



**HAL**  
open science

# Real-time Mode-Aware Dataflow: A Dataflow Model to Specify and Analyze Mode-dependent CPSs under Relaxed Timing Constraints

Guillaume Roumage, Selma Azaiez, Cyril Faure, Stéphane Louise

► **To cite this version:**

Guillaume Roumage, Selma Azaiez, Cyril Faure, Stéphane Louise. Real-time Mode-Aware Dataflow: A Dataflow Model to Specify and Analyze Mode-dependent CPSs under Relaxed Timing Constraints. 2025. hal-04879612

**HAL Id: hal-04879612**

**<https://hal.science/hal-04879612v1>**

Preprint submitted on 10 Jan 2025

**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.

# Real-time Mode-Aware Dataflow: A Dataflow Model to Specify and Analyze Mode-dependent CPSs under Relaxed Timing Constraints

Guillaume Roumage<sup>†</sup>, Selma Azaiez<sup>‡</sup>, Cyril Faure<sup>‡</sup>, Stéphane Louise<sup>‡</sup>

<sup>†</sup>guillaume.roumage.research@proton.me

<sup>‡</sup>Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

<sup>‡</sup>firstname.lastname@cea.fr

**Abstract**—Modern Cyber-Physical Systems (CPS) often exhibit both relaxed real-time constraints and a mode-dependent execution. Relaxed real-time constraints mean that only a subset of the processes of a CPS have real-time constraints, and a mode-dependent CPS has conditional execution branches. Static analysis tools, such as the PolyGraph model (a formalism extending the Cyclo-Static Dataflow model with real-time constraints), can specify and analyze systems with relaxed real-time constraints. However, PolyGraph is limited in its ability to specify and analyze mode-dependent CPSs. This paper extends PolyGraph with routing actors, yielding the Routed PolyGraph model. This model is further extended to the Real-time Mode-Aware Dataflow (RMDF), which both leverages routing actors and incorporates a new dataflow actor to specify mode-dependent CPSs under relaxed real-time constraints. This paper also extends the static analyses of PolyGraph to RMDF. We showcase the application of RMDF with a specification and an analysis (derivation of timing constraints at the job-level and a feasibility test) of the vision processing system of the Ingenuity Mars helicopter.

**Index Terms**—Dataflow Model, Real-time Systems, Mode-Dependent Execution, Timing Analysis

## I. INTRODUCTION

Cyber-Physical Systems (CPSs) are reactive systems that detect environmental shifts through sensors, process this information using computational processes, and then use the output to control actuators. CPSs range from digital signal processing systems to embedded/cloud infrastructures, soft/hard real-time systems, and even a mix of all the above. These complex systems must operate reliably without threatening their internal processes. For example, a failure of the actuator in an autonomous car can lead to catastrophic consequences such as a car crash or a pedestrian accident.

An increasing number of systems exhibit dynamic behavior, allowing them to switch between various modes of operation at runtime, e.g., a switch between the nominal and the degraded execution mode. The procedure that initiates these mode switches may rely on runtime data, limiting the static analyzability of the system. This paper presents an extension of an existing static analysis tool to specify and analyze such systems with dynamic behavior.

The design tools of interest in this paper are Dataflow Models of Computation and Communication (DF MoCC). These tools can formally prove the correctness of some classes of CPSs regarding safety and temporal properties, such as

consistency, liveness, latency, and throughput. This paper presents a DF MoCC called Real-time Mode-aware Dataflow (RMDF), which extends the *routed PolyGraph model*, this latter being itself an extension of the DF MoCC PolyGraph [1]. While PolyGraph is tailored to specify and analyze CPSs with *relaxed* real-time constraints, that is, a CPS with real-time constraints on only a subset of its processes, it cannot specify and analyze *mode-dependent* CPSs.

1) *Contribution*: We have developed an extension of PolyGraph, called RMDF, to specify and analyze *mode-dependent* CPSs under relaxed real-time constraints. As a dataflow model is a trade-off between its expressiveness and its analyzability [2], the challenge here is to enhance the former without reducing the latter. Within the scope of this paper, a *mode-dependent* CPS has a set of conditional execution branches. The consistency analysis (checking the existence of a memory-bounded execution) and the liveness analysis (checking the existence of a deadlock-free execution) of PolyGraph are extended to RMDF. In addition, the timing analysis described in [3] for PolyGraph is extended to RMDF, and a feasibility test is proposed.

2) *Paper organization*: The paper starts by presenting the background and the terminology of the PolyGraph model in Section II. The PolyGraph model is extended to the *routed PolyGraph model* in Section III, which is further extended to the RMDF model in Section IV. The extension of the static analysis of PolyGraph to RMDF is detailed in Section V. Section VI presents the timing analysis of an RMDF specification. This timing analysis permits a feasibility test presented in Section VII. Section VIII applies the RMDF model to specify and analyze a dataflow specification of the Ingenuity vision processing system. Section IX discusses the mode change protocol of RMDF. Finally, Section X presents discussions and related works, and Section XI concludes the paper.

## II. BACKGROUND AND TERMINOLOGY

RMDF is an extension of PolyGraph [4], which is itself a superset of the Synchronous Dataflow (SDF) model [5]. In this section, we briefly present SDF, and we explain how the PolyGraph model expands it. We also introduce *disjoint directed meshes* to explain further how RMDF extends PolyGraph.

### A. The SDF Model

An SDF specification is an oriented graph  $G = (V, E)$  where  $V$  is a finite set of *actors* and  $E$  is a finite set of *channels*. An SDF specification is characterized by a *topology matrix*  $G_\Gamma$  of size  $|E| \times |V|$ . An *actor* is a computational unit that both produces and consumes data tokens every time it is executed. The atomic amount of data exchanged is known as a *token*. The entry  $(i, j)$  of  $G_\Gamma$  is the number of tokens produced or consumed by the actor  $v_j$  on the channel  $c_i$  each time the actor  $v_j$  is executed (this number is positive if the tokens are produced, and negative otherwise). An execution of an actor is called a *job*.

An actor  $v \in V$  has (optional) input and output ports connected to input and output channels of  $E$ . A channel  $c_i = (v_j, v_k, n_{ij}, n_{ik}, init_i) \in E$  connects an output port of the actor  $v_j \in V$  to an input port of the actor  $v_k \in V$ . This channel also has a production rate  $n_{ij} \in \mathbb{N}^*$  (which is the entry  $(i, j)$  of  $G_\Gamma$ ), a consumption rate  $n_{ik} \in \mathbb{N}^*$  (the entry  $(i, k)$  of  $G_\Gamma$ ), and a number of initial tokens  $init_i \in \mathbb{N}$ . An actor's job produces/consumes tokens to/from the buffer of its output/input channels according to the production/consumption rate. For the sake of simplicity in the rest of this paper, we denote  $[c_i]$  the number of initial tokens of the channel  $c_i$ .

### B. The PolyGraph Model

1) *Syntax and semantics*: The PolyGraph model is a superset of the SDF model. It expands it in two ways: the production and consumption rates become periodic sequences (as in CSDF [6]), and a subset of actors may have timing constraints. The terminology of the static analysis of an SDF specification, i.e., consistency, liveness, consistent and live execution, and hyperperiod [5], is extended to a PolyGraph specification. Consistency ensures the system can be executed within a bounded memory, while liveness guarantees deadlock-free execution. A hyperperiod is a partially ordered set of actors' jobs that returns the system to its initial state within a given time frame.

2) *Rational production and consumption rates*: Let  $G = (V, E)$  be a PolyGraph specification and let  $G_\Gamma = (\gamma_{ij}) \in \mathbb{Q}^{|E| \times |V|}$  be its topology matrix. The production and consumption rates are rational<sup>1</sup>: a channel  $c_i \in E$  is a tuple  $(v_j, v_k, \gamma_{ij}, \gamma_{ik}, [c_i])$  such that  $\gamma_{ij} \in \mathbb{Q}^*$  and  $\gamma_{ik} \in \mathbb{Q}^*$ , and  $[c_i] \in \mathbb{Q}$  ( $\mathbb{Q}^* = \mathbb{Q} \setminus \{0\}$ ). Rational rates imply that the number of tokens produced/consumed can differ for each job. Specifically, a token is produced and stored in a channel if and only if sufficient fractional token parts are symbolically produced. Indeed, only an integer number of tokens may be produced or consumed. In other words, an actor with a rational production rate such as  $1/n$  in the dataflow specification actually produces data once every  $n$  jobs during runtime, thus validating the precedence constraint toward the consumer once every  $n$  of its jobs.

<sup>1</sup>A channel of a PolyGraph specification as defined in [4] has at least one integer rate. However, for the sake of simplicity, we consider in this paper that both rates are rational.

The rational production and consumption rates are natural consequences of the periodicity of real-time actors (with specified and imposed frequencies) together with the consistency of the system's model imposed by the topology matrix. The rational rates are a compact notation for a CSDF equivalency [6].

3) *Timing constraints*: The actors of a PolyGraph specification may have a frequency constraint. Furthermore, they may also have a phase if they have a frequency constraint. The phase usually models the system's end-to-end latency. The frequency dictates the actor's execution at the specified rate, while the phase postpones its initial execution. Actors with frequency constraints are referred to as *timed actors*. An execution of a consistent and live PolyGraph specification is an infinite repetition of its first hyperperiod. Although the number of jobs is not bounded during an execution, their timing constraints are cyclic over a hyperperiod.

### C. Disjoint Directed Meshes

Let  $G = (V, E)$  be a directed graph. Let us define two maps  $in, out : V \rightarrow 2^E$  that associate to each actor  $v \in V$  the set<sup>2</sup> of input and output channels, respectively.  $G$  is a *disjoint directed mesh* if:

- 1) there is a single source actor, i.e., there exists a unique actor  $v_s \in V$  such that  $in(v_s) = \emptyset$ ;
- 2) there is a single sink actor, i.e., there exists a unique actor  $v_t \in V$  such that  $out(v_t) = \emptyset$ ;
- 3) there exists at least one path from  $v_s$  to  $v_t$ , and a path is called a *branch*;
- 4) the intersection of all the branches is the empty set.

A disjoint directed mesh is illustrated in Fig. 1.




Fig. 1: An illustration of a disjoint directed mesh with three branches.  $n_1$  to  $n_3$  are the (bounded positive) number of actors in each branch.

## III. THE ROUTED POLYGRAPH MODEL: SYNTAX AND SEMANTICS

The RMDF model permits the specification of CPSs with conditional execution branches. This feature requires routing capabilities -specific actors to explicitly define how data tokens are routed into the conditional execution branches. We will introduce further in the paper such actors, called *controlled splitters* and *controlled joiners*. Controlled splitters and controlled joiners rely on actors called *splitters*, *joiners*, *duplicaters*, and *discard* [7], [8]. Those actors are not yet integrated into the PolyGraph model (only an implicit reference to them is used

<sup>2</sup> $2^E$  is the powerset of  $E$ , that the set of all subset of  $E$ . For instance, if  $E = \{c_1, c_2\}$ , then  $2^E = \{\emptyset, \{c_1\}, \{c_2\}, \{c_1, c_2\}\}$

in [1]). This section solves this issue by integrating these actors into the PolyGraph model, resulting in the *routed PolyGraph model*.

#### A. Routing Actors: Joiners, Splitter, Duplicaters, and Discards

The actors *Splitter*, *Joiner*, *Duplicater* and *Discard* were introduced in [7]. These actors are routing actors. Instead of engaging in computational tasks, they primarily render data communication patterns explicit. This explicitness allows for compile-time optimizations [7], especially reducing the memory requirements and improving overall CPS performance.

A *splitter* consumes tokens from a single input and distributes them to multiple outputs. A predefined number of tokens is sent to the first channel in the lexicographic order of the output channels, then to the second channel, and so on. The number of tokens sent to an output channel is the numerator of the production rate of that channel. We enforce that an execution of a splitter consumes one token at a time; the reason behind this enforcement will be detailed in the next paragraph. Hence, the splitter of Fig. 2a sends two tokens to  $c_1$  in two jobs and one token to  $c_2$  in one job. The process is repeated cyclically. Similarly, a *joiner* consumes tokens from multiple inputs and produces them to a single output. As for the splitter, the execution of a joiner produces one token at a time. Thus, the joiner of Fig. 2b receives two tokens from  $c_1$  in two jobs and one token from  $c_2$  in one job. The process is repeated cyclically. A *duplicater*, as illustrated in Fig. 2c, consumes tokens from a single input and sends them to multiple outputs. A *discard* (cf. Fig. 2d) is essentially a sink actor that consumes tokens from a single input without producing any. It is used to discard tokens that are no longer needed and is typically employed at the output of a splitter.

In order to ease the extension of static analysis from PolyGraph to the routed PolyGraph and then to RMDF, we enforce that an execution of a splitter or a joiner consumes or produces one token at a time. To that end, the production rate of the input/output channel of a splitter/joiner is equal to 1. In addition, the sum of the production/consumption rates of the output/input channels of a splitter/joiner is equal to 1. For instance, the production rates of the splitter of Fig. 2a are  $2/3$  and  $1/3$ , whose sum equals 1.

#### B. Static Analysis of a Routed PolyGraph Specification

There is an equivalence between a PolyGraph specification and a routed PolyGraph specification. Algorithm 3 details the transformation of a routed PolyGraph specification to a semantically equivalent PolyGraph specification.

**Definition 1** (Floor and ceil functions). *The floor function  $\lfloor x \rfloor$  returns the greatest integer less than or equal to  $x$ . The ceil function  $\lceil x \rceil$  returns the smallest integer greater than or equal to  $x$ .*

To illustrate, let  $G = (V, E)$  be a routed PolyGraph specification with a splitter  $v \in V$  that has  $n$  output channels  $c_1, \dots, c_n$  with the numerator of their production rate denoted as  $\gamma_1, \dots, \gamma_n$ . The splitter  $v$  is removed by modifying the

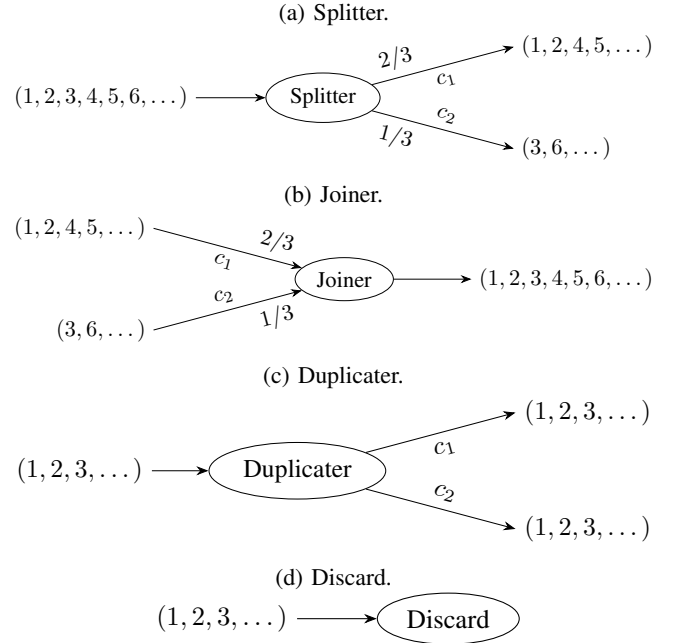


Fig. 2: Illustration of a splitter, a joiner, a duplicater, and a discard with their input and output traces. Rates of 1 and initial tokens of 0 are omitted for clarity.

producer actor of the channel  $c_i$  from  $v$  to the predecessor of  $v$ . At this point, the predecessor of  $v$  may send more than one token at some executions and no tokens at others. In order to reproduce the expected behavior of the splitter, one token has to be sent each time the predecessor of  $v$  executes. This is done by shifting the production rate with initial tokens; a challenge here is to find the correct amount of initial tokens.

There is a bijection from the production rate and the initial tokens of a channel to the sequence of tokens sent to this channel. This bijection can be constructed by exhaustively enumerating all sequences generated by a production rate and all values of initial tokens (cf. Algorithm 1). For instance, the production sequence  $[1, 1, 0]$  is associated with the production rate  $2/3$ . The question is “how many initial tokens are needed to generate this production sequence?”. The production sequence  $[1, 1, 0]$  is associated with the production rate  $2/3$ , so possible initial tokens are  $[2/3, 1/3, 0]$ . The production rate  $2/3$  with no initial token generates the production sequence  $[0, 1, 1]$ , the production rate  $2/3$  with  $1/3$  initial token generates the production sequence  $[1, 0, 1]$ , and the production rate  $2/3$  with  $2/3$  initial token generates the production sequence  $[1, 1, 0]$ . We can see that the production sequence  $[1, 1, 0]$  is associated with  $2/3$  initial tokens.

The production sequence can be easily computed (cf. Line 5 to Line 9 of Algorithm 3). By following the procedure described in Algorithm 1, the production rate and the required initial tokens can be derived. Initial tokens are associated to all output channels of a splitter with this approach before removing the splitter. As a consequence, the behavior of the splitter is preserved while removing it.

The same logic applies to removing a joiner using Algorithm 2. A discard is removed while removing the splitter, as a discard is always an output of a splitter. Removing a duplicater is trivial. Fig. 3 illustrates the transformation of a routed PolyGraph specification to a semantically equivalent PolyGraph specification by removing a splitter (cf. Fig. 3a), a joiner (cf. Fig. 3b), a duplicater (cf. Fig. 3c), and a discard (cf. Fig. 3d).

---

**Algorithm 1:** Compute the rate and the number of initial tokens that generate a production sequence.

---

**Input:** A production sequence  $s$  of length  $n$ .

**Output:** The production rate and the number of initial tokens generating the input production sequence.

```

1  $prod\_rate \leftarrow \sum_{i=1}^n s.get(i)/n$  ;
2  $init\_tkn \leftarrow 0$  ;
3 while  $init\_tkn \neq 1$ :
4    $prod\_seq \leftarrow []$  ;
5   for  $j \leftarrow 1$  to  $n$ :
6      $prod\_seq.get(i) \leftarrow$ 
7        $[n \cdot \gamma + init\_tkn] - [(n-1) \cdot \gamma + init\_tkn]$  ;
8   if  $s = seq$  then
9     Result:  $prod\_rate, init\_tkn$ 
10   $init\_tkn \leftarrow init\_tkn + 1/n$  ;
```

---



---

**Algorithm 2:** Compute the rate and the number of initial tokens that generate a consumption sequence.

---

**Input:** A consumption sequence  $s$  of length  $n$ .

**Output:** The consumption rate and the number of initial tokens generating the input consumption sequence.

```

1  $cons\_rate \leftarrow \sum_{i=1}^n s.get(i)/n$  ;
2  $init\_tkn \leftarrow 0$  ;
3 while  $init\_tkn \neq 1$ :
4    $cons\_seq \leftarrow []$  ;
5   for  $j \leftarrow 1$  to  $n$ :
6      $cons\_seq.get(i) \leftarrow$ 
7        $[n \cdot \gamma - init\_tkn] - [(n-1) \cdot \gamma - init\_tkn]$  ;
8   if  $s = cons\_seq$  then
9     Result:  $cons\_rate, init\_tkn$ 
10   $init\_tkn \leftarrow init\_tkn + 1/n$  ;
```

---

The resulting PolyGraph specification is less explicit as there are no routing actors. However, the static analyses of PolyGraph apply to the transformed PolyGraph specification.

#### IV. RMDF MODEL: SYNTAX AND SEMANTICS

This section explains how splitters and joiners are extended to the controlled splitters and controlled joiners. Those two latter actors will be used to route data tokens into conditional execution branches in an RMDF specification.

---

**Algorithm 3:** Removal of routing actors from a routed Polygraph specification.

---

**Input:** A routed Polygraph specification  $G = (V, E)$ .

**Output:** A semantically equivalent Polygraph specification without routing actors.

```

1 for  $v \in V$ :
2   if  $v.is\_splitter$  then
3     for  $c \in v.output\_channels$ :
4        $cons\_seq \leftarrow []$  ;
5       for  $c' \in v.output\_channels$ :
6         if  $c'.index < c.index$  or  $c'.index >$ 
7            $c.index$  then
8            $cons\_seq.append(0)$  ;
9         else
10         $cons\_seq.append(1)$  ;
11     $_, init\_tkn \leftarrow$ 
12       $Algorithm\ 2(cons\_seq, len(cons\_seq))$  ;
13     $E = E \cup$ 
14       $\{(v.pred, c.cons, c.prod\_rate, 1, [init\_tkn])\}$ ;
15
16  if  $v.is\_joiner$  then
17    for  $c \in v.input\_channels$ :
18       $prod\_seq \leftarrow []$  ;
19      for  $c' \in v.input\_channels$ :
20        if  $c'.index < c.index$  or  $c'.index >$ 
21           $c.index$  then
22           $prod\_seq.append(0)$  ;
23        else
24         $prod\_seq.append(1)$  ;
25     $_, init\_tkn \leftarrow$ 
26       $Algorithm\ 1(prod\_seq, len(prod\_seq))$  ;
27     $E = E \cup$ 
28       $\{(c.prod, v.succ, 1, c.cons\_rate, [init\_tkn])\}$ ;
29
30  if  $v.is\_duplicater$  then
31    for  $c \in v.output\_channels$ :
32       $E = E \cup \{(v.pred, c.cons, 1, 1, [0])\}$ ;
33
34 for  $v \in V$ :
35   if  $v.is\_splitter$  or  $v.is\_duplicater$  then
36     for  $c \in v.output\_channels$ :
37        $E = E \setminus \{c\}$ ;
38      $V = V \setminus \{v\}$ ;
39
40   if  $v.is\_joiner$  then
41     for  $c \in v.input\_channels$ :
42        $E = E \setminus \{c\}$ ;
43      $V = V \setminus \{v\}$ ;
44
45   if  $v.is\_discard$  then
46     for  $c \in v.input\_channels$ :
47        $E = E \setminus \{c\}$ ;
48      $V = V \setminus \{v\}$ ;
```

**Result:**  $G = (V, E)$

---

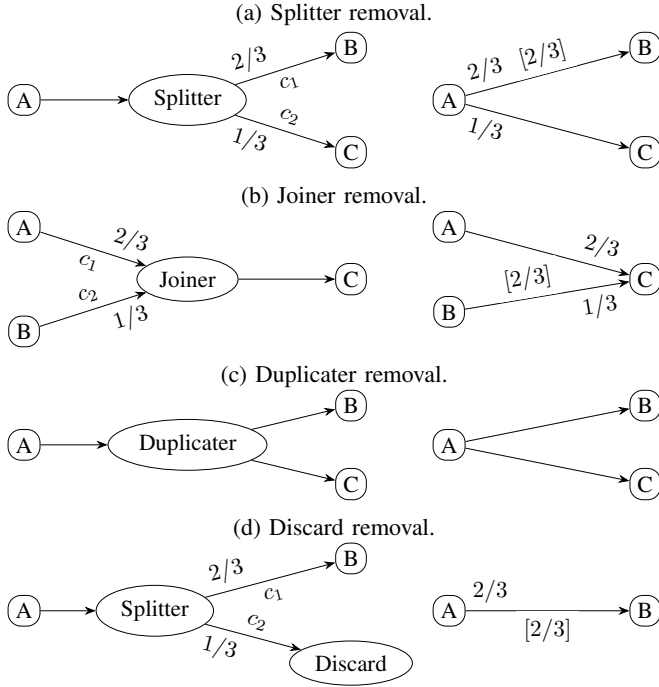


Fig. 3: Removal of routing actors while preserving the semantics of data communication pattern. Rates of 1 and initial tokens of 0 are omitted for clarity.

#### A. Controlled Routing Actors and Mode-Dependent Execution

A controlled splitter/joiner is a splitter/joiner with an additional input *control channel*. A *control channel* is where transit *control token*. A controlled splitter/joiner, which is illustrated in Fig. 4a and Fig. 4b, alternates between consuming a control token and consuming/producing a data token. The control token decides which channel the next data token will be sent/received, leading to a *mode-dependent* execution. In addition, the production/consumption rates of the output/input channels of a controlled splitter/joiner are parameters that take the value 0 or 1. For instance, parameters<sup>3</sup>  $m_1$  and  $m_2$  in Fig. 4a and Fig. 4b take the value 0 or 1. The actual value of  $m_1$  and  $m_2$  is determined at runtime by the value carried by control tokens consumed from  $c_2$  in Fig. 4a and Fig. 4b.

As for the non-controlled splitter/joiner, we want to ensure that executing a controlled splitter/joiner consumes/produces one token at a time. To that end, we enforce that a parameter assignment (i.e., transforming  $m_1$  and  $m_2$  into 0 or 1 at runtime) to adhere to the following constraints: exactly one parameter must be equal to 1, and the other parameters must be equal to 0. In Fig. 4, we have either  $m_1 = 1$  and  $m_2 = 0$  or  $m_1 = 0$  and  $m_2 = 1$ . We will detail later in the paper how this constraint is formally defined.

#### B. Mode Decider

As controlled splitters and joiners consume control tokens, we need an actor to produce them. A *mode decider* consumes

<sup>3</sup> $m_1$  and  $m_2$  are symbolic rates which are assigned a value at runtime

a data token from its input channels and produces a control token on its output control channel. The algorithm inside the mode decider that determines the value of the control token is usually a switch case, but it can be more complex depending on the application's needs. A mode decider has exactly one input data channel<sup>4</sup> (the channel  $c_1$  in Fig. 4c) and exactly one output control channel (the channel  $c_2$  in Fig. 4c), with this latter being connected to a duplicator. This duplicator is connected to one or many controlled splitters and controlled joiners. Fig. 4c illustrates a mode decider.

#### C. A Complete RMDF Specification

Fig. 5 illustrates an RMDF specification with three conditional execution branches ( $C_1$  to  $C_{n_1}$ ,  $D_1$  to  $D_{n_2}$  and  $E_1$  to  $E_{n_3}$  with  $n_1, n_2, n_3 \in \mathbb{N}^*$ ). An RMDF specification is a tuple  $(V, E, M)$  where  $V$  is a finite set of actors,  $E$  is a finite set of channels, and  $M$  is a set of functions that enforce the production and consumption of one token at a time for the controlled splitters and joiners. This set of functions of the RMDF specification of Fig. 5 is  $M = \{M_1, M_2, M_3\}$  as defined in the right bottom corner of Fig. 5. Each function specifies an execution mode. In Fig. 5, the actor B is a mode decider. The control token it produces is sent to a duplicator and then to a controlled splitter and a controlled joiner. The value carried by the control token (produced by B) will be used at runtime to determine the value of parameters  $m_1$ ,  $m_2$  and  $m_3$ .

### V. STATIC ANALYSIS OF AN RMDF SPECIFICATION

#### A. Control Areas

In order to ease the extension of the static analysis of the PolyGraph model to the RMDF model, we enforce that an execution of a controlled splitter/joiner consumes/produces one token at a time. In addition, we impose conditions on using the mode deciders, controlled splitters, and controlled joiners in an RMDF specification. Let us define *control areas* (Definition 2) on which restrictions (Property 1) will be applied. A control area is a disjoint directed mesh where the source and sink actors are a controlled splitter and controlled joiner.

**Definition 2** (Control area of a mode decider). *Let  $G = (V, E, M)$  be an RMDF specification, let  $v_{md} \in V$  be a mode decider; let  $dup(v_{md})$  be the duplicator connected to  $v_{md}$ , let  $cs(v_{md})$  be the controlled splitters connected to  $dup(v_{md})$  and let  $cj(v_{md})$  be the controlled joiners connected to  $dup(v_{md})$ . The control area of  $v_{md}$ , defined as  $control(v_{md})$ , is the set of actors conditioned by the execution mode output by  $v_{md}$ . It is the intersection of all the successors of controlled splitters belonging to  $cs(v_{md})$  and all the predecessors of the controlled joiners belonging to  $cj(v_{md})$ . Formally, we have  $control(v_{md}) = \{v \in V \mid \exists (v_1, \dots, v_i, \dots, v_n), v_1 \in cs(v_{md}), v_i = v, v_n \in cj(v_{md}), \forall 1 \leq j < n : v_{j+1} \in succ(v_j)\}$  with  $succ(v)$  the set of successors of  $v$ .*

<sup>4</sup>It would not be difficult to consider more than one input data channel for mode deciders.

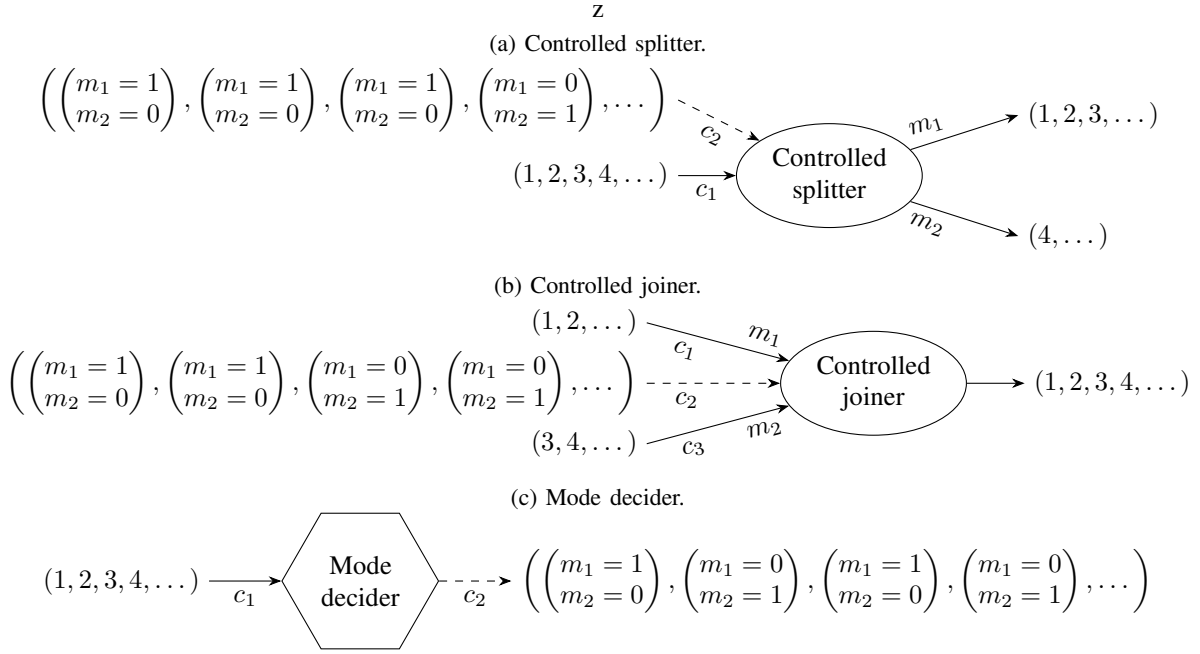


Fig. 4: An illustration of a controlled splitter and a controlled joiner with their input and output traces. Dashed lines represent control channels, and plain lines represent data channels. Rates of 1 and initial tokens of 0 are omitted for clarity.

For instance, in Fig. 5, the actor **B** is a mode decider. We have  $dup(B) = \{Duplicater\ 2\}$ ,  $cs(B) = \{Controlled\ splitter\}$  and  $cj(B) = \{Controlled\ joiner\}$ . The control area of the actor  $B$  is  $\{C_1, \dots, C_{n_1}, D_1, \dots, D_{n_2}, E_1, \dots, E_{n_3}\}$ . The following property imposes restrictions on the control areas of an RMDF specification.

**Property 1** (Restrictions on controls areas). *Let  $G = (V, E, M)$  be an RMDF specification:*

- 1) *Execution of actors in a control area must depend on exactly one execution mode. Algorithm 4 return **True** if this condition is fulfilled, **False** otherwise (see Fig. 6a for an RMDF specification that does not fulfill this requirement).*
- 2) *Let  $v_{md} \in V$  be a mode decider. There is no channel from an actor  $v_1 \in control(v_{md}) / v_1 \notin control(v_{md})$  to an actor  $v_2 \notin control(v_{md}) / v_2 \in control(v_{md})$  (cf. Fig. 6b).*
- 3) *If there are timed actors in a control area, they must have the same frequency (cf. Fig. 6c).*
- 4) *The production and consumption rate inside a control area and the parametric rates associated with the controlled joiner and controlled splitter of that control area must be equal to 1 or 0 (cf. Fig. 6d).*
- 5) *Let  $v_{md} \in V$  be a mode decider. When all parametric rates of a controlled splitter are assigned a value, their sum is equal to 1. The same condition holds for the parametric rates of a controlled joiner (cf. Fig. 6e).*

**Definition 3** (Mode-coherent RMDF specification). *Let  $G$  be*

---

**Algorithm 4:** Checking that actors' execution in control areas of an RMDF specification depends on a single execution mode. Some work remains to handle loops within a branch.

---

**Input:** An RMDF specification  $G = (V, E, M)$ .

**Output:** **True** if execution of the actors of  $G$  depends on a single mode, **False** otherwise.

```

1 def label_successor(v, associated_mode):
2   for c ∈ v.output_channels:
3     if mode ∉ c.conditioning_mode then
4       | c.conditioning_mode.append(mode) ;
5     if len(c.conditioning_mode) = 2 then
6       | Result: False
7     if c.consumer.is_controlled_joiner then
8       | Result: True
9     else
10      | label_successor(c.consumer, mode) ;
11 for v ∈ V:
12   if v.is_mode_decider then
13     for v ∈ cs(v):
14       mode = 0 ;
15       for c ∈ v.output_channels:
16         if
17           label_successors(c.consumer, mode) =
18             False then
19           | Result: False
20           mode ← mode + 1 ;
21 Result: True

```

---

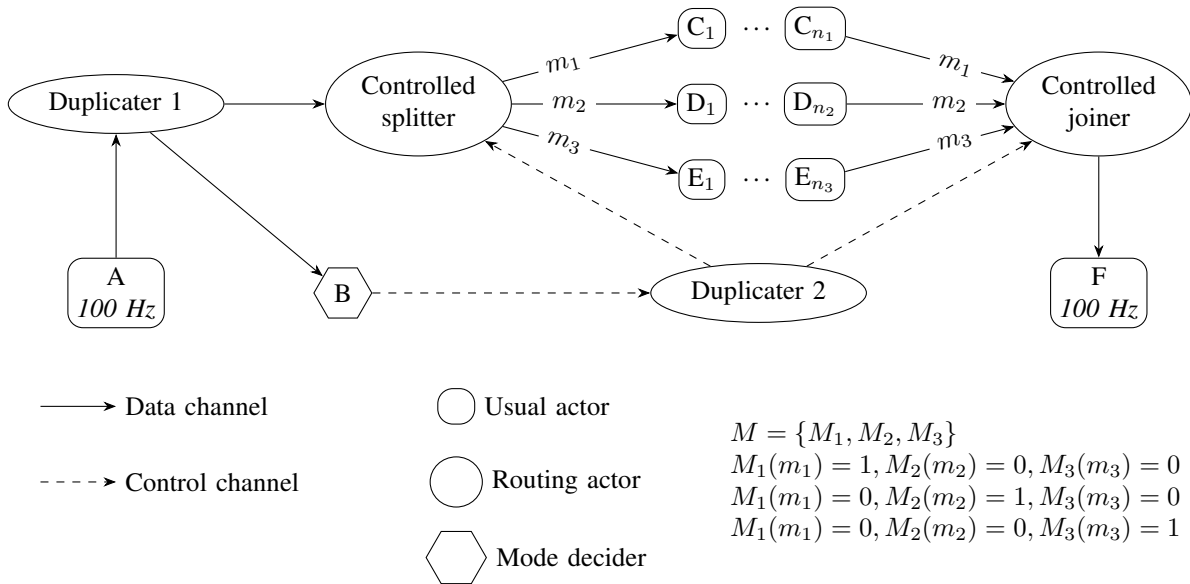


Fig. 5: An example of an RMDF specification. A usual actor is an actor which is neither (non-) controlled splitter or joiner, nor a mode decider.

an RMDF specification.  $G$  is mode-coherent if it satisfies the restrictions of Property 1.

B. Consistency and Liveness Analysis of an RMDF Specification

The consistency and liveness analysis of an RMDF specification is based on the static analysis of the exhaustive set of conditional execution branches of the RMDF specification. An RMDF specification with a single conditional execution branch can be analyzed as a routed PolyGraph specification. For instance, Fig. 7 presents the PolyGraph specifications obtained from an RMDF specification.

**Theorem 1** (Consistent and live RMDF specification). *Let  $G = (V, E, M)$  be a mode-coherent RMDF specification.  $G$  is consistent if the PolyGraph specification yield when parametric rates are assigned a value is consistent.  $G$  is consistent and live if the PolyGraph specifications yield when parametric rates are assigned a value is consistent and live.*

*Proof of Theorem 1.* Let  $G = (V, E, M)$  be a mode-coherent RMDF specification. A symbolic execution of  $G$  is a sequence of symbolic execution of PolyGraph specifications. Indeed, executions of mode deciders assign values to parameters, transforming an RMDF specification into a PolyGraph specification. If this latter is consistent and live,  $G$  returns to its initial state without deadlocking. Note that the mode-coherence property of  $G$  ensures that the execution of a control area is consistent and live. Once back to the initial state, the process repeats with another assignation of parameters, possibly yielding a different PolyGraph specification because another conditional execution branch can be taken.  $G$  also returns to its initial state. It remains to verify that all PolyGraph specifications are consistent and live. Such techniques are defined in [1].  $\square$

## VI. TIMING ANALYSIS OF AN RMDF SPECIFICATION

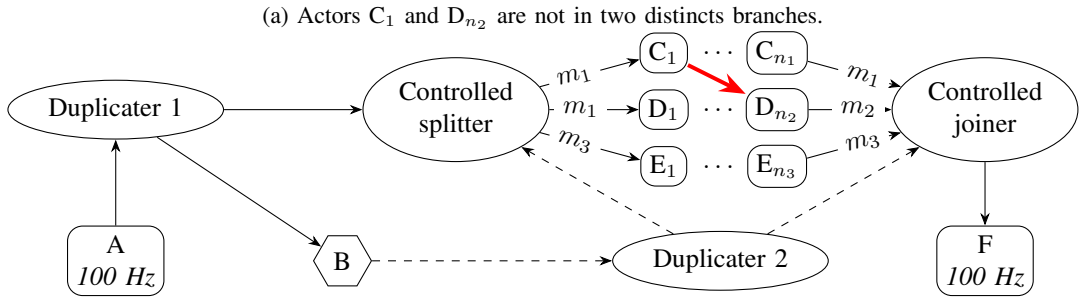
### A. Pre-processing of an RMDF Specification

The RMDF model, as it is based on the routed PolyGraph model, allows system designers to define an arbitrary number of initial tokens on each channel. This feature is helpful for modeling data in memory, such as the initial state of a system, which is already available when the system starts. Therefore, an actor may have all its input channels filled with initial tokens, meaning that it can start its execution immediately. Such actor can undergo “offline execution” and be *pre-processed* before the system starts without waiting for the first data produced by sensors. In other words, “offline jobs” may be executed to reduce the computational load of the system when it starts. It also eases the computation of the system’s timing constraints, as we will see further in this paper.

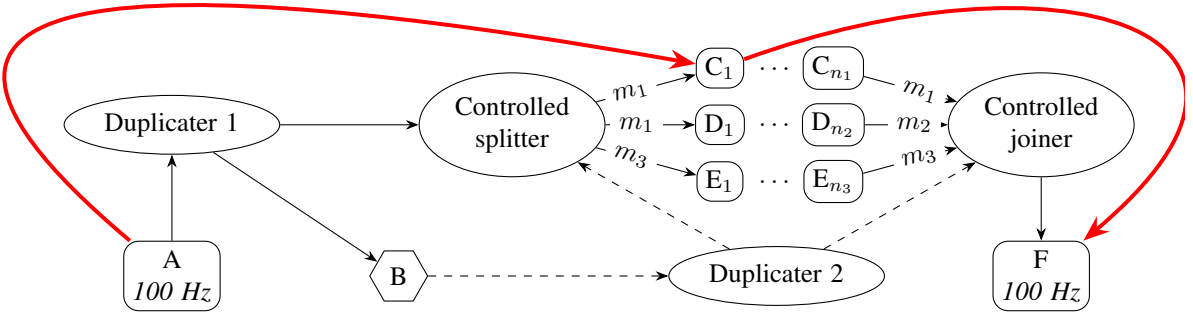
The pre-processing of an RMDF specification is detailed in Algorithm 6 (located on Page 11), which in turn relies on Algorithm 5. This latter converts a rational rate into the number of tokens produced or consumed at a specific job instance. Indeed, a rational rate of a PolyGraph specification only gives a number of tokens produced or consumed over a set of consecutive jobs, but not at every job instance.

Algorithm 6 pre-processes as many actors as possible using the initial tokens available in the channels. Actors with sufficient initial tokens in all their input channels are pre-processed. As each actor pre-processing produces tokens in their output channels, they may produce sufficient tokens for their successors to be pre-processed as well. Thus, each pre-processing must be followed by a check of the complete specification. Algorithm 6 terminates when no jobs can be executed without depending on at least one job instance from a sensor (which jobs are always at runtime). In other words,

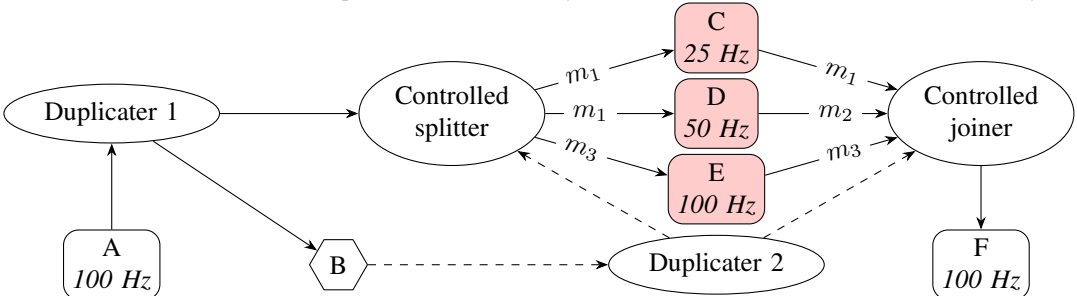




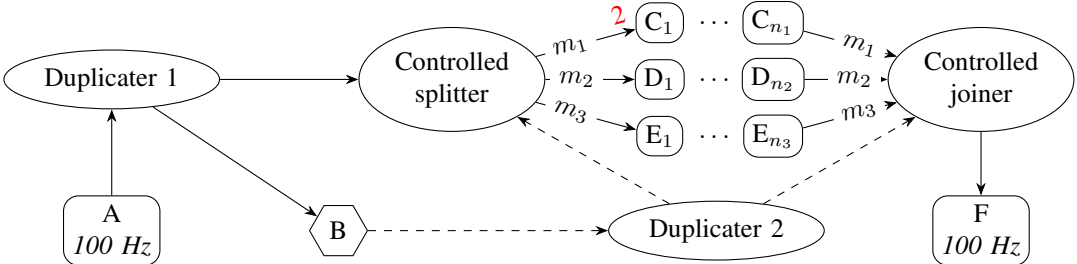
(b) There is a channel from an actor outside the control area to an actor inside the control area, and conversely.



(c) Actors C, D, and E have different frequencies. For readability, each conditional execution branch has exactly one actor.



(d) At least one production or consumption rate of the control area is not equal to 0 or 1.



(e) The sum of the assigned values of the production rates/consumption rates of the controlled splitter/controlled joiner is not equal to 1.  $m_1$  and  $m_2$  are assigned to 1 and  $m_3$  to 0.

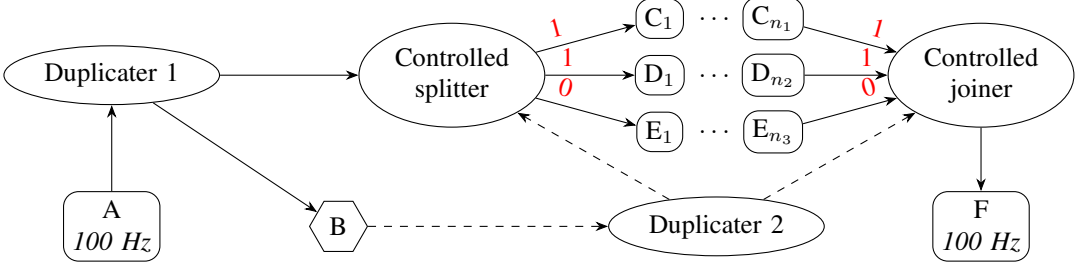


Fig. 6: RMDF specifications that do not meet the required restrictions presented in Property 1.

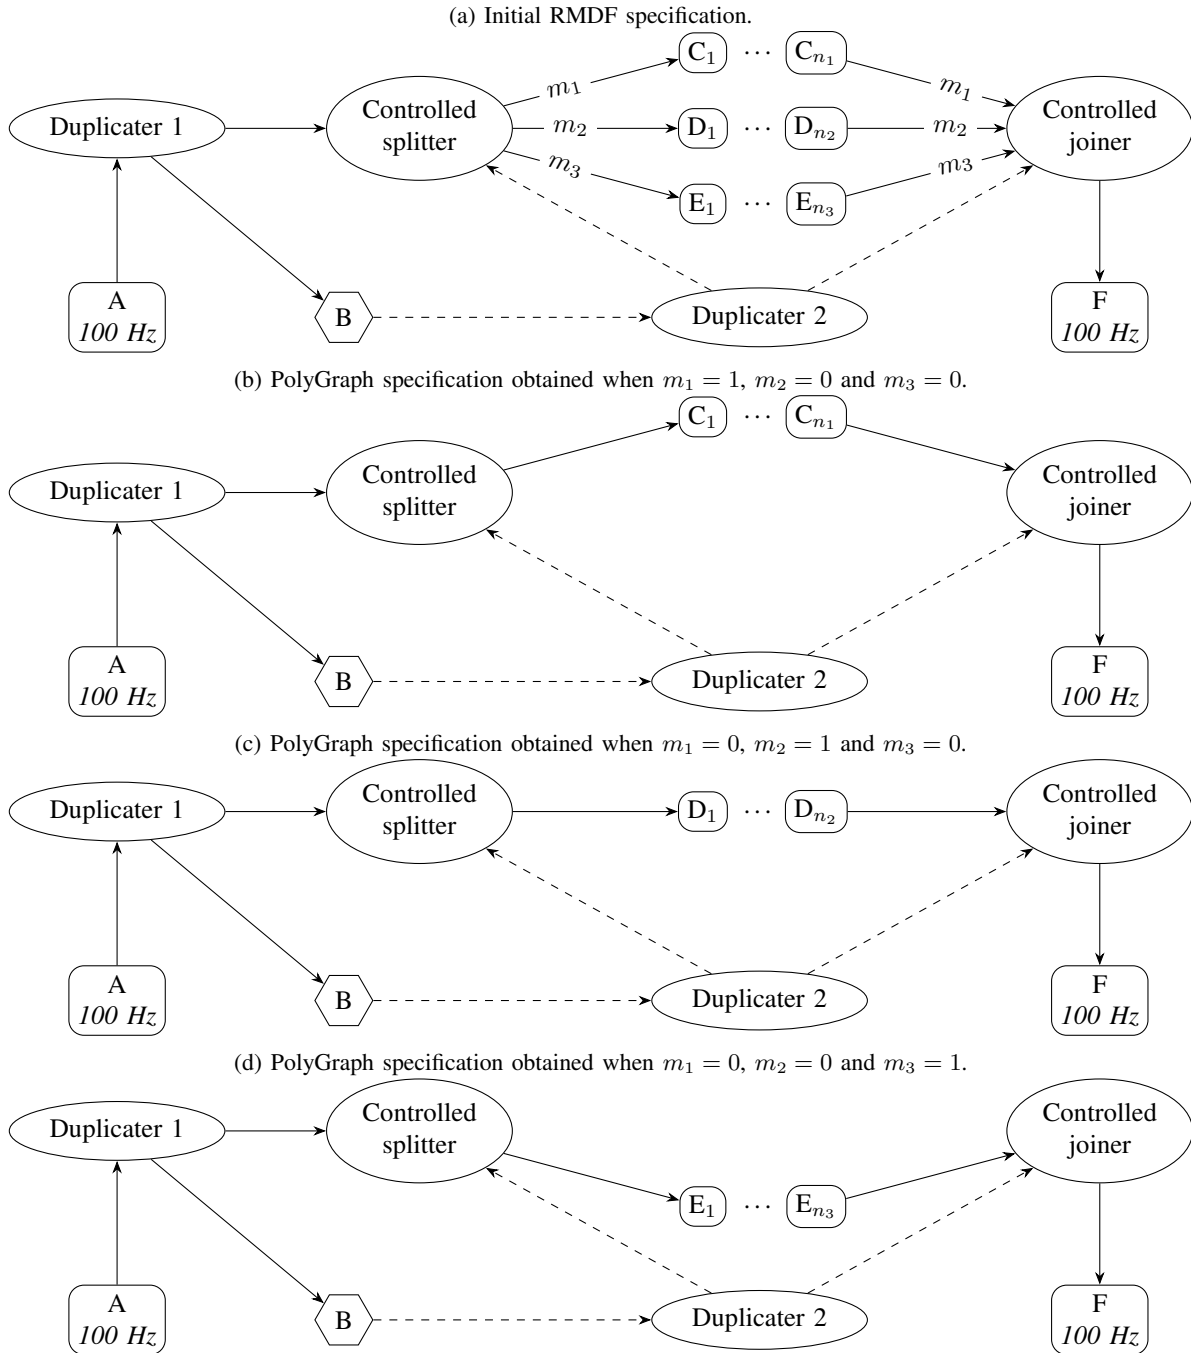


Fig. 7: An illustration of the PolyGraph specification obtained from an RMDF specification. All PolyGraph specifications are obtained by assigning values to  $m_1$ ,  $m_2$  and  $m_3$ .

all actors (except sensors) have a direct or indirect data dependency.

TABLE I: Notations and variables used in Algorithm 6.

Notation/Variable	Description
$tkn\_state$	A data structure that stores tokens produced by the actors. It is an array of queues with a FIFO policy. The length of this array is equal to the number of channels of the RMDF specification being pre-processed.
$ctrl\_tkn\_channel$	The channel on which the controlled splitter/joiner that consume control token will produce/consume its next token.

**Algorithm 5:** Conversion of a rational rate to the number of tokens produced or consumed at the job level (Algorithm 2 in [1]).

**Input:** The  $n$ -th job instance of an actor  $v \in V$  of a PolyGraph specification  $G = (V, E)$ , a channel  $c \in E$  from/to which  $v$  consumes/produces tokens with a rate  $\gamma = \frac{p}{q} \in \mathbb{Q}^*$  and with  $[c] \in \mathbb{Q}$  initial tokens.

**Output:** The number of tokens consumed/produced by the  $n$ -th job of  $v$  from/to  $c$ .

```

1  $r \leftarrow [c] - \lfloor [c] \rfloor$  ;
2 if  $\gamma > 0$  then
  | Result:  $\lfloor n \cdot \gamma + r \rfloor - \lfloor (n-1) \cdot \gamma + r \rfloor$ 
3 else
  | Result:  $\lceil n \cdot \gamma - r \rceil - \lceil (n-1) \cdot \gamma - r \rceil$ 

```

A consistent and live RMDF specification returns to its initial state after a hyperperiod. We claim that the timing constraints of a consistent, live and pre-processed RMDF specification (i.e., pre-processed with Algorithm 6) are cyclic over the first hyperperiod if the sensors and actuators of the specification have a frequency. Indeed, after pre-processing, all actors but the sensors have a direct or indirect data dependency with at least one sensor. Sensors are timed actors, so they release jobs periodically after the start of the execution. Consequently, they also periodically release their successors' jobs, ensuring cyclic timing constraints. The same reasoning applies to the actuators, which yield periodic deadline constraints that are propagated to their predecessor through data dependencies.

### B. Propagation of Timing Constraints

The arithmetic relation of [3], tailored initially for a PolyGraph specification, can be applied to the RMDF specification. To do that, we must define a similar notion as *well-defined PolyGraph specification* of [3] for an RMDF specification.

**Definition 4** (Well-defined PolyGraph/RMDF specification). *A well-defined PolyGraph/RMDF specification is a weakly connected graph<sup>5</sup> that is both consistent and live. Actors without predecessors/successors (i.e., the sensors/actuators) are timed actors, meaning they have a frequency constraint. All*

<sup>5</sup>A directed graph is weakly connected when its undirected induced graph is connected [9].

actors within the specification have both a BCET and a WCET, such that the BCET is less than or equal to the WCET. A well-defined PolyGraph/RMDF specification is also pre-processed, meaning that it has been processed with Algorithm 6. Note that for a PolyGraph specification, Line 10 to Line 35 of Algorithm 6 are not executed.

We remind the main result of [3] that allows the derivation of the execution windows of a PolyGraph specification. This result will be extended to an RMDF specification.

**Theorem 2** (Derivation of the execution windows of a PolyGraph specification - Theorem 1 of [3]). *Let  $G = (V, E)$  be a well-defined PolyGraph specification, let  $v_j, v_k \in V$  be two actors of  $G$ , and let  $c_i \in E$  be a channel of  $G$  from  $v_j$  to  $v_k$ . We define  $G_\Gamma = (\gamma_{ij}) \in \mathbb{Q}^{|E| \times |V|}$  to be the topology matrix of  $G$ , and  $r_i = [c_i] - \lfloor [c_i] \rfloor$  with  $\lfloor [c_i] \rfloor$  being the floor function applied to  $[c_i]$ . We define  $period_v$  and  $phase_v$  to be the period and phase of an actor  $v \in V$  if  $v$  is a timed actor. Let  $bcet_v/wcet_v$  be the BCET/WCET of an actor  $v \in V$ . Then, the following holds:*

- 1) *Let  $v_k \in V$  be a non-timed actor of  $G$ . Then, assuming  $release_{v_j}^n$  returns the release of the  $n$ -th job of an actor  $v_j \in V$ , the release of the  $p$ -th job of  $v_k$  is computed as follows:*

$$release_{v_k}^p = \max_{\substack{i \in [0, |E|] \\ j \in [0, |V|] \\ st \ \gamma_{ij} < 0}} release_{v_j}^{\alpha_1(p, i, j, k)} + bcet_{v_j} + (p - \beta_1(p, i, j, k)) \cdot bcet_{v_k}$$

with

$$\alpha_1(p, i, j, k) = \lceil \frac{p \cdot |\gamma_{ik}| - r_i - [c_i]}{\gamma_{ij}} \rceil$$

and

$$\beta_1(p, i, j, k) = 1 + \lfloor \frac{(\alpha_1(p, i, j, k) - 1) \cdot \gamma_{ij} + [c_i] + r_i}{|\gamma_{ik}|} \rfloor$$

- 2) *Let  $v_k \in V$  be a timed actor of  $G$ . Then, by denoting  $release_{v_k}^p$  the release computed by Item 1, the release of the  $p$ -th job of  $v_k$  is computed as follows:*

$$release_{v_k}^p = \max(release_{v_k}^p, period_{v_k} \cdot (p-1) + phase_{v_k})$$

- 3) *Let  $v_j \in V$  be a non-timed actor of  $G$ . Then, assuming  $deadline_{v_k}^p$  returns the deadline of the  $p$ -th job of an actor  $v_k \in V$ , the deadline of the  $n$ -th job of  $v_j$  is computed as follows:*

$$deadline_{v_j}^n = \min_{\substack{i \in [0, |E|] \\ k \in [0, |V|] \\ st \ \gamma_{ik} > 0}} deadline_{v_k}^{\alpha_2(n, i, j, k)} - wcet_{v_k} - (\beta_2(n, i, j, k) - n) \cdot wcet_{v_j}$$

with

$$\alpha_2(n, i, j, k) = 1 + \lfloor \frac{\lfloor (n-1) \cdot \gamma_{ij} + [c_i] \rfloor + r_i}{|\gamma_{ik}|} \rfloor$$

---

**Algorithm 6:** Pre-processing of an RMDF specification.

---

**Input:** An RMDF specification  $G = (V, E, M)$ .

**Output:** The RMDF specification where all jobs have a direct or indirect successor with at least one sensor's job.

```
1 def consume( $ts, v, c, n$ ):
2   forall  $i \in [0, \dots, n - 1]$  do  $consumed\_token.get(i) \leftarrow ts.get(v.c).dequeue()$ ;
3   Result:  $consumed\_token$ 
4 def produce( $ts, v, c, tkns$ ):
5   forall  $t \in tkns$  do  $ts.get(v.c).enqueue(t)$ ;
6  $should\_restart \leftarrow \mathbf{True}$ ;  $execute\_actor \leftarrow \mathbf{True}$ ;  $tkn\_state \leftarrow []$ ;
7 forall  $v \in V$  do if  $v.is\_ctrl\_splitter$  or  $v.is\_ctrl\_joiner$  then  $v.has\_consumed\_ctrl\_tkn \leftarrow \mathbf{False}$ ;
8 while  $should\_restart$ :
9    $should\_restart \leftarrow \mathbf{False}$ ;
10  for  $v \in V$ :
11    if  $v.is\_mode\_decider$  and  $ts.get(v.input\_channel.get(0)).size() \geq 1$  then
12       $tkn \leftarrow consume(ts, v, v.input\_channel.get(0), 1)$ ;
13       $ctrl\_tkn.channel \leftarrow v.decide\_mode(tkn)$ ;
14       $produce(ts, v, v.ctrl\_channel, ctrl\_tkn)$ ;
15       $should\_restart \leftarrow \mathbf{True}$ ;
16      continue
17    if  $v.is\_ctrl\_splitter$  or  $v.is\_ctrl\_joiner$  then
18      if  $v.has\_consumed\_ctrl\_tkn = \mathbf{False}$  and  $ts.get(v.ctrl\_channel).size() \geq 1$  then
19         $ctrl\_tkn \leftarrow consume(ts, v, v.ctrl\_channel, 1)$ ;
20         $v.has\_consumed\_ctrl\_tkn \leftarrow \mathbf{True}$ ;
21         $should\_restart \leftarrow \mathbf{True}$ ;
22        continue
23      else
24        if  $ts.get(v.ctrl\_tkn.channel).size() \geq 1$  then
25          if  $v.is\_controlled\_splitter$  then
26             $tkn \leftarrow consume(ts, v, v.input\_channel.get(0), 1)$ ;
27             $produce(ts, v, v.output\_channel.get(ctrl\_tkn.channel), tkn)$ ;
28             $v.has\_consumed\_ctrl\_tkn \leftarrow \mathbf{False}$ ;
29             $should\_restart \leftarrow \mathbf{True}$ ;
30            continue
31          else
32             $tkn \leftarrow consume(ts, v, v.input\_channel.get(ctrl\_tkn.channel), 1)$ ;
33             $produce(ts, v, v.output\_channel.get(0), tkn)$ ;
34             $v.has\_consumed\_ctrl\_tkn \leftarrow \mathbf{False}$ ;
35             $should\_restart \leftarrow \mathbf{True}$ ;
36            continue
37    for  $c \in v.input\_channels$ :
38      if  $ts.get(c).size() < \text{Algorithm 5}(v.job\_instance, c, c.cons\_rate, [c])$  then  $execute\_actor \leftarrow \mathbf{False}$ ;
39    if  $execute\_actor = \mathbf{True}$  then
40      for  $c \in v.input\_channels$ :
41         $n \leftarrow \text{Algorithm 5}(v.job\_instance, c, c.cons\_rate, [c])$ ;
42         $v.consumed\_tkns.get(c).append(consume(ts, v, c, n))$ ;
43       $v.produced\_tkns \leftarrow execute\_actor(v)$ ;
44      for  $c \in v.output\_channels$ :  $produce(ts, v, c, v.produced\_tkns.get(c))$ ;
45       $v.job\_instance \leftarrow v.job\_instance + 1$ ;
46       $should\_restart \leftarrow \mathbf{True}$ ;
```

---

and

$$\beta_2(n, i, j, k) = \left\lceil \frac{[\alpha_2(n, i, j, k) \cdot |\gamma_{ik}| - r_i] - [c_i]}{\gamma_{ij}} \right\rceil$$

- 4) Let  $v_j \in V$  be a timed actor of  $G$ . Then, by denoting  $\overline{deadline}_{v_j}^n$  the deadline computed by Item 3, the deadline of the  $n$ -th job of  $v_j$  is computed as follows:

$$deadline_{v_j}^n = \min(\overline{deadline}_{v_j}^n, period_{v_j} \cdot n + phase_{v_j})$$

The interested reader can find the proof of Theorem 2 in [3].

**Theorem 3** (Derivation of the execution windows of an RMDF specification). *Let  $G = (V, E, M)$  be an RMDF specification where all parameters are set to 1. Under the assumptions of Theorem 2, except that  $G$  is a well-defined RMDF specification instead of a well-defined PolyGraph specification, the result of Theorem 2 can be extended to  $G$ .*

*Proof of Theorem 3.* An RMDF specification as defined in the assumptions of Theorem 2, i.e., with all its parameters set to 1, can be analyzed as a well-defined PolyGraph specification, which is the assumption of Theorem 2.  $\square$

Theorem 3 extends the timing constraints propagation of a PolyGraph specification to an RMDF specification by considering all parameters to be set to 1. This implies a *conservative approach* for the execution windows assignment of the controlled splitters and controlled joiners. Deadlines are set to the minimum of the deadlines of their output branches. Similarly, the releases are set to the maximum of the releases of their input branches.

## VII. A FEASIBILITY TEST OF AN RMDF SPECIFICATION

The timing constraint propagation described in the previous section provides a timing assessment of an RMDF specification, serving as a feasibility test. Let  $G = (V, E)$  be a well-defined RMDF specification and  $G_{rv}(v)$  the repetition vector of  $G$ , that is, the number of executions of  $v$  in a hyperperiod of  $G$ . A *necessary* condition for the system to be feasible is to assert if the execution window length of all jobs is greater or equal to the WCET of the actor:

$$\forall v \in V, \forall i \in G_{rv}(v) : wct_v \leq deadline_v^i - release_v^i$$

The proposed timing assessment reduces the domain search space of dataflow specifications, as unfeasible RMDF specifications can be detected as early as possible. Fig. 8 illustrates the design flow we proposed with RMDF.

## VIII. APPLICATION TO THE VISION SYSTEM OF INGENUITY

As a reminder, a dataflow model is a trade-off between its expressiveness and analyzability. In Section IV, we have created the RMDF model, which extends the expressiveness of PolyGraph while maintaining its consistency and liveness analyses as detailed in Section V. In Section VI and Section VII, the analyzability of RMDF was extended while maintaining its expressiveness. The current section applies the previous results

to a dataflow specification of the vision processing system of Ingenuity.

### A. An Enhanced Model Fidelity with RMDF

We have specified the vision processing system of the Ingenuity Mars helicopter [10] in Fig. 9 with both the PolyGraph model (Fig. 9a) and the RMDF model (Fig. 9b). The vision processing system detects visual features on camera frames and either stores them as reference features (modeled as *Pseudo landmarks*) or tracks them from one frame to another. After a filtering procedure, the tracked visual features are matched with pseudo landmarks, enabling the computation of a visual drift to enhance position estimation.

New pseudo landmarks are generated if certain conditions are not met, and some of those conditions depend on runtime data. Thus, we have to simplify the pseudo landmarks generation procedure in the PolyGraph specification and assume that 1 image over 10 is used to generate new pseudo landmarks<sup>6</sup>.

The expressiveness of the RMDF model allows us to model the pseudo landmark generation with a mode decider actor (which is *Label decider* in Fig. 9b), which is more accurate than the simplification made in the PolyGraph specification. Indeed, the selection algorithm inside the *Label decider* may depend on runtime data, which is impossible to specify in the PolyGraph specification.

### B. Lazy Evaluation Propagation of Timing Constraints

The timing propagation presented in [3] and reminded in the previous section applies to an RMDF specification. Let us apply it to a slightly modified RMDF specification of the Ingenuity vision processing system of Fig. 9b. This modification, presented in Fig. 9c, is the following: 1) remove the frequency of *Feature Match*, 2) add an actor *Motors* with a frequency of 500 Hz and 3) add a channel between *Feature Match* and *Motors* with a production rate of 1, a consumption rate of 3/50 and 1/50 initial token. This modification resulted in a more interesting outcome for the timing propagation, as there is a fractional rate (the non-modified specification only has rates of 1). Note that the actor *Motors* specified the rest of the Ingenuity system, especially the motors triggering the helicopter's movements at a frequency of 500 Hz, as defined by the textual specification [11]. The results of the propagation of timing constraints is presented in Table II, in which we assumed that BCETs and WCETs of all actors are 0.12 and 0.20 ms, respectively.

The results of Theorem 2 and its extension to RMDF allow a timing constraints propagation with a lazy evaluation. Specifically, arithmetic formulas enable us to compute directly the job of interest, e.g., the successor job that consumes the token produced by an actor. As a result, only the necessary and sufficient computations are performed to derive the timing constraints. This lazy evaluation is particularly useful when

<sup>6</sup>In the real system, at least 1 image over 10 is used to generate pseudo landmarks. This condition is not the only one used in the real system. However, it is the only one that can be modeled statically. All other conditions depend on runtime data [10].

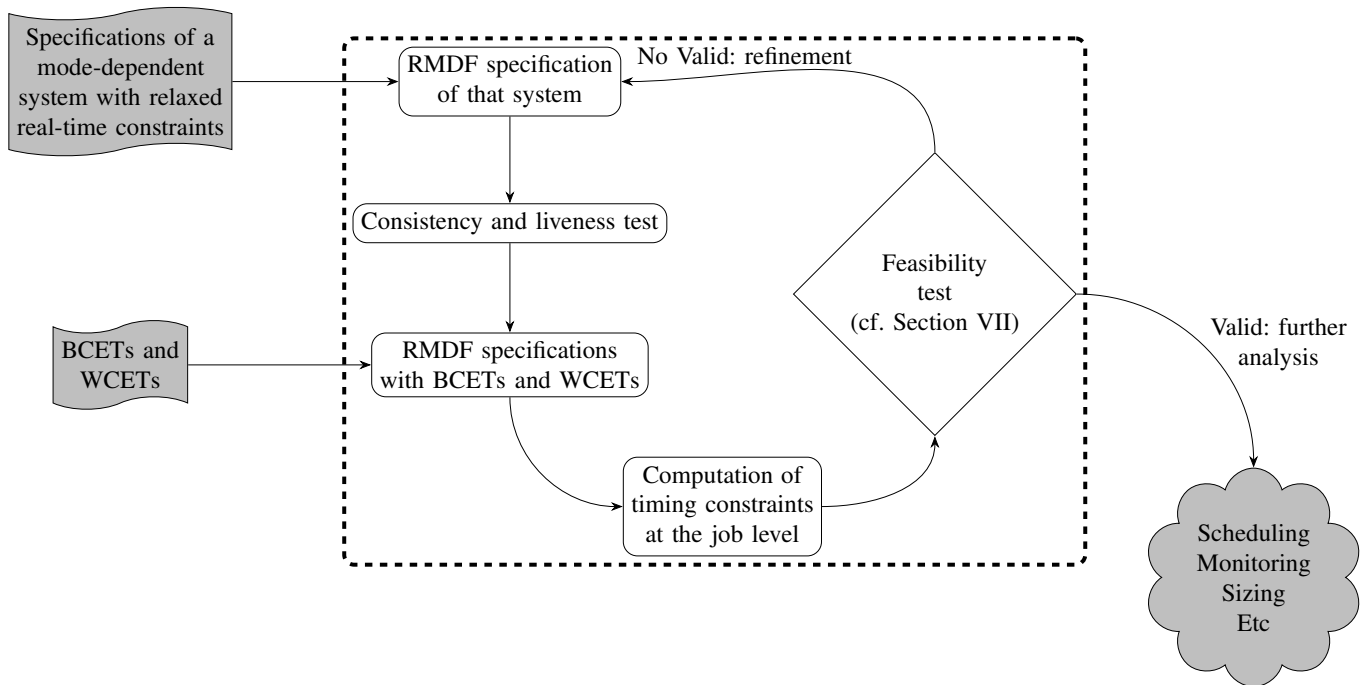


Fig. 8: A design flow with RMDF. The central dotted rectangle is the workflow we proposed.

timing constraints are computed for a subset of actors or even a subset of jobs.

### C. A Feasibility Test

To put into context the feasibility test presented in section Section VII, the result of Table II can be interpreted as follows. It is necessary that the WCET of all actors of Ingenuity are less than the minimum execution window length presented in column *Execution window* of Table II for Ingenuity to be feasible. Those values are summarized in Table III.

A schedulability test has been proposed for the PolyGraph model in [12]. While in this paper we do not consider a scheduling policy, the authors of [12] consider a preemptive context and interferences between actors. A response time analysis is performed to determine the worst-case response time of actors, considering the worst-case execution time of actors and the interferences between actors. Another difference is the timing evaluation procedure. While our approach is a lazy evaluation, the approach of [12] recursively updates the response time of actors until a fixed point is reached. Both approaches achieve the same result, and the lazy evaluation requires less computation if all timing constraints are unnecessary.

## IX. THE MODE CHANGE PROTOCOL OF RMDF

Mode-oriented execution relies on *Mode Change Protocols* (MCP) [13]. A *Mode Change Protocol* defines how the transition from one mode to another is handled. A *Mode Change Request* (MCR) is an event that triggers a mode change during a *transition phase*.

In the context of RMDF, an MCR is the production of a control token by a mode decider. Following the criteria

presented in [13], an MCP is classified according to three criteria: *overlapping*, *periodicity*, and *retirement*. In our context, those criteria are defined as follows:

- An *overlapping* protocol allows actors from two distinct execution branches to be executed at the same time.
- A *periodic* protocol allows actors outside conditional execution branches to not be interrupted by the mode change.
- A *late-retirement* protocol allows actors from an old mode, i.e., a previous execution branch, to continue their execution for a given amount of time.

The MCP proposed in this paper for RMDF is non-overlapping, periodic, and late-retirement. In order to explain the non-overlapping property, let us remind restriction 5 of control areas (cf. Property 1). This restriction states that whenever parametric rates of a control area are assigned a value, their sum is equal to 1. In conjunction with restriction 4, which states that production and consumption rates of a control area are equal to 0 or 1, it results in a single branch being executed at a time, resulting in the non-overlapping property. The proposed MCP is periodic as a mode change only affects the execution of actors in a control area. Finally, the late-retirement is ensured by the controlled joiner, which prevents the consumption of the token produced by the second branch until the first branch terminates. In other words, actors are allowed to finish their execution without compromising the expected order of tokens.

TABLE II: Releases, deadlines and execution windows length of the  $n$ -th job for each actor of the Ingenuity vision processing system, assuming BCET and WCET of all actors are 0.12 and 0.20 ms, respectively, except for the *Controlled Splitter* and *Controlled Joiner* which have constant execution time of 0 ms.

Actor	Timing constraint of the $n$ -th job		
	Release	Deadline	Execution window
Camera	$\begin{cases} 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 100/3 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 200/3 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 7/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 177/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 337/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 7/5 = 1.4 & \text{ms} \\ 31/15 \simeq 2.1 & \text{ms} \\ 11/15 \simeq 0.73 & \text{ms} \end{cases}$
Feature Detection	$\begin{cases} 3/25 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 2509/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 5009/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 8/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 178/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 338/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 37/25 = 1.48 & \text{ms} \\ 161/75 \simeq 2.1 & \text{ms} \\ 61/75 \simeq 0.81 & \text{ms} \end{cases}$
Label Decider	$\begin{cases} 6/25 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 2518/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 5018/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 9/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 179/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 339/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 39/25 = 1.56 & \text{ms} \\ 167/75 \simeq 2.2 & \text{ms} \\ 67/75 \simeq 0.87 & \text{ms} \end{cases}$
Controlled Splitter	$\begin{cases} 9/25 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 2527/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 5027/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 2 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 36 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 68 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 41/25 = 1.64 & \text{ms} \\ 173/75 \simeq 2.3 & \text{ms} \\ 73/75 \simeq 0.97 & \text{ms} \end{cases}$
Feature Tracking	$\begin{cases} 12/25 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 2536/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 5036/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 11/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 181/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 341/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 43/25 = 1.72 & \text{ms} \\ 179/75 \simeq 2.4 & \text{ms} \\ 79/75 \simeq 1.0 & \text{ms} \end{cases}$
Filtrng Procedure	$\begin{cases} 3/5 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 509/15 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 1009/15 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 12/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 182/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 342/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 9/5 = 1.8 & \text{ms} \\ 37/15 \simeq 2.5 & \text{ms} \\ 17/15 \simeq 1.1 & \text{ms} \end{cases}$
Pseudo Landmarks	$\begin{cases} 12/25 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 2536/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 5036/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 12/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 182/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 342/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 48/25 = 1.92 & \text{ms} \\ 194/75 \simeq 2.6 & \text{ms} \\ 94/75 \simeq 1.2 & \text{ms} \end{cases}$
Controlled Splitter	$\begin{cases} 18/25 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 2554/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 5054/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 13/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 183/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 343/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 47/25 = 1.88 & \text{ms} \\ 191/75 \simeq 2.5 & \text{ms} \\ 91/75 \simeq 1.2 & \text{ms} \end{cases}$
Feature Match	$\begin{cases} 21/25 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 0 \\ 2563/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 1 \\ 5063/75 + 100 \cdot \lfloor (n-1)/3 \rfloor & \text{if } (n-1) \bmod 3 = 2 \end{cases}$	$\begin{cases} 14/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 184/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \\ 344/5 + 100 \cdot \lfloor (n-1)/3 \rfloor \end{cases}$	$\begin{cases} 49/25 = 1.96 & \text{ms} \\ 197/75 \simeq 2.6 & \text{ms} \\ 97/75 \simeq 1.3 & \text{ms} \end{cases}$
Motors	$1 + 2 \cdot (n - 1)$	$1 + 2 \cdot n$	2 ms

TABLE III: Maximum WCETs of the actors of the Ingenuity vision processing system that constitute a necessary condition for its feasibility.

Actors	WCET	Actors	WCET
Camera	0.73 ms	Filtering Procedure	1.1 ms
Feature Detection	0.81 ms	Pseudo Landmarks	1.2 ms
Label Decider	0.87 ms	Controlled Splitter	1.2 ms
Controlled Splitter	1.2 ms	Feature Match	1.3 ms
Feature Tracking	1.0 ms	Motors	2 ms

## X. DISCUSSIONS AND RELATED WORKS

### A. Comparison of RMDF with LUSTRE and PRELUDE

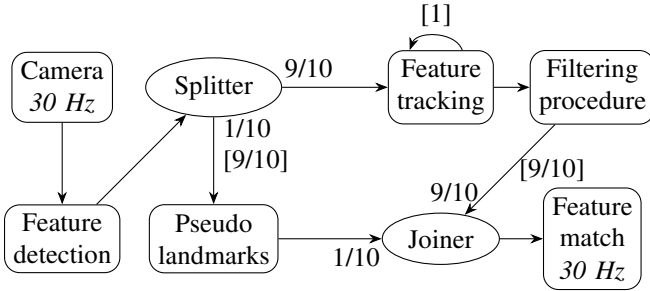
Synchronous languages such as LUSTRE [14] and PRELUDE [15] share the same goal as RMDF, that is to specify and analyze real-time systems. However, differences remain; the state of the art of LUSTRE and PRELUDE from [16] and [17] are used in the following paragraphs to explain

those differences. An RMDF specification is composed of *actors*, which communicate by exchanging data tokens on *channels*, and a system specified with LUSTRE is composed of *nodes*. Variables in nodes represent an infinite sequence of values, i.e., a *dataflow* or, more simply, a *flow*. Nodes are composed of *equations*, which can be understood as mathematical equations: each equations define an output flow from input flow(s).

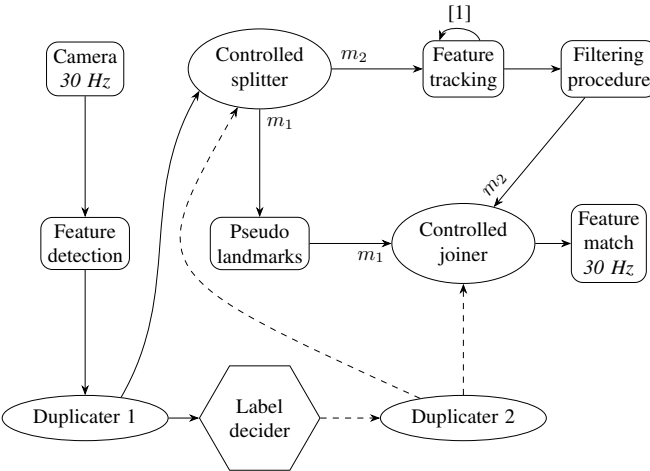
LUSTRE is structured around *clocks*: each dataflow has a clock and is present only when the clock is present, i.e., “it produces a tick”. A static analysis named *clock calculus* verifies the consistency of clocks, i.e., all combinations of flows can be evaluated without synchronization mechanisms such as buffering. This is different from RMDF (and dataflow models in general). In RMDF, produced data tokens are always present until they are consumed.

LUSTRE cannot directly specify real-time constraints, e.g., periods and deadlines. To overcome that, PRELUDE extends

(a) A PolyGraph specification of the vision processing system of Ingenuity.



(b) An RMDF specification of the vision processing system of Ingenuity.



(c) A modified RMDF specification of the vision processing system of Ingenuity to have more interesting result for the timing constraints propagation.

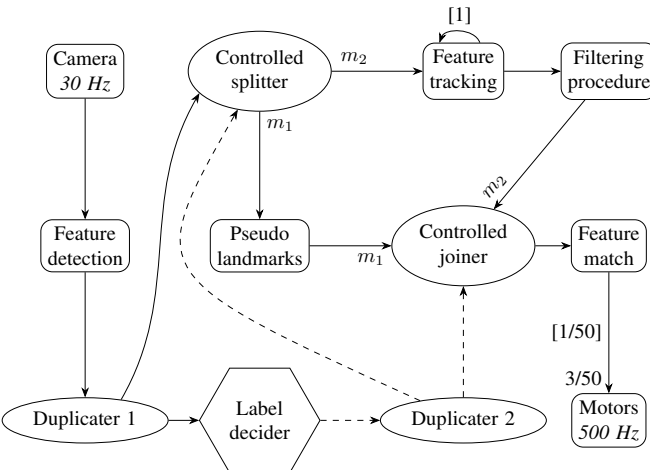


Fig. 9: Specifications of the Ingenuity vision processing system with PolyGraph and RMDF. All those specifications are consistent and live. Rates of 1 and initial tokens of 0 are omitted for clarity.

LUSTRE to allow system designers to explicitly define the period of a flow (i.e., the period of its clock) or let the compiler infer it. As an illustration, yet *without proving any semantic equivalence*, the same logic in RMDF would be to specify the period of production and consumption of data tokens instead of the period of the actors.

A system specified with PRELUDE can then be composed of flows with different periods; those flows are not *synchronous* in the sense of [14]. In order to combine non-synchronous flows, *rate-transition operators* allow the production of a flow that is faster or slower than the input flow. As an illustration, yet again *without proving any semantic equivalence*, the rational rates RMDF greater/lower than 1 allow for specifying faster/slower data tokens production and consumption.

### B. Comparing RMDF with an Extended PRELUDE

The static analysis of the proposed RMDF relies on restrictions applied to control areas (cf. Property 1 for the exhaustive list of restrictions). Among those restrictions, the third asserts that “if there are timed actors in a control area, they must have the same frequency” and the fourth asserts that “the production rate inside a control area and the parametric rates associated to the controlled joiner and controlled splitter of that control area must be equal to one”.

LUSTRE has been extended in [18] to allow the specification of multi-mode system. A switch-like statement allows to switch between sets of equations according to some enumerated value and an automaton construction. In addition, in the work of [18], all flows within a state must have the same period. The mode-dependent systems considered in this paper have a similar behavior because of the third and fourth restrictions presented in Property 1.

The gap between and PRELUDE and [18] is addressed in [13]. The best of both works is used to support the specification of states with multi-periodic states in PRELUDE. The work of [13] addressed the third and fourth restrictions of Property 1. However, this work is in the context of synchronous languages, and work remains to be done to know if the same logic can be applied to RMDF and dataflow models.

### C. Dataflow Models with Conditional Execution

Dynamic reconfiguration has been introduced in the dynamic PolyGraph model [19]. In this reference, a set of actors can select an execution mode. In addition to its data, each token is labeled with an execution mode. When an actor consumes tokens, it reads the mode of the token and decides how to process the data. Note that this processing occurs as a black box inside the actor’s implementation. The processing of input tokens can also include discarding input tokens. A key advantage of this procedure is that the consistency and liveness checking algorithms of PolyGraph can be directly extended to dynamic PolyGraph. Mode execution is transmitted transitively from actor to actor, meaning all actors of a dynamic PolyGraph model must have a specific implementation to process execution mode. In the RMDF model, the execution modes are carried only from mode deciders to controlled



splitters and controlled joiners. In other words, actors in a control area of an RMDF specification do not have a specific implementation that considers the execution mode. Therefore, RMDF allows for the easier integration of actors inside a control area. Our approach puts the complexity at the model level, while in dynamic PolyGraph, the complexity is at the implementation level.

TPDF [20] is a generalization of CSDF where some actors are *control actors*. Those latter are similar to our mode decider. A control actor outputs a control token, which a subsequent actor consumes to decide how this actor consumes or produces its tokens. For instance, it can select one of the input or output tokens or select the available input token with the highest priority. The expressivity of the different execution modes available is greater in TPDF than in RMDF. The key advantage of RMDF is that it is easier to specify data communication patterns and timing constraints than TPDF.

## XI. CONCLUSION AND FUTURE WORKS

We have presented the RMDF model, a dataflow model that extends the PolyGraph model. This latter is well-suited to specify CPSs with relaxed real-time constraints. Besides keeping this advantage, RMDF allows for the specification of CPSs with a mode-dependent execution, i.e., with a set of conditional execution branches. In such CPSs, some actors can select from/to which input/output channel tokens are consumed/produced, and this choice may depend on runtime data. We have extended the static analysis algorithm of the PolyGraph model to the RMDF model to verify the consistency and liveness of an RMDF specification. The timing behavior analysis of the PolyGraph model presented in [3] is also extended to RMDF, and we presented a feasibility test. Static analyses of the RMDF model have also been implemented.

A dataflow model is a trade-off between expressiveness and analyzability. The expressiveness of RMDF could be improved by relaxing restrictions imposed on control areas, while the analyzability could be enhanced by incorporating mode dependencies within an RMDF specification. This approach is reasonable, as execution modes in real systems do not occur arbitrarily. More generally, it would be valuable to explore how the feasibility test proposed for RMDF can be applied to other design models beyond dataflow models, such as MatLab/Simulink. To that end, it is necessary to define semantically equivalent model transformations from those models to RMDF.

## REFERENCES

- [1] P. Dubrulle, C. Gaston, N. Kosmatov, A. Lapitre, and S. Louise, "A Data Flow Model with Frequency Arithmetic," in *Fundamental Approaches to Software Engineering - 22nd International Conference*, 2019, pp. 369–385.
- [2] G. Roumage, S. Azaiez, and S. Louise, "A Survey of Main Dataflow MoCCs for CPS Design and Verification," in *2022 IEEE 15th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc)*, 2022, pp. 1–9.
- [3] G. Roumage, S. Azaiez, C. Faure, and S. Louise, "Static Timing Analysis of Cyber-Physical Systems with Relaxed Real-Time Constraints," in *SEAA 2024: 50th Euromicro Conference Series on Software Engineering and Advanced Applications*, 2024.
- [4] P. Dubrulle, N. Kosmatov, C. Gaston, and A. Lapitre, "PolyGraph: A Data Flow Model with Frequency Arithmetic," *International Journal on Software Tools for Technology Transfer*, vol. 23, no. 3, pp. 489–517, 2021.
- [5] E. A. Lee and D. Messerschmitt, "Synchronous Data Flow," *Proceedings of the IEEE*, vol. 75, no. 9, pp. 1235–1245, 1987.
- [6] G. Bilsen, M. Engels, R. Lauwereins, and J. A. Peperstraete, "Cyclo-Static Dataflow," *IEEE Transactions on Signal Processing*, vol. 44, no. 2, pp. 397–408, 1996.
- [7] P. De Oliveira Castro, S. Louise, and D. Barthou, "Reducing Memory Requirements of Stream Programs by Graph Transformations," in *Proceedings of the 2010 International Conference on High Computing & Simulation*, 2010, pp. 171–180.
- [8] S. Louise, "Graph Transformations and Derivation of Scheduling Constraints Applied to the Mapping of Real-Time Distributed Applications," in *2019 IEEE 13th International Symposium on Embedded Multicore/Many-core Systems-on-Chip*, Singapore, Singapore, 2019, pp. 295–303.
- [9] A. Honorat, K. Desnos, S. S. Bhattacharyya, and J.-F. Nezan, "Scheduling of Synchronous Dataflow Graphs with Partially Periodic Real-Time Constraints," in *Proceedings of the 28th International Conference on Real-Time Networks and Systems*, 2020, pp. 22–33.
- [10] D. S. Bayard, D. T. Conway, R. Brockers, J. H. Delaune, L. H. Matthies, H. F. Grip, G. B. Merewether, T. L. Brown, A. M. San Martin, "Vision-Based Navigation for the NASA Mars Helicopter," in *AIAA Scitech 2019 Forum*, 2019.
- [11] H. Grip, J. N. Lam, D. Bayard, D. T. Conway, G. Singh, R. Brockers, J. Delaune, L. Matthies, C. Malpica, T. Brown, A. Jain, M. San Martin, and G. Merewether, "Flight Control System for Nasa's Mars Helicopter," in *AIAA Scitech*, 2019.
- [12] E. Hamelin, A. Berne, M. Boudiaf, and P. Dubrulle, "Performance and Confidence in Feasibility Analysis of Real-Time Multi-Core Distributed Systems," 2024.
- [13] F. Fort and J. Forget, "Synchronous semantics of multi-mode multi-periodic systems," in *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing*. Virtual Event: ACM, Apr. 2022, pp. 1248–1257.
- [14] N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud, "The Synchronous Data Flow Programming Language Lustre," *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1305–1320, 1991.
- [15] J. Forget, F. Boniol, D. Lesens, and C. Pagetti, "A Multi-Periodic Synchronous Data-Flow Language," in *2008 11th IEEE High Assurance Systems Engineering Symposium*. Nanjing, China: IEEE, 2008, pp. 251–260.
- [16] F. Fort, "Programing adaptive real-time systems," Ph.D. dissertation, 2022.
- [17] J. Forget, "Programming and Analysis of Critical Real-Time Systems," HDR Thesis, 2023.
- [18] J.-L. Colaço, B. Pagano, and M. Pouzet, "A conservative extension of synchronous data-flow with state machines," in *Proceedings of the 5th ACM International Conference on Embedded Software*. Jersey City NJ USA: ACM, Sep. 2005, pp. 173–182.
- [19] P. Dubrulle, C. Gaston, N. Kosmatov, and A. Lapitre, "Dynamic Recon-figurations in Frequency Constrained Data Flow," in *Integrated Formal Methods*, vol. 11918, 2019, pp. 175–193.
- [20] X. Khanh Do, S. Louise, and A. Cohen, "Transaction Parameterized Dataflow: A Model for Context-Dependent Streaming Applications," in *2016 Design, Automation & Test in Europe Conference & Exhibition*, 2016, pp. 960–965.