrightarrow[X]{+} x_1 \xrightarrow[\sigma(rhs)]{+} x_2 \xrightarrow[X]{+} \cdots \xrightarrow[\sigma(rhs)]{+} x_n \xrightarrow[X]{+} x \\
x \xrightarrow[X]{+} x_1 \xrightarrow[\sigma(rhs)]{+} x_2 \xrightarrow[X]{+} \cdots \xrightarrow[X]{+} x_n \xrightarrow[\sigma(rhs)]{+} x \\
x \xrightarrow[\sigma(rhs)]{+} x_1 \xrightarrow[X]{+} x_2 \xrightarrow[\sigma(rhs)]{+} \cdots \xrightarrow[\sigma(rhs)]{+} x_n \xrightarrow[X]{+} x \\
x \xrightarrow[\sigma(rhs)]{+} x_1 \xrightarrow[X]{+} x_2 \xrightarrow[\sigma(rhs)]{+} \cdots \xrightarrow[X]{+} x_n \xrightarrow[\sigma(rhs)]{+} x
\end{array}$$

Actors x, x_1, \dots, x_n belong to X : $x \in X$ by hypothesis and each x_i is either the initial or terminal vertex of an edge in X . Subpaths in X , $x_i \xrightarrow[X]{+} x_j$, are unchanged by tr and therefore occur also in G . For subpaths in $\sigma(rhs)$, $x_i \xrightarrow[\sigma(rhs)]{+} x_j$, we know that $x_i \in X$ and $x_j \in X$. Recall that an actor in $\sigma(rhs)$ is either a fresh actor created by tr , or belongs also to $\sigma(lhs)$. Since $x_i \in X$ and $x_j \in X$, then x_i and x_j must also belong $\sigma(lhs)$. In that case, Cond. (\mathbf{C}^{live}) enforces that the path $x_i \xrightarrow[\sigma(lhs)]{+} x_j$ exists. Therefore, in each of the above cases, we have $x \xrightarrow[G]{+} x$.

Case $x \notin X$. The path $x \xrightarrow[tr(G)]{+} x$ can take one of the two following forms:

$$\begin{array}{c}
x \xrightarrow[\sigma(rhs)]{+} x_1 \xrightarrow[X]{+} x_2 \xrightarrow[\sigma(rhs)]{+} \cdots \xrightarrow[X]{+} x_n \xrightarrow[\sigma(rhs)]{+} x \\
x \xrightarrow[\sigma(rhs)]{+} x
\end{array}$$

In the first case, we apply the same reasoning as before. All actors x_i (except x) belong to X and $x_1 \xrightarrow[G]{+} x_n$. We also have $x_n \xrightarrow[\sigma(rhs)]{+} x_1$ with $x_1 \in X$ and $x_n \in X$. Since x_1 and x_n also belong to $\sigma(lhs)$, Cond. (\mathbf{C}^{live}) ensures that $x_n \xrightarrow[\sigma(lhs)]{+} x_1$. It follows that $x \xrightarrow[G]{+} x$.

The second case is impossible. Indeed, Cond. (\mathbf{C}^{live}) enforces rhs to be live, and since tr can only manipulate edges without initial tokens, $\sigma(rhs)$ must be acyclic. □

We now return to the proof of Theorem 3. A consistent SDF graph admits an SAS (or a flat SAS following the terminology of [1]) if and only if all its cycles have a *saturated* edge, that is, an edge with enough initial tokens to permit its destination actor to complete all its firings in this SAS for one iteration.

Since the rule tr does not introduce new cycles (Lemma 2), nor removes (matches) any edge with initial tokens, nor changes the solution of actors (Theorem 2), all cycles remain with a saturated edge in $tr(G)$. We can therefore conclude that $tr(G)$ is live and admits an SAS. □