

# *From multi-clock constraints to multi-rate GALS executives*

Dumitru Potop-Butucaru — Robert de Simone — Yves Sorel

**N° 6021**

Novembre 2006

Thème COM



*Rapport  
de recherche*



## From multi-clock constraints to multi-rate GALS executives

Dumitru Potop-Butucaru , Robert de Simone , Yves Sorel

Thème COM — Systèmes communicants

Projet AOST

Rapport de recherche n° 6021 — No

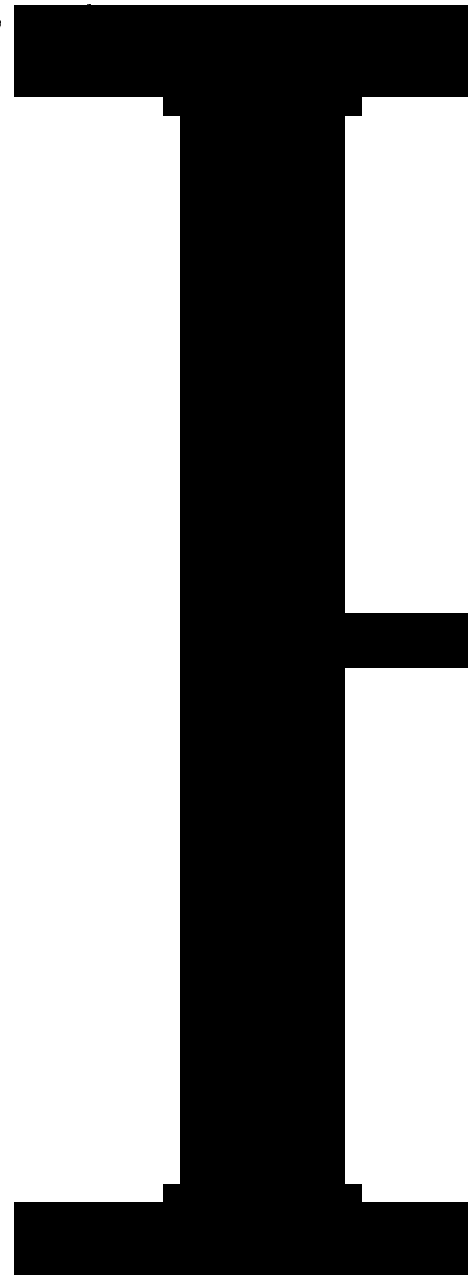
**Abstract:** We define a method for synthesizing the asynchronous executives that are driving the synchronous modules of a globally asynchronous, locally synchronous (GALS) system. The technique takes as input high-level synchronization constraints under the form of multi-clock modular synchronous reactive (S/R) specifications. For each asynchronous module, our technique produces a multi-rate executive that drives the module using the clock of the component using a mixed static/dynamic scheduling policy. The resulting GALS system is predictable and functionally correct with respect to the original synchronous specification. The approach is based on the theory of weakly endochronous systems and relies on a notion of atomic reaction which allow us to exploit the concurrency in order to improve the communication efficiency of the executives.

**Key-words:** Synchronous, Asynchronous, GALS, Multi-clock, Endochronous, Endochrony, Automatic distribution

## Synthèse d'exécutifs GALS à partir de contraintes multi-horloges

**Résumé :** Nous définissons une méthode pour la synthèse des exécutifs asynchrones qui contrôlent l'exécution des modules synchrones d'un système globalement asynchrone, localement synchrone (GALS). Notre technique prend en entrée des contraintes de synchronisation de haut ni eau, sous la forme de spécifications synchrones réactives modulaires multi-horloges. Pour chaque module synchrone, notre technique produit un exécutif multi-rythmes qui contrôle les entrées, les sorties, et l'horloge du module en utilisant une politique d'ordonnancement mixte statique/dynamique. Le système GALS résultant est prédictible, correct et complet par rapport à la spécification synchrone d'entrée. L'approche est fondée sur la théorie des systèmes faiblement endochrones (weakly endochronous systems), qui nous permet d'exploiter la concurrence de la spécification pour améliorer l'efficacité des exécutifs.

**Mots-clés :** Synchrone, Asynchrone, GALS, Multi-horloges, Multi-rythme, faible, Répartition automatique



## 1 Introduction

Development techniques based on synchronous formalisms are today common in various fields, such as digital circuit design or (safety-critical) embedded software development [3]. Synchronous specifications are used in these fields to model concurrent reactive systems. The synchronous model allows the explicit representation of the reaction to signal absence, and thus supports a notion of *deterministic concurrency* which facilitates functional modelling and analysis. Implementing synchronous specifications is often hard. This is mainly due to the global synchronization mechanisms of the model, which need to be preserved, at least in part, to ensure that absence information is not lost when it can influence the behavior of the system.

In this paper, we address the synthesis of the asynchronous executives that are driving the synchronous modules of a GALS system. The topic draws much attention today, when distributed embedded systems and complex systems-on-chip (SoC) are commonplace in the industry. Our synthesis technique is mathematically founded on the theory of *weakly endochronous systems*, due to Potop, Caillaud, and Benveniste [16]. Weak endochrony gives criteria establishing that a synchronous presentation hides a behavior where the absence information is never needed, so that the synchronous specification can safely be executed in an asynchronous environment, with predictable results. Weak endochrony extends to a synchronous framework the Mazurkiewicz traces [12]. Thus, the current paper bridges between the fields of synchronous language compilation and classical concurrency theory.

The implementation technique is defined as an *assisted synthesis technique* that takes as input high-level synchronization constraints structured in a multi-clock synchronous specification. The transformation is performed in two steps. The first one checks whether the specifications of the synchronous modules are weakly endochronous. If they are not, intuitive diagnostics allow the user to incrementally improve the signalling protocols. When the specification is weakly endochronous, the second translation step automatically synthesizes the GALS executives.

**Outline.** The paper is organized as follows: Section 2 intuitively presents our problem, previous work, and the desired solution. Section 3 defines the formalism that will support our presentation. Section 4 is on weak endochrony: the original theory of [16] and algorithms to determine if a specification is Weakly Endochronous. Section 5 shows how the GALS executives are generated for Weakly Endochronous specifications. Section 6 shows the synthesis of the resulting GALS executives.

## 2 Overview

Our goal is the construction of correct and efficient GALS executives for GALS systems. In this paper, we consider this problem in the context of real-time constraints. We aim to account or ensure the satisfaction of real-time constraints in the synthesis of the GALS executives. This is a non-trivial task. The work presented in this paper is a first step towards this goal.

## 2.1 Distributed implementation

Our basic model of distributed system is similar to that of Kahn Process Networks (KPN) [10]. A system is formed of computing *components*. The components communicate through *message passing* along *lossless order-preserving asynchronous lines*. For simplicity, our examples only use simple point-to-point communication lines—*asynchronous FIFOs*.

Reading or writing a message on a communication line is blocking and non-interruptible. Blocking reads are part of the basic KPN model. Blocking writes are needed to ensure flow regulation in the absence of timing or synchronization information on the environment. With blocking writes, we can use real-life bounded-memory communication lines instead of the unbounded FIFOs of the basic Kahn model.

## 2.2 Modular synchronous specification

To synthesize our executives, we start with the set of synchronous modules of the GALS system: *off-the-shelf synchronous IPs* or compiled synchronous code, under the form of software *reaction functions*. Following the terminology of [5], we call these implementation modules *pearls*.

The [redacted] of its synchronization properties. The [redacted] to the pearls. They should control the [redacted] of the asynchronous operations, can, and should be abstracted [redacted] to facilitate analysis.

The input of our [redacted] technique is the modular synchronous [redacted] obtained by putting in parallel [redacted] descriptions of all the pearls. [redacted] which will support our [redacted] small sub-set of the Signal language. [redacted] means that the specification is [redacted] composition of a number of *processes*.

We use a small [redacted] example to present our problem, the desired [redacted] the main implementation is [redacted] example, pictured in Fig. 1, is a simple [redacted] adder, where two independent [redacted] ALUs can be used either independently [redacted] synchronized to form a double- [redacted] choice between synchronized and [redacted] mode is done using the [redacted] to simplify the figure, we compacted both [redacted] of each adder under a single [redacted] (I1 and I2, respectively).

The carry between [redacted] adders is propagated through the [redacted] SYNC is present. The two [redacted] ALUs are the processes of our specification.

As we shall see [redacted] paper, the high-level specification [redacted] example can be functionally incorrect [redacted] safely abstract away the [redacted] of the adders, and the integer data.

### 2.2.1 Synchronous processes

Like any synchronous [redacted], we follow a discrete model [redacted] are executions are sequences of *reactions* [redacted] by a *global clock*. Table 1 [redacted] the execution of

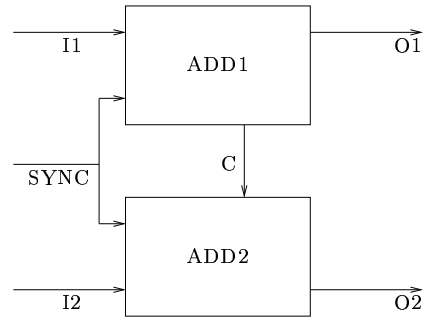


Figure 1: Configurable adder: global dataflow

Clock	1	2	3	4	5	6	7
I1	(1,2)	⊥	(9,9)	(9,9)	⊥	(2,5)	⊥
O1	3	⊥	8	8	⊥	7	⊥
SYNC	⊥	⊥	⊥	⊥	⊥	⊥	⊥
C	⊥	⊥	1	⊥	⊥	0	⊥
I2	⊥	⊥	(0,0)	(0,0)	⊥	(1,4)	(2,3)
O2	⊥	⊥	1	0	⊥	5	5

Table 1: Sample synchronous run of the example

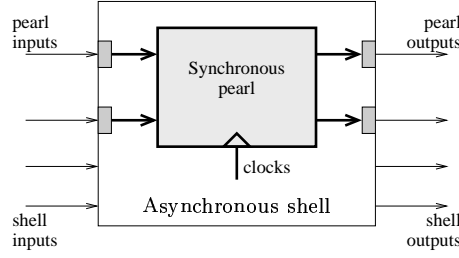


Figure 2: Desired component structure

our example. A reaction is a valuation of the input, output, and internal *signals* of the process. We shall denote with  $\mathcal{V}$  the finite set of signals of a process. In our example,  $\mathcal{V} = \{1, 2, YNC, O1, O2, C\}$ .

All signals are typed. We denote with  $\mathcal{D}_S$  the domain of a signal  $S$ . Not all signals need to have a value in a reaction, to model cases where only parts of the process compute. We will say that a signal is *present* in a reaction when it has a value in  $\mathcal{D}_S$ . Otherwise, we say that it is *absent*. Absence is simply represented with value  $\perp$ , which is appended to all domains  $\mathcal{D}_S^\perp = \mathcal{D}_S \cup \{\perp\}$ . Formally, a reaction of the process is a partial valuation of its signals. We denote with  $\mathcal{R}$  the set of all such valuations. The *support* of a reaction  $r$ , denoted  $\text{supp}(r)$ , is the set of present signals. For instance, the support of reaction 4 in Table 1 is  $\{1, 2, O1, O2\}$ .

In many cases we are only interested in the presence or absence of a signal, because it transmits no data, or because we are only interested in synchronization aspects. To represent such signals, the Signal language uses a dedicated *event* type of domain  $\mathcal{D}_{\text{event}} = \{\top\}$ . In our example, SYNC has type *event*.

To represent reactions, we use a *set-like convention* and omit signals with value  $\perp$ . In Fig. 1, the signal set is  $\{YNC : \text{event}, O1, O2 : \text{integer}, 1, 2 : \text{integer\_p}, r\}$ . In Table 1, reaction 4 is denoted  $(1^{(9,9)}, O1^8, 2^{(0,0)}, O2^0)$ . The *stuttering reaction* assigning  $\perp$  to all signals is denoted  $\perp$ . In Table 1, reaction 5 is a stuttering reaction.

### 2.3 Structure of a computing component

To interface between the asynchronous communication lines and the synchronous pearls, we structure our computing components as pictured in Fig. 2.

We are interested here in the asynchronous executive, also called *shell*, which controls the asynchronous I/O and drives the synchronous pearl of the component. The shell is basically an asynchronous automaton<sup>1</sup>. It takes as input the asynchronous FIFOs transmitting the inputs of the pearl, and it can also take additional control inputs added by the synthesis process to ensure correct synchronization. When enough input signals are present, the shell

<sup>1</sup>But its implementation can be synchronous, like in [5], or simply sequential.



I1	(1,2)	(9,9)	(9,9)	(2,5)
O1	3	8	8	7
SYNC		$\top$	$\top$	
C		1	0	
I2	(0,0)	(0,0)	(1,4)	(2,3)
O2	1	0	5	5

Table 2: Corresponding asynchronous run. Correctly reconstructing synchronous inputs from the asynchronous ones is impossible.

triggers a synchronous reaction of the pearl. To do this, the shell provides the pearl with the needed inputs, and marks all other pearl inputs as *absent* ( $\perp$ ) whenever values have been received for them. Then, it triggers the computation of the reaction. In software, this amounts to calling the reaction function. In hardware, the pearl clock is enabled. Once the pearl completes its computation: in hardware, the pearl clock is disabled; in software, the state update protocol is performed. The outputs computed by the pearl are transmitted onto the output FIFOs. The input values used in the reaction are acknowledged as read, so that new ones can be accepted from the corresponding FIFOs.

## 2.4 The synthesis problem

We need to synthesize the (hardware or software) shells that will drive the execution of the pearls on the components of the architecture.

The main task of the shell is to construct in a consistent way inputs for the synchronous pearl. The input of a reaction must assign a value, or declare absent, every input of the pearl. With this hypothesis, code generation for the pearl is highly simplified<sup>2</sup>, for it does not have to infer or enforce input presence or absence, nor buffer inputs or outputs to match the asynchronous transmission protocols.

Reconstructing reactions from asynchronous messages must be done in a deterministic fashion, regardless of the message arrival order. This is not always possible. Assume, like in Table 2, that we consider the inputs of Table 1 without the synchronization information. The module ADD1 will then receive the first value (1 2) on the input channel I1 and  $\top$  on SYNC. Depending on the arrival order, which cannot be precised, any of the reactions ( $1^{(1,2)} O1^3 \text{ SYNC}^\top C^0$ ) or ( $1^{(1,2)} O1^3$ ) can be triggered, leading to divergent computations. The problem is that the two reactions are not independent (they do not commute), yet the choice between them cannot be done over the value of a message on a given channel, so a shell has no means of choosing between them. A similar problem occurs in ADD2.

Our approach is to transform the synchronous specification in such a way as to make reaction reconstruction deterministic upto commutation of independent reactions. One such implementation of our example is presented in Fig. 3, with an execution trace in Table 4.

<sup>2</sup>The code can be *exochronous*, in the sense of [14].

Transformation is done by adding communication lines. To discriminate between the non-independent transitions of ADD1, we introduce the Boolean signal SYNC1. Value 1 on SYNC1 indicates that the two adders synchronize, so that SYNC and C are present. Value 0 indicates an independent computation on ADD1. Symmetrically, SYNC2 is added to ADD2. Note that signal SYNC becomes useless, as it corresponds to choices of either SYNC1 and SYNC2. Therefore, we can delete it from the dataflow scheme. Note, in Table 4, that the two processes can be now safely run on separate clocks in the GALS implementation.

## 2.5 Previous solutions

A particular class of solutions to our problem has already been thoroughly studied: The case where every pearl reads all inputs and writes all outputs at each reaction (none can be absent). We shall call this the mono-clock case, for reasons that will become clear in Section 3. In this case, implementations have been given in both software and hardware:

- In SynD [1], the output of optimized multi-processors is statically scheduled on a common clock.
- In Latency-insensitive technology-independent synthesis [5], the model serves as basis for the synthesis of communication protocols.

Putting our example in mono-clock asynchronous lines the absence of synchronization may be inefficient when we want to perform multi-rate computation.

The first extension towards multi-rate pearls were the *endochronous systems* [17]. Here, we assume that the presence and absence of all signals is known incrementally inferred starting from the signals that are always present. Table 3 presents a run of an endochronous system. Figure 4 shows the transformation of our example into one that carries out multi-rate computation. Figure 5 shows the transformation of ADD2 executing alone, 2 for both adders, and 3 for the synchronized execution. The presence/absence of signals determines the presence/absence of transitions. The compilation of the Signal language is based on a pearl [1].

The problem with endochronous systems is that they do not provide synchronous modules. The reaction records of the system are not necessarily synchronized w.r.t. the signals. In our example, an endochronous system would not allow concurrency of independent computations of the two adders. If we cannot allow implementation of ADD1 and ADD2 are implemented on separate clocks, we can implement them through added signals. This leads to over-synchronizing the inputs,

would consist in explicitly transmitting the signals, *i.e.*, transmitting all the symbols in the system. This may optimize communication, or when we want to perform multi-rate computation.

Figure 4 shows the transformation of our example into one that carries out multi-rate computation. Figure 5 shows the transformation of ADD2 executing alone, 2 for both adders, and 3 for the synchronized execution. The presence/absence of signals determines the presence/absence of transitions. The compilation of the Signal language is based on a pearl [1].

The problem with endochronous systems is that they do not provide synchronous modules. The reaction records of the system are not necessarily synchronized w.r.t. the signals. In our example, an endochronous system would not allow concurrency of independent computations of the two adders. If we cannot allow implementation of ADD1 and ADD2 are implemented on separate clocks, we can implement them through added signals. This leads to over-synchronizing the inputs,

Clock	1	2	3	4	5
I1	(1,2)	(9,9)	(9,9)	(2,5)	$\perp$
O1	3	8	8	7	$\perp$
SYNC	0	3	2	3	1
C	$\perp$	1	$\perp$	0	$\perp$
I2	$\perp$	(0,0)	(0,0)	(1,4)	(2,3)
O2	$\perp$	1	0	5	5

Table 3:

Clock1	1			
I1	(1,2)	(9,9)	(9,9)	
O1	3	8	8	
SYNC1	0	1	0	
C		1		
SYNC2		1		0
I2		(0,0)		(2,3)
O2		1		5
Clock2		1		4

Table 4: Weakly endochron

endochrony is not compositional, which makes the incremental development of processes difficult.

## 2.6 Our approach

In this paper, we explain how to allow both synchronous and asynchronous (interleaved) computations to be realized by a single compiler. In our approach, the user can decide at pearl synthesis time (maybe through an annotation) how to group the processes onto the distributed components, while knowledge of the efficient shells are automatically generated. For instance, the two processes of a component can be implemented by separate pearls, or by a single one.

The main contribution is given in Section 4. We explain how to determine whether a process is weakly endochronous. This is based on determining a minimal set of reactions that generate all other reactions. The process is weakly endochronous when all generators satisfy the *atomicity hypothesis* – they are either asynchronous, commutable, or non-interferent. When a process is not weakly endochronous, we need to transform it further by eliminating non-confluent choices that are not determined by a generator (not by absence). Given the reaction lines, we do not fully

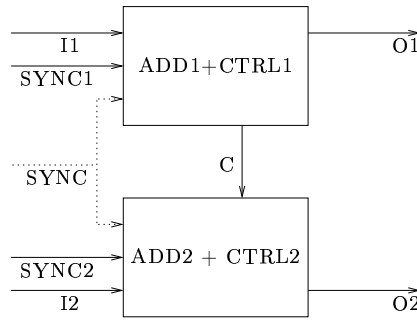


Figure 3: Weakly endochronous solution

automate it. Instead of providing solutions that might not match the global architecture of the system, our analysis algorithm simply points out through intuitive error messages why the system is not weakly endochronous. Once the assisted transformation step is completed, the asynchronous shells are automatically generated, as explained in Section 5.

### 2.6.1 Stateless abstraction

To keep the presentation simple, we present here a simple version of the technique, which does not take into account the process state. The approach can be extended to deal with specifications where each process has a state.

## 3 Synchronous Dataflow in Signal

We use a small sub-set of the high-level synchronous language Signal to formally present our technique. The Signal multi-clock constraint language allows a simple representation of the synchronization constraints we use. However, the results can be simply applied to any synchronous dataflow specification accompanied by synchronization constraints. In particular, SynD [10] can be extended using our results.

### 3.1 Process structure

A Signal process is a *hierarchical* specification corresponding to the configuration of a system.

A process is formed of a header, an interface specification, and local declarations. **EXAMPLE.** Its interface defines

its *interface specification*. Fig. 4 gives the Signal specification of the system in Fig. 1.

Its name, an interface specification, and local declarations. In our example, the top process is defined as follows: `(SYNC1, I1, O1) <- ADD1+CTRL1`

```

1 process EXAMPLE =
2   (? event SYNC ; event I1,I2 ;
3   ! event O1,O2 ; )
4   ( | (O1,C) := ADD1 (SYNC,I1)
5     | O2 := ADD2 (SYNC,I2,C) | )
6   where boolean C ;
7
8   process ADD1 =
9     (? event SYNC ; event I1 ;
10    ! event O1 ; event C ; )
11    ( | I1 ^= O1 | SYNC ^< I1 | C ^= SYNC | ) ;
12
13   process ADD2 =
14     (? event SYNC ; event I2 ; event C ;
15    ! event O2 ; )
16    ( | I2 ^= O2 | SYNC ^< I2 | C ^= SYNC | ) ;
17 end ;

```

Figure 4: The Signal process of the configurable adder in Fig. 1

its set of input signals, with  $\mathcal{O}(P)$  its set of output signals, and with  $\mathcal{V}(P)$  the set of all its signals. By definition,  $\mathcal{I}(P) \cap \mathcal{O}(P) = \emptyset$  and  $\mathcal{I}(P) \cup \mathcal{O}(P) \subseteq \mathcal{V}(P)$ .

The dataflow specification of `EXAMPLE` consists of two equations, which define the interconnections between `ADD1`, `ADD2`, and the environment. The local definition section defines the internal signal `C`, and the processes `ADD1` and `ADD2`.

Our synthesis technique will only refer directly to the first two levels of the process hierarchy. The unique top-level process is the specification, and it only contains the global dataflow. Its children are the specifications of the synchronous modules, which can contain more complex behavioral Signal definitions. Lower-level processes cannot be distinguished from the specification of the module that contains them.

### 3.1.1 Finite abstraction

As explained in section Section 2.2, complex data operations that are irrelevant to the synthesis of the shells should be abstracted away in the synchronous specification. For instance, we model the integer inputs and outputs of the two adders with `event` signals. Data abstraction is mandatory for all infinite types (e.g., `integer`, `float`), to make decidable the analysis of the weak endochrony. We say that the synchronous specification is a *finite stateless abstraction* of the corresponding pearl behavior.

### 3.2 Dataflow specification

The dataflow specification of a process is formed of *equations* defining *constraints* between the signals of the process. Any reaction satisfying all the equations of a process  $P$  is a valid reaction of  $P$ . We denote with  $\mathcal{R}(P)$  the set of all the reactions of  $P$ . Some of the constraints (such as as si n e t) a e f

```

process EXAMPLE_RESULT =
(? boolean SYNC1,SYNC2; event I1,I2 ;
 ! event O1,O2 ; )
(| (O1,C) := CTRL1 (SYNC1,I1)
 | O2 := CTRL2 (SYNC2,I2,C) |)
where boolean C ;

process CTRL1 =
(? boolean SYNC1 ; event I1 ;
 ! event O1 ; event C ; )
(| SYNC1 ^= I1 ^= O1 | C ^= when SYNC1=1
 | (O1,C) := ADD1(when SYNC1=1,I1) |)
where process ADD1 = ... end ;

process CTRL2 =
(? boolean SYNC2 ; event I2 ; event C ;
 ! event O2 ; )
(| SYNC2 ^= I2 ^= O2 | C ^= when SYNC2=1
 | O2 := ADD2(when SYNC2=1,I2,C) |)
where process ADD2 = ... end ;
end ;

```

Figure 5: Weakly endochronous process corresponding to Fig. 3. Processes ADD1 and ADD2 are copied from Fig. 4.

The operator `default` merges two signals of the same type, giving priority to the first. The signal “X default Y” is present whenever one of X or Y is present. It is equal to X whenever X is present, and is equal to Y otherwise.

A Signal process can instantiate other Signal processes inside its dataflow. In Fig. 4 EXAMPLE instantiates ADD1 and ADD2. In Fig. 5 CTRL1 instantiates ADD1.

### 3.3 Notations

A synchronous trace of a process  $P$  is any finite succession of reactions of  $P$ . The set of traces of  $P$  is  $T(P) = \mathcal{R}(P)^*$ .

On the values of any signal  $s$ , we introduce a partial order, given by  $\perp \leq$ , for all  $v \in \mathcal{D}_S$ . The order on signal values induces a product partial order  $\leq$  on reactions. Note that  $r_1 \leq r_2$  if and only if  $s_{\downarrow}pp(r_1) \subseteq s_{\downarrow}pp(r_2)$  and  $r_1(\sigma) = r_2(\sigma)$  for all  $\sigma \in s_{\downarrow}pp(r_1)$ .

We say of two reactions  $r_1$  and  $r_2$  that they are *non-contradictory*, written  $r_1 \not\sim r_2$ , if  $r_1(\sigma) = r_2(\sigma)$  for all  $\sigma \in s_{\downarrow}pp(r_1) \cap s_{\downarrow}pp(r_2)$ . Otherwise, we say that the reactions are *contradictory*, written  $r_1 \sim r_2$ . Given a set of reactions  $\mathcal{R}$ , we shall say that it is non-contradictory, denoted  $\mathcal{R} \not\sim$  if any two reactions of  $\mathcal{R}$  are non-contradictory. In this case,

we can define  $\vee = \bigvee_{r \in K} r$ . Two reactions  $r_1$  and  $r_2$  are *non-interferent* if  $s_{\blacktriangledown}pp(r_1) \cap s_{\blacktriangledown}pp(r_2) = \emptyset$ .

On non-contradictory reactions, we define the union  $\vee$  and intersection  $\wedge$  operators as the least upper bound and greatest lower bound induced by  $\leq$ . We also define the difference  $r_1 \setminus r_2$ , which has support  $s_{\blacktriangledown}pp(r_1) \setminus s_{\blacktriangledown}pp(r_2)$  and equals  $r_1$  on its support.

## 4 Weak endochrony

### 4.1 Basic theory

The theory of *weakly endochronous (WE) systems* [16], gives criteria establishing that a synchronous presentation hides a behavior that is fundamentally asynchronous and deterministic. Absence information is not needed, which guarantees the deterministic implementability of the synchronous specification in an asynchronous environment. Weak endochrony extends to a synchronous framework the Mazurkiewicz traces [12].

Weak endochrony is defined in an automata-theoretic framework. We simplify it here according to our stateless abstraction:

**Definition 1 (stateless weak endochrony)** *We say that process  $P$  is weakly endochronous if for all  $r_1, r_2 \in \mathcal{R}(P)$*

$$r_1 \setminus r_2 \Rightarrow r_1 \vee r_2 \setminus r_2 \wedge r_1 \setminus r_2 \setminus r_1 \in \mathcal{R}(P)$$

We already saw, in Section 2.4, that running a synchronous process in an asynchronous environment may be a non-deterministic process. Weak endochrony gives very general sufficient conditions ensuring that the execution of the process always gives the same result, regardless of the asynchronous input arrival order.

The intuition behind weak endochrony is that we are looking for systems where (1) all causality is implied by the sequencing of messages on communication channels, and (2) all choices are visible as choices over the value (and not present/absent status) of some message. As explained in [15], the axioms of weak endochrony can be traced down to the fundamental result of Keller [11] on the deterministic operation of a system in an asynchronous environment. From a different perspective, Weakly Endochronous Systems (WES) are similar to deterministic results.

#### 4.1.1 Atoms

From our point of view oriented towards code generation, an interesting consequence of these axioms is that any behavior of a WES is the smallest<sup>3</sup> in  $\mathcal{R}(P)$  different from

<sup>3</sup>In the sense of  $\leq$ .



Atomic transitions have very nice properties. Two atoms  $a_1$  and  $a_2$  are either non-interferent (of disjoint support, cf. Section 3.3), or contradictory. Furthermore, any transition of the system can be uniquely decomposed into a set of non-interferent atoms.

**Theorem 1 (Generation)** *Let  $P$  be a weakly endochronous stateless process, and  $r \in \mathcal{R}(P)$ . Then:*

$$r = \bigvee_{a \in \text{Atoms}(P), a \leq r} a$$

## 4.2 Checking weak endochrony

In this section we explain how to check the weak endochrony of the process corresponding to a synchronous module of the GALs system. The analysis is based on the identification of a set of reactions which generate by union all other reactions. The generator set is computed incrementally from the equations of the specification. The process is weakly endochronous if and only if the generator set has the properties of an atom set. If the process is weakly endochronous, we will directly refer to Section 5 to generate the asynchronous

### 4.2.1 Partial reactions

The domain of signal valuations, defined in Section 2.2.1, allows the representation of complete reactions of processes, as well as the manipulation of reactions of weakly endochronous processes, where absence information is not needed. However, they do not represent absence synchronizations constraints, such as clock exclusions. To handle such constraints, we enrich the domain  $\mathcal{D}_S^\perp$  of each signal with a new value  $\perp$ .  $\mathcal{D}_S^\perp = \mathcal{D}_S^\perp \cup \{\perp\}$ , with  $\perp \leq \perp$ .

A partial reaction  $r$  is any valuation of the signals or their extensions. We denote with  $\mathcal{R}^\perp$  the set of all such valuations. All the operators of Section 2.2.1 are extended to  $\mathcal{R}^\perp$ , according to the new domain structure. The support of a partial reaction is the set of signals with value different from  $\perp$ .

A partial reaction  $r$  sets signal  $s$  to  $\perp$  to represent the fact that  $s$  is forced to be absent. For instance, if “S^#T” and  $r(s) \in \mathcal{D}_S$ , we need to set  $r(s) = \perp$ . A partial reaction  $(v \perp)$  cannot be united (using  $\vee$ ) with other partial reactions if  $v(s) \neq \perp$  is present. This is different from  $(v \perp)$ , which can be composed, for instance, with  $(\perp v)$ .

Any reaction is a partial reaction, and any partial reaction  $r$  has an associated reaction  $[r]$  which changes all  $\perp$  values of  $r$  into  $\perp$  values. We denote with  $\mathcal{R}^\perp$  the set of partial reactions  $\{r \in \mathcal{R}(P)\}$ .

### 4.2.2 Generators

In this section, we define the non-interferent generator set of a process, and how to compute it. Non-interferent generators of the reactions of a process can be seen as the prime implicants of a logic formula – they are reactions of smallest support

all other reactions. The main difference is that our algebraic structure is richer than that of the Boolean logic (it has a domain structure), allowing us to exploit concurrency to preserve fewer generators (exponentially fewer, in certain cases). Also, it must be noted that  $\perp$  is always a solution, in our case.

**Definition 2 (Generator set)** *Let  $P$  be a process. A set  $C \subseteq \mathcal{R}^\perp(P)$  of partial reactions is a generator set of  $\mathcal{R}(P)$  if  $\mathcal{R}(P) = \{[\bigvee_{r \in K} r] \mid K \subseteq C \wedge \dots\}$ .*

**Definition 3 (Non-interferent generator set)** *Let  $P$  be a process. A generator set  $C$  is called a non-interferent generator set of  $P$  if:*

$$\forall r_1, r_2 \in C : \left. \begin{array}{l} r_1 \quad r_2 \\ r_1 \neq r_2 \end{array} \right\} \Rightarrow s_{\mathbf{v}}pp([r_1]) \cap s_{\mathbf{v}}pp([r_2]) = \emptyset \quad (\mathcal{NI})$$

Non-interference implies a strong form of minimality for a generator set, also facilitating the manipulation of the generators.

**Theorem 2** *Any process has a unique non-interferent generator set. When the process is weakly endochronous, it is  $\mathbf{O}$ ,  $\mathbf{MD}$ ,  $\mathbf{F}$ ,  $\mathbf{F}$ ,  $\mathbf{F}$ ,  $\mathbf{FF}$ ,  $\mathbf{I}$ ,  $\mathbf{y}$ .*



We also give here generators sets of other, non-primitive equations:

$$\begin{aligned}
\mathcal{G}(P \text{ x} := Y) &= \{(X^v \ Y^v) \mid v \in \mathcal{D}_X \cup \{\perp\}\} \cup \\
&\quad \{(\dot{W}^v) \mid v \in \mathcal{D}_X \ \dot{W} \notin \{X \ Y\}\} \\
\mathcal{G}(P \text{ x}^\wedge = Y) &= \{(X^v \ Y^w) \mid w \in \mathcal{D}_X\} \cup \{(X^\perp \ Y^\perp)\} \cup \\
&\quad \{(\dot{W}^v) \mid v \in \mathcal{D}_X \ \dot{W} \notin \{X \ Y\}\} \\
\mathcal{G}(P \text{ x}^\wedge \# Y) &= \{(X^v \ Y^\perp) \ (X^\perp \ Y^v) \mid v \in \mathcal{D}_X\} \cup \\
&\quad \{(\dot{W}^v) \mid v \in \mathcal{D}_X \ \dot{W} \notin \{X \ Y\}\} \\
\mathcal{G}(P \text{ x}^\wedge \ Y) &= \{(X^v \ Y^w) \mid v \in \mathcal{D}_X \cup \{\perp\} \ w \in \mathcal{D}_Y\} \cup \\
&\quad \{(\dot{W}^v) \mid v \in \mathcal{D}_X \ \dot{W} \notin \{X \ Y\}\}
\end{aligned}$$

**Theorem 3 (Merging)** *Let  $P$  be a process and  $p$  and  $q$  be sets of equations of  $P$ . Then:*

1. *If  $p$  is a single equation, then  $\mathcal{G}(P \ p)$  defined above give the non-interferent generator set of  $p$  in  $P$ .*
2. *Procedure 1 (MergeGeneratorSets) computes  $\mathcal{G}(P \ p \mid q)$ , the unique non-interferent generator set of  $p \mid q$  in  $P$ .*

Procedure 1 (MergeGeneratorSets) computes in  $\mathcal{G}(P \ p \mid q)$  all the minimal reactions of  $p \mid q$ . It explores every minimal combination of compatible atoms in  $p$  and

---

**Procedure 1 MergeGeneratorSets**


---

**Input:**  $\mathcal{G}(P \ p)$ ,  $\mathcal{G}(P \ q)$ : reaction set

**Output:**  $\mathcal{G}(P \ p \ | \ q)$ : reaction set

$\mathcal{G}_p \leftarrow \emptyset$ ;  $\mathcal{G}'_p \leftarrow \emptyset$ ;

**for all**  $g \in \mathcal{G}(P \ p)$  **do**

**if**  $[g] = \perp$  **then**  $\mathcal{G}'_p \leftarrow \mathcal{G}'_p \cup \{g\}$

### 4.2.3 Weak endochrony check

A process is weakly endochronous as soon as we do not need the forced absence information to distinguish between elements of  $\mathcal{G}(P \ p)$ .

**Theorem 4 (Weak endochrony)** *Let  $P$  be a process and let  $p$  be the set of its equations. Then,  $P$  is WE iff for all  $p \ q \in \mathcal{G}(P \ p)$ ,  $p \neq q$ :*

$$\left. \begin{array}{l} s_{\downarrow}pp(p) \cap s_{\downarrow}pp(q) \neq \emptyset \\ [p] [q] \neq \perp \end{array} \right\} \Rightarrow \exists \ : \mathcal{D}_S \ni p(\ ) \neq q(\ ) \in \mathcal{D}_S \quad (\text{WE} - \text{nos})$$

In this case, the atom set of  $P$  is:

$$A_{\text{nos}}(P) = \{[g] \mid g \in \mathcal{G}(P \ p) \ [g] \neq \perp\}$$

**Proof sketch:** If  $P$  is WE, then for any reaction, and therefore generator, we can build the corresponding generator (with the needed properties).

If  $\mathcal{G}(P \ p)$  has property (WE - nos) and the properties of generators imply we

Checking the weak endochrony of a system is not weakly endochronous, in fact  $\in \mathcal{G}(P \ p)$ ,  $\neq$ , with  $s_{\downarrow}pp(\ ) \cap s_{\downarrow}pp(\ ) = \emptyset$  and  $[\ ]$  and  $[\ ]$  are reactions of  $P$  that are not asynchronous environment. The user must be contradictory, for instance by addition of Automatic synthesis of weak endochrony is

### 4.2.4 Examples

The generator set of the process in Fig. 4

$$\{(1^\top O1^\top C^\perp YNC^\perp), (1^\top O1^\top 2^\top C^\perp YNC^\perp)\}$$

As expected, the process is not weakly endochronous and  $[(1^\top O1^\top YNC^\perp)]$  are neither comparable.

The generator set of the transformed process

$$\{(1^\top O1^\top C^\perp YNC2^0), (1^\top O1^\top 2^\top O2^\top YNC2^1)\}$$

The process is weakly endochronous.

## 5 Executive generation

Once we determined that a process is weakly endochronous we can generate for it the shell (executive) described in Section 2. We shall present here a simple technique for generating the shell directly from the atom set.

To further simplify code generation, we also assume here that the weakly endochronous process  $P$  describing our shell is I/O deterministic.<sup>4</sup> Formally, we shall assume that for any two different  $r \in \text{Atos}(P)$  we have:

$$r \Rightarrow \mathcal{I}(P) / \mathcal{I}(P)$$

where  $r \mathcal{I}(P)$  is the restriction of  $r$  on the input signals.

The executive we generate for  $P$  takes advantage of atom independence properties. For each atom  $a$ , independently, we generate a function that cyclically:

1. Awaits the arrival of a new value  $v(a)$  for each signal  $a \in \text{supp}(\mathcal{I}(P))$ .
2. When all needed input arrived, trigger a reaction of the pearl, with  $\mathcal{I}(P)$  as input, and then output the values of  $\mathcal{O}(P)$ .

The actual executive is formed of all the functions associated with the atoms, running in parallel. In practice, parallelism can be simulated using a simple event-driven approach, where each incoming event  $v \in \mathcal{D}_S$  triggers computations in the functions associated with the atoms having  $v = \cdot$ .

The actions of step 2 above are complex, and must be executed atomically. In a purely asynchronous shell, no two atoms must be executing step 2 at the same time (it is a mutual exclusion region). The operations performed during step 2 are the following:

1. [Redacted]
2. [Redacted]
3. Trigger the completion of the pearl by enabling the clock (in hardware) or calling the reaction function (in software).
4. When the reaction is completed: in hardware, disable the clock; in software, update the pearl state.
5. Send the data to the output ports. This may block the computation until the communication lines are free.
6. Signal the fact that the input  $\mathcal{I}(P)$  has been read. Remove these signals from the await list of other atoms (new signals can arrive on the input ports of  $P$ ).
7. [Redacted]

<sup>4</sup>In the literature, this property is often referred to as "I/O determinism" in the specification.



## 6.1 Causality and completeness

In our approach, the causality of computations is determined by the pearl reactions, and by the ability of the shells to control the pearls. In a reaction of the pearl, all input must come before the computation takes place (clock enable/disable cycle in hardware, reaction function call in software), and before any output can be produced. The finest reactions, which give the causal dependencies, are the atoms.

By consequence, some reactions of the synchronous model may be unfeasible in the GALS implementation. Consider, for instance, the GALS system synthesized from the following specification:

```
process EXAMPLE2 =
  (? event I ; ! event O ;)
  ( | (A,B) := P1(I) | O:=P2(A,B) | )
where event A,B ;
  process P1 =
    (? event I,A ; ! event B ;) ( | A ^= B ^= I | );
  process P2 =
    (? event A ; ! event B,O ;) ( | A ^= B ^= O | );
end ;
```

Module  $m_1$  has exactly one atom:  $\alpha_1 =$  Asynchronous reaction of EXAMPLE2.  $m_1$  cannot perform it, because the dependency on  $B$  requires the input  $B$  that  $\alpha_1$  is executed before the dependency on  $B$  is satisfied.

We need to ensure that the GALS implementation indeed implements the functions of the synchronous model. To do this, we need to ensure that the system is acyclic in all possible configurations. This problem is similar to the problem of compilation in all possible configurations. This problem is similar to the problem of compilation in all possible configurations. This problem is similar to the problem of compilation in all possible configurations.

In our case, the problem is ensuring the correctness of the atom system. More complex analyses can be used to ensure the correctness of the atom system. More complex analyses can be used to ensure the correctness of the atom system.

## 6.2 Weak

The weak end-to-end causality ensures deterministic behavior of the system. The specified behaviors of the system are feasible. We ensure now that the behaviors are feasible than the specified behaviors. More precisely, we ensure that the synchronous behaviors are feasible than the specified behaviors.

This is not the case. Consider, for instance, the GALS system synthesized from the following specification:



```

process EXAMPLE2 =
  (? boolean I ; ! event 0 ;)
  ( | (A,B) := PROD(I) | 0:=CONS(A,B) | )
where event A,B ;
  process PROD =
    (? boolean I ; ! event A,B ;)
    ( | A:=when I=0 | B:=when I=1 | );
  process CONS =
    (? event A,B ;) ( | A^=B | ) ;
end ;

```

The atoms of PROD are  $(^0 A^\top)$  and  $(^1 B^\top)$ . The only atom of CONS is  $(A^\top B^\top)$ . The only behavior of process EXAMPLE2 is  $\perp$ . Therefore, it has no atom.

However, if PROD and CONS are compiled using the technique of Section 5, the resulting GALIS system has non-oid asynchronous behaviors. For instance, after reading  $^0$  and  $^1$ , it produces  $O^\top$ . This is due to the fact that the events  $A^\top$  and  $B^\top$  are produced in different reactions of PROD, but read by CONS in a single reaction. The reverse can also happen, where signals produced synchronously are read in different instants.

The criterion ensuring that such stray behaviors cannot occur is *weak isochrony*. After a few notations, we give the basic definition of Potop, Caillaud, and Beneniste [16], as simplified by our stateless framework.

**Definition 4 (asynchronous prefixes)** *Let  $P$  be a process. Then,  $\text{pref}_d(P)$  is the set of all  $\mu \in \mathcal{R}$  such that there exists  $\nu \in T(P)$  with  $(\mu, \nu)$  being the first value on signal in  $\mu$ , or  $\perp$ , if no such value exists.*

In other terms,  $\text{pref}_d(P_i)$  is the set of all asynchronous prefixes of depth 1 of executions of  $P$ . With this notation:

**Definition 5 (weak isochrony)** *Let  $P$  be a process with sub-processes  $P_1 \dots P_n$ . We say that the processes  $P_1 \dots P_n$  are weakly isochronous if, by definition, any execution of the GALIS implementation can be started with a synchronous reaction.*

*Formally, for all  $r_i \in \text{pref}_d(P_i)$ , such that any two  $r_i, r_j$  are equal on  $\mathcal{V}(P_i) \cap \mathcal{V}(P_j)$  we have:*

$$\exists \bar{r}_i \in \mathcal{R}(P_i) \ 1 \leq i \leq n : \begin{cases} \bar{r}_i \leq r_i \\ \bigvee_{i=1}^n \bar{r}_i \in \mathcal{R}(P) \\ \bigvee_{i=1}^n \bar{r}_i \neq \perp \end{cases}$$

Checking weak isochrony in its basic form is decidable, but can be very complex, due to the computation of  $\text{pref}_d(P_i)$ . We propose here a sufficient property adapted to our atom-based approach.

**Theorem 5 (sufficient conditions)** *Let  $P$  be a process with sub-processes  $P_1 \dots P_n$ . Then, a sufficient condition for  $P_1 \dots P_n$  to be isochronous is that the input of no atom of a component can be covered by atoms of other components, unless these atoms are independent.*

*Formally:* We denote with  $\mathcal{A}_i$  the union of all  $\text{Atoms}(P_j)$  with  $j \neq i$ . Then,  $P_1, \dots, P_n$  are isochronous if: For all  $1 \leq i \leq n$ , and all  $\alpha \in \text{Atoms}(P_i)$ , if  $\alpha_1, \dots, \alpha_k \in \mathcal{A}_i$  such that for all  $\beta \in \text{support}(\mathcal{I}(P))$  there exists  $j$  with  $\beta(j) = \alpha$ , then  $\text{support}(\alpha_i)$  are mutually disjoint.

**Proof sketch:** Consider  $r_i \in \text{Atoms}(P_i)$   $1 \leq i \leq n$  satisfying the hypothesis of Definition 5. There exists  $\alpha$  and  $\beta \in \text{Atoms}(P_i)$  such that  $\beta \leq r_i$ . Using the hypothesis of the theorem, we can incrementally build a common reaction by appending atoms to one or the other of the  $P_j$ .  $\diamond$

Note that the criterion offered by Theorem 5 offers a simple computational criterion based on the atoms already computed in previous sections. Determining isochrony is done through a simple traversal of atom co-ocurrences. The criterion is an over-approximation of weak isochrony. It can be improved, for instance by taking into account causality information.

## 7 Conclusion

We have defined a method for synthesizing the asynchronous executions that are driving the synchronous modules of a (GALS) system. The technique takes as input high-level synchronization constraints. The resulting GALS system is predictable and functionally correct and complete with respect to the initial synchronous specification, and regardless of the size of the communication lines. The technique allows the synthesis of executions for all specifications whose modules are stateless weakly endochronous.

**Future work.** The current paper only concerns correctness in an untimed setting. Our long-term objective is to take into account and/or guarantee real-time requirements, such as periodicity, end-to-end, or throughput constraints. This involves the definition of timing analysis and scheduling techniques compatible with our executions. On the other hand, the executions could be simplified under specific timing hypothesis (for instance, the FIFO protocols can be simplified if the reader is faster than the writer).

The various steps of our technique can be improved and extended in a variety of ways, already mentioned in Sections 5 and 6.

## References

- [1] P. Amagbégnon, L. Besnard, and P. L. Guernic. Implementation of the data-flow synchronous language signal. In *Proceedings PLDI'95*, La Jolla, CA, USA, June 1995.
- [2] A. Ben eniste, B. Caillaud, and P. L. Guernic. Compositionality in dataflow synchronous languages: Specification and distributed code generation. *Information and Computation*, 163:125 – 171, 2000.
- [3] A. Ben eniste, P. Caspi, S. A. ... ne. The synchronous languages 15 ... 2003.

- 
- [4] G. Butazzo. *Hard real-time computing systems. Predictable scheduling, algorithms and applications*. Kluwer, 2002.
  - [5] L. Carloni, K. McMillan, and A. Sangioanni-Vincentelli. Theory of latency-insensitive design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 20(9):18, Sep 2001.
  - [6] P. Caspi and M. Pouzet. Synchronous Kahn networks. In *International Conference on Functional Programming*, pages 226–238, 1996.
  - [7] T. Grandpierre and J.-F. Sorel. From algorithm and architecture specification to auto-



---

Unité de recherche INRIA Rocquencourt  
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes  
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique  
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

---

Éditeur  
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)  
<http://www.inria.fr>  
ISSN 0249-6399