



HAL
open science

Compositionality in dataflow synchronous languages: specification & distributed code generation

Albert Benveniste, Benoit Caillaud, Paul Le Guernic

► **To cite this version:**

Albert Benveniste, Benoit Caillaud, Paul Le Guernic. Compositionality in dataflow synchronous languages: specification & distributed code generation. *Information and Computation*, 2000, 163 (1), pp.125-171. 10.1006/inco.2000.9999 . hal-00543297

HAL Id: hal-00543297

<https://hal.science/hal-00543297>

Submitted on 6 Dec 2010

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.

Compositionality in dataflow synchronous languages : specification & distributed code generation ^{*†‡}

Albert Benveniste Benoît Caillaud Paul Le Guernic [§]

October 16, 2001

^{*}This paper is a significantly revised version of a preliminary report which appeared under the same title in the Proceedings of 1997 Malente Workshop on Compositionality, organized by W.P. de Roever and H. Langmaack; these proceedings will be published in the LNCS, Springer Verlag.

[†]This work is or has been supported in part by the following projects: Eureka-SYNCHRON, Esprit R&D -SACRES (Esprit project EP 20897), Esprit LTR-SYRF (Esprit project EP 22703).

[‡]In addition to the listed authors, the following people have indirectly, but strongly, contributed to this work: the STS formalism has been shamelessly borrowed from Amir Pnueli, the background on labelled partial orders is mostly acknowledged to Paul Caspi.

[§]Irisa/Inria, Campus de Beaulieu, 35042 Rennes cedex, France; email: first-name.lastname@irisa.fr

Abstract

Modularity is advocated as a solution for the design of large systems, the mathematical translation of this concept is often that of *compositionality*. This paper is devoted to the issues of compositionality for modular code generation, in dataflow synchronous languages.

As careless reuse of object code in new or evolving system designs fails to work, we first concentrate on what are the additional features needed to abstract programs for the purpose of code generation: we show that a central notion is that of *scheduling specification* as resulting from a *causality analysis* of the given program. Using this notion, we study separate compilation for synchronous programs. An entire section is devoted to the formal study of causality and scheduling specifications.

Then we discuss the issue of distributed implementation using an asynchronous medium of communication. Our main results are that it is possible to characterize those synchronous programs which can be distributed on an asynchronous architecture without losing semantic properties. Two new notions of *endochrony* and *isochrony* are introduced for this purpose. As a result, we derive a theory for synthesizing additional schedulers and protocols needed to guarantee the correctness of distributed code generation.

Corresponding algorithms are implemented in the framework of the DC+ common format for synchronous languages, and the V4-release of the SIGNAL language.

Keywords : synchronous languages, modularity, distributed code generation, separate compilation, desynchronization.

Contents

1	Rationale	4
2	Specification	5
2.1	The essentials of the synchronous paradigm	6
2.2	Synchronous Transition Systems (STS)	7
3	Compositionality in code generation: informal analysis	11
3.1	What is the problem?	11
3.2	Scheduling specifications	13
3.3	Causality analysis: examples	18
3.4	Generating scheduling for separate modules	20
3.5	Relaxing synchrony	22
3.6	Modular design, GALS architectures	24
4	Formal study of desynchronization	25
4.1	Desynchronizing STS, and two fundamental problems	26
4.2	Endochrony and re-synchronization	29
4.2.1	Formal results	29
4.2.2	Practical consequences	33
4.3	Isochrony, and synchronous and <i>asynchronous</i> compositions	34
4.4	Getting GALS architectures	42
4.5	Handling endo/isochrony in practice	42
4.5.1	Checking endo/isochrony	43
4.5.2	Enforcing endo/isochrony	44
5	Formal study of causality	45
5.1	Encoding scheduling specifications using an algebraic domain	46
5.2	Circuitfree schedulings	47
5.3	Deriving scheduling specifications as causality constraints	51
5.4	Correct programs	54
6	Conclusion	57

1 Rationale

Modularity is advocated as the ultimate solution for the design of large systems, and this holds in particular for embedded systems, for both software and architecture. Modularity allows the designer to scale down design problems, and facilitates the reuse of pre-existing modules.

The mathematical translation of the concept of modularity is often that of *compositionality*. Paying attention to the composition of specifications [Manna and Pnueli 1992] is central to any system model involving concurrency or parallelism. More recently, significant effort has been devoted toward the introduction of compositionality in verification, which aims at deriving proofs of large programs from partial proofs involving (abstractions of) components [Manna and Pnueli 1995]. See also the whole volume [de Roever et al., Eds, 1998] where a number of papers are devoted to this topic.

Compilation and code generation has been given less attention from this very same point of view. This is unfortunate, as it is critical for the designer to scale down the design of large systems by 1/ storing modules like black-box “procedures” or “processes” with minimal interface description, and 2/ generating code which uses these modules only on the basis of their interface description, while preserving in any case the correctness of the design. This paper is devoted to the issues of compositionality of dataflow synchronous languages, aimed at modular code generation.

Dataflow synchrony is rather a paradigm than a set of concrete languages or visual formalisms [Benveniste and Berry, 1991], hence it is desirable to abstract from such and such particular language. Thus we have chosen to work with Synchronous Transition Systems (STS), a lightweight formalism proposed by Amir Pnueli, general enough to capture the essence of the synchronous paradigm. This is the topic of section 2. Using this formalism, we study in section 2 the composition of specifications.

Most of our effort is then devoted to issues of compositionality that are critical to code generation. Section 3 contains an informal discussion of this problem. It is known that careless storing of object code for further reuse in systems design fails to work. Hence we first concentrate on the additional features that are required to abstract programs for the purpose of code generation and reuse: we show that a central notion is that of scheduling specification as resulting from a causality analysis of the given program. Related issues of compositionality are investigated. Then we show that there

is some appropriate level of “intermediate code”, which at the same time allows us to scale down code generation for large systems, and still maintains correctness at the system integration phase. Finally we discuss the side issue of distributed implementation using an asynchronous medium of communication.

In section 4 we formally study desynchronization. We first formalize what we mean by desynchronization. Our theory requires that the communication medium or operating system: 1/ shall not lose messages, and 2/ shall preserve the total ordering of messages, *for each flow individually* (but, of course, not globally). These assumptions are typically satisfied by services offered by reliable communication media or operating system. Our main result is that *it is possible to check, directly on the original synchronous specification, whether semantic properties will or will not be preserved after desynchronization*. The two fundamental notions are *endochrony*, which guarantees that, for a single STS, desynchronization is a “reversible” transformation, and *isochrony*, which guarantees that, for a pair of STS, desynchronizing communications is also a “reversible” transformation. In some sense formalized in section 4, semantics is preserved by desynchronization when these conditions are satisfied.

Then section 5 is devoted to a formal study of causality. In many respects, this formal study is important. First, it is instrumental in getting executable, deterministic code from a given STS specification. Then, it is a cornerstone of proper abstractions for separate compilation and reuse. We pay strong attention to this study, using a technique not unlike the one used for analyzing causality in Esterel [Berry, 1995]. Our analysis encompasses the case of arbitrary data types, and suitable abstractions are used for this purpose.

In the conclusion we discuss how our views on compositionality are modified by this study. We sketch the resulting system design methodology, and we briefly mention the implementation resulting from this theory, mostly developed in the framework of the Esprit-SACRES project.

2 Specification

This section discusses compositionality aspects of specifications, first informally, and then formally.

2.1 The essentials of the synchronous paradigm

There have been several attempts to characterize the essentials of the synchronous paradigm [Berry, 1989] [Benveniste and Berry, 1991] [Halbwachs, 1993]. With some experience, we feel that the following features are indeed essential and sufficient for characterizing this paradigm :

1. Programs progress via an infinite sequence of *reactions*, informally written :

$$P = R^\omega$$

where R denotes the set of legal reactions¹.

2. Within a reaction, decisions can be taken on the basis of the *absence* of some events, as exemplified by the following typical statements, taken from ESTEREL, LUSTRE, and SIGNAL respectively :

```
present S else 'stat'  
y = current x  
y := u default v
```

The first statement is self-explanatory. The “**current**” operator delivers the most recent value of x at the clock of the considered node, it thus has to test for absence of x before producing y . The “**default**” operator delivers its first argument when it is present, and otherwise its second argument.

3. Communication is performed via instantaneous broadcast. In other words, when it is defined, parallel composition is always given by the conjunction of associated reactions :

$$P_1 || P_2 = (R_1 \wedge R_2)^\omega$$

The above formula is a perfect definition of parallel composition when the intention is specifying. In contrast, if producing executable code was the intention, then this definition has to be compatible with an operational semantics. This very much complicates the “when it is defined” prerequisite².

¹In fact, “reaction” is a slightly restrictive term, as we shall see in the sequel that “reacting to the environment” is not the only possible kind of interaction a synchronous system may have with its environment.

²For instance, most of the effort related to the semantics of ESTEREL has been directed toward solving this issue satisfactorily [Berry, 1995].

Of course, such a characterization of the synchronous paradigm makes the class of “synchrony-compliant” formalisms much larger than usually considered. However it has been our experience that these were the key features of the techniques we have developed so far.

Clearly, this calls for the simplest possible formalism comprizing the above features, and on which fundamental questions should be investigated. This is one of the objectives of the STS formalism described next.

2.2 Synchronous Transition Systems (STS)

Synchronous Transition Systems (STS).

We assume a vocabulary \mathcal{V} which is a set of typed variables. All types are implicitly extended with a special element \perp to be interpreted as “absent”. Some of the types we consider are the type of *pure signals* with domain $\{\mathsf{T}\}$, and *booleans* with domain $\{\mathsf{T}, \mathsf{F}\}$ (recall both types are extended with the distinguished element \perp).

We define a *state* s to be a type-consistent interpretation of \mathcal{V} , assigning to each variable v a value $s[v]$ over its domain. We denote by S the set of all states. For a subset of variables $V \subseteq \mathcal{V}$, we define a V -state to be a type-consistent interpretation of V .

We define a *Synchronous Transition System* (STS) to be a triple

$$\Phi = \langle V, \Theta, \rho \rangle$$

consisting of the following components :

- V is a finite set of typed *variables*,
- Θ is an assertion characterizing the set of *initial states*: $\{s \mid s \models \Theta\}$.
- $\rho \subseteq S \times S$ is the *transition relation* relating past and current states denoted by s^- and s respectively³. For example the assertion $x = x^- + 1$ states that the value of x in s is greater by 1 than its value in s^- . If $(s^-, s) \models \rho$, we say that state s^- is a ρ -*predecessor* of state s .

³Usually, states and *primed* states are used to refer to current and *next* states. This is equivalent to our present notation. We have preferred to consider s^- and s , just because the formulas we shall write mostly involve current variables, rather than past ones. Using the standard notation would have resulted in a burden of primed variables in the formulas.

Runs.

A *run* $\sigma : s_0, s_1, s_2, \dots$ is a sequence of states such that

$$s_0 \models \Theta \quad \bigwedge \quad \forall i > 0, (s_{i-1}, s_i) \models \rho \quad (1)$$

Composition.

The *composition* of two STS $\Phi = \Phi_1 \parallel \Phi_2$ is defined as follows :

$$\begin{aligned} V &= V_1 \cup V_2 \\ \Theta &= \Theta_1 \wedge \Theta_2 \\ \rho &= \rho_1 \wedge \rho_2, \end{aligned}$$

the composition is thus the pairwise conjunction (denoted by \wedge) of initial and transition relations. Composition is thus commutative and associative. Note that, in STS composition, interaction occurs through common variables only.

Notations for STS.

For the convenience of specification, STS have a set of *declared* variables, written $V_{\mathbf{d}}$, implicitly augmented with associated *auxiliary* variables : the whole constitutes the set V of variables. We shall use the following generic notations in the sequel :

- b, c, v, w, \dots denote STS declared variables, and b, c are used to refer to variables of boolean type.
- for v a declared variable, $h_v \in \{\top, \perp\}$ denotes its *clock* :

$$[h_v \neq \perp] \Leftrightarrow [v \neq \perp]$$

- for v a declared variable, ξ_v denotes its associated *state-variable*, defined by :

$$\begin{aligned} \mathbf{if} \quad h_v \quad \mathbf{then} \quad & \xi_v = v \\ & \mathbf{else} \quad \xi_v = \xi_v^- \end{aligned} \quad (2)$$

Values can be given to $s_0[\xi_v]$ as part of the initial condition. Then, ξ_v is always present after the 1st occurrence of v . Note that $\xi_{\xi_v} = \xi_v$, thus only state variables of declared variables have to be considered.

Stuttering.

As modularity is desirable, an STS should be permitted to do nothing while its environment is possibly working. This feature has been yet identified in the literature and is known as *stuttering* invariance or robustness [Lamport, 1983a, Lamport, 1983b]. Stuttering invariance of an STS Φ is defined as follows: if

$$\sigma : s_0, s_1, s_2, \dots$$

is a run of Φ , so is

$$\sigma' : s_0, \underbrace{\perp_{s_0}, \dots, \perp_{s_0}}_{0 \leq \#\{\perp_{s_0}\} < \infty}, s_1, \perp_{s_1}, \dots, \perp_{s_1}, s_2, \perp_{s_2}, \dots, \perp_{s_2}, \dots, \quad (3)$$

where, for every state s , symbol \perp_s denotes the *silent state* associated with s , defined by:

$$\forall v \in V_{\mathbf{d}} : \begin{cases} \perp_s[v] = \perp \\ \perp_s[\xi_v] = s[\xi_v] \end{cases} .$$

This means that state variables are kept unchanged, whenever their associated declared variables are absent. Note that stuttering invariance allows for runs possessing only a finite number of present states.

We require in the sequel that all STS we consider are stuttering invariant. They should indeed satisfy:

$$\left[(s^-, s) \models \rho \right] \Rightarrow \left[(s^-, \perp_{s^-}) \models \rho \right] \wedge \left[(\perp_{s^-}, s) \models \rho \right] \quad (4)$$

By convention, we shall simply write \perp when mentioning a particular state s is not required.

Examples of Transition Relations:

- A selector:

$$\mathbf{if} \ b \ \mathbf{then} \ z = u \ \mathbf{else} \ z = v . \quad (5)$$

Note that the “**else**” part corresponds to the property “ $[b = \mathbf{F}] \vee [b = \perp]$ ”.

- A register :

$$\mathbf{if } h_z \mathbf{ then } v = \xi_z^- \mathbf{ else } v = \perp . \quad (6)$$

where ξ_z is the state variable associated with z as in (2), and ξ_z^- denotes its past value. The more intuitive interpretation of this statement is : $v_n = z_{n-1}$, where index “ n ” denotes the instants at which both v and z are present (their clocks are specified to be equal). Decrementing a register would simply be specified by :

$$\mathbf{if } h_z \mathbf{ then } v = \xi_z^- - 1 \mathbf{ else } v = \perp , \quad (7)$$

where z is of integer type. Note that both statements (6,7) imply the equality of clocks :

$$h_z = h_v .$$

- Testing for a property :

$$\mathbf{if } h_v \mathbf{ then } b = (v \leq 0) \mathbf{ else } b = \perp . \quad (8)$$

Note that a consequence of this definition is, again,


$$h_v = h_b .$$

- A synchronization constraint :

$$(b = \top) = (h_u = \top) , \quad (9)$$

meaning that the clock of u is the set of instants where the boolean variable b is true.

Putting (5,7,8,9) together yields the STS :



$$\begin{aligned} & \mathbf{if } b \mathbf{ then } z = u \mathbf{ else } z = v \\ \wedge & \mathbf{if } h_z \mathbf{ then } v = \xi_z^- - 1 \mathbf{ else } v = \perp \\ \wedge & \mathbf{if } h_v \mathbf{ then } b = (v \leq 0) \mathbf{ else } b = \perp \\ \wedge & h_v = h_z = h_b \\ \wedge & (b = \top) = (h_u = \top) \end{aligned}$$

A run of this STS for the variable z is depicted on the figure above. Each time u is received, z is set to the value of u . Then z is decremented by one at each activation cycle of the STS, until it reaches the value 0. Immediately after this, a fresh u can be read, and so on. Note the schizophrenic nature of the “inputs” of this STS. While the value carried by u is an input, the instant at which u is read is not: reading of the input is on demand-driven mode. This is reflected by the fact that inputs of this STS are the pair {activation clock h , value of u when it is present}.

Using the primitives (5,6,8,9), dataflow synchronous languages such as LUSTRE [Halbwachs, 1993] and SIGNAL [LeGuernic et al., 1991] are easily encoded. Note that primitives (5,6,8,9) and their composition are stuttering invariant STS, i.e., they satisfy condition (4).


3 Compositionality in code generation: informal analysis

In this section, we informally discuss issues of compositionality aiming at code generation. After a brief review of the problems, we acknowledge the importance of extending our basic STS model with preorders; preorders are useful to capture causality, to specify schedulings, and to model communications in a distributed environment. Also, preorders are instrumental in handling abstractions. Then we discuss causality analysis and we analyse a few simple examples. Separate compilation is discussed, using preorders: we show that separate compilation requires a new level of intermediate code which allows us to store and reuse modules in a correct way. Finally we discuss the issue of distributed code generation on an asynchronous architecture.

3.1 What is the problem?

Basically, the problem is twofold: 1/ brute-force separate compilation can be the source of deadlock, and 2/ generating distributed code is generally not compatible with maintaining strict compliance with the synchronous model of computation. We illustrate briefly these two issues next.


Naive separate compilation may be dangerous. This is illustrated in the following picture:



The first diagram depicts the “dependencies” associated with some STS specification: the 1st output needs the 1st input for its computation, and the 2nd output needs the 2nd input for its computation. The second diagram shows a possible scheduling, corresponding to the standard scheduling: 1/ read inputs, 2/ compute reaction, 3/ emit outputs. This gives a correct sequential execution of the STS. In the third diagram, an additional dependency is enforced by setting the considered STS in some environment which reacts with no delay to its inputs: a deadlock is created. In the last diagram, however, it is revealed that this additional dependency caused by the environment indeed was compatible with the original specification, and no deadlock resulted from applying it. Here, deadlock was caused by the actual implementation of the specification, not by the specification itself.

The traditional answer to this problem by the synchronous programming school has been to refuse considering separate compilation: modules for further reuse should be stored as source code, and combined as such before code generation. We shall later see that this does not need to be the case, however.

Desynchronization. This is illustrated in the following picture:



This figure depicts a communication scenario: two processors, modelled as sequential machines, exchange messages using an asynchronous medium for

their communications. The natural structure of time is that of a *partial order*, as derived from the directed graph composed of 1/ linear time on each processor, and 2/ communications. This structure for time does not match the linear time corresponding to the infinite sequence of reactions which is the very basis of synchronous paradigm.

The need for reasoning about causality, schedulings, and communications. This need emerges from the above discussion. In the next subsection, we shall introduce a unique framework to handle these diverse aspects : the formalism of *scheduling specifications*.

3.2 Scheduling specifications

Causality relations have been investigated for several years in the past in the area of models of distributed systems and computations. The classical approach considers a classical automaton, in which concurrency is modelled via an “independence” equivalence relation among the labels of the transitions. Since independence is generally not a symmetric relation (actions of writing and reading are not symmetric), the theory of traces [Aabelsberg and Rozenberg, 1988] has been extended to so-called “semi-commutations” [Clerbout and Latteux, 1987], and this technique has been recently applied to the implementation of reactive automata on distributed architectures [Caillaud et al., 1997]. Causality preorder relations have also been used in a different way in [LeGuernic and Gautier, 1991], and also in [Benveniste Caspi et al., 1994], from which we borrow the essentials of the present technique. In addition to modelling causality relations, preorders can be used to specify scheduling requirements, they can also be used to model send/receive type of communications.

STS with scheduling specifications

We consider a set V of variables. A *preorder* on the set V is a relation (generically denoted by \preceq) which is reflexive ($x \preceq x$) and transitive ($x \preceq y$ and $y \preceq z$ imply $x \preceq z$). To \preceq we associate the equivalence relation \asymp , defined by $x \asymp y$ iff $x \preceq y$ and $y \preceq x$. If equivalence classes of \asymp are singletons, then \preceq is a *partial order*. Preorders are naturally specified via (*possibly cyclic*) directed graphs, denoted :

$$x \rightarrow y \text{ for } x, y \in V , \tag{10}$$

by defining $x \preceq z$ iff there is a path originating from x and terminating in z . The *supremum* of two preorders, written

$$\preceq_1 \vee \preceq_2 , \quad (11)$$

is the least preorder which is an extension of \preceq_1 and \preceq_2 . The set of all preorders on V is denoted Λ_V .

A *labelled preorder* on V is a preorder on V , together with a value $s[v]$ for each $v \in V$ over its domain. A *state* \vec{s} is a labelled preorder. The set of all states is denoted \vec{S} . As before for STS, we denote by S the set of all type consistent interpretations of V . Thus $\vec{S} = S \times \Lambda_V$, and a state \vec{s} decomposes as

$$\vec{s} = (s, \preceq_V) . \quad (12)$$

An STS *with scheduling specifications* is a triple $\vec{\Phi} = \langle V, \Theta, \vec{\rho} \rangle$, where V, Θ are as before, and

$$\vec{\rho} \subset S \times \vec{S} = S \times S \times \Lambda_V , \quad (13)$$

i.e., $\vec{\rho}$ relates the value for the tuple of previous variables to the current state.

By convention, transition relation $\vec{\rho}$ is trivially extended to a transition on \vec{S} , i.e., a subset of $\vec{S} \times \vec{S}$, and runs are sequences s_0, s_1, s_2, \dots that are consistent with transition relation (13).

We shall denote by ρ the transition relation on S obtained by projecting $\vec{\rho}$ on $S \times S$, i.e., by ignoring the preorder component. Note that $\Phi = \langle V, \Theta, \rho \rangle$ is an ordinary STS. The *composition* of two STS with scheduling specifications

$$\vec{\Phi} = \vec{\Phi}_1 \parallel \vec{\Phi}_2 , \quad (14)$$

is defined as follows :

1. Associated underlying STS (without scheduling specifications) are simply composed :

$$\Phi = \Phi_1 \parallel \Phi_2 . \quad (15)$$

Then we need to define how preorders are combined.

2. For s a state for Φ , for $i = 1, 2$ let s_i be the restriction of s to V_i , we know that s_i is a state for Φ_i . Let $\vec{s}_i = (s_i, \preceq_{V_i})$ be the corresponding state for $\vec{\Phi}_i$, cf (12). Define

$$\preceq_V \stackrel{\text{def}}{=} \preceq_{V_1} \vee \preceq_{V_2} \quad (\text{cf. (11)},) \quad (16)$$

$$\vec{s} \stackrel{\text{def}}{=} (s, \preceq_V) . \quad (17)$$

Thus (15,16,17) define how states of the components $\vec{\Phi}_i$ are combined together, building up the states and runs of $\vec{\Phi} = \vec{\Phi}_1 \parallel \vec{\Phi}_2$. Again, composition \parallel as extended to STS with scheduling specifications, is commutative and associative.

Notations for scheduling specifications

We now introduce convenient notations for the graphs generating the above introduced preorders. The notation $u \longrightarrow v$ corresponds to the edge (10). For b a variable of type $bool \cup \{\perp\}$, and u, v variables of any type, the following generic conjunct will be used to specify preorders :

$$\mathbf{if} \ b \ \mathbf{then} \ u \longrightarrow v \quad , \text{ resp. } \ \mathbf{if} \ b \ \mathbf{else} \ u \longrightarrow v$$

also written :


$$u \xrightarrow{b} v \quad \text{resp.} \quad u \xrightarrow{\bar{b}} v$$

In subsection 5.1, it is shown that scheduling specifications have the following properties :

$$x \xrightarrow{b} y \parallel y \xrightarrow{c} z \Rightarrow x \xrightarrow{b \wedge c} z \quad (18)$$

$$x \xrightarrow{b} y \parallel x \xrightarrow{c} y \Rightarrow x \xrightarrow{b \vee c} y \quad (19)$$

Properties (18,19) can be used to compute input/output abstractions of scheduling specifications :



In this figure, the diagram on the left depicts a scheduling specification involving local variables. These are hidden in the diagram on the right, using rules (18,19).

Inferring scheduling specifications from causality analysis

We now provide a technique for inferring schedulings from causality analysis for STS specified as conjunctions of the particular set of generic conjuncts we have introduced so far. Considering this restricted set of generic conjuncts is justified by the fact that 1/ all known synchronous languages can be encoded using this set of basic conjuncts, and even more, 2/ these primitives allow to express the most general synchronization mechanisms that are compatible with the paradigm of perfect synchrony [Benveniste et al., 1992]. We recall next this set of basic conjuncts for the sake of clarity :

$$\begin{array}{l}
 \mathbf{if} \ b \ \mathbf{then} \ w = u \\
 \qquad \mathbf{else} \ w = v \\
 \\
 u \xrightarrow{b} w
 \end{array} \tag{20}$$

$$\left. \begin{array}{l}
 w = f(u_1, \dots, u_k) \\
 h_w = h_{u_1} = \dots = h_{u_k}
 \end{array} \right\}$$

In addition to the set (20) of primitives, state-variable ξ_v associated to variable v can be used on the right hand side of each of the above primitive statements. The third primitive involves a conjunction of statements that are considered jointly. Later on, in the examples, we shall freely use nested expressions such as “**if** b **then** $w = expr$ ”, where “*expr*” denotes an expression built on the same set of primitives. It is understood that such expressions need to be expanded prior to applying the rules of formulas (21) given next.

In formulas (21), each primitive statement has a scheduling specification associated with it, given on the corresponding right hand side of the table. Given an STS specified as the conjunction of a set of such statements, for each conjunct we add the corresponding scheduling specification to the considered STS. Since, in turn, scheduling specifications themselves have scheduling specifications associated with them, this mechanism of adding scheduling specifications must be applied until fixpoint is reached. Note that applying these rules until fixpoint is reached takes at most two successive passes. In

formulas (21), labels of schedulings are expressions involving variables in the domain $\{\perp, \text{F}, \text{T}\}$ ordered by $\{\perp < \text{F} < \text{T}\}$; with this in mind, expressions involving the symbols “ \wedge ” (min) and “ \vee ” (max) have a clear meaning.

$$\begin{aligned}
\text{(R-1)} \quad & \forall u \quad h_u \longrightarrow u \\
\text{(R-2)} \quad & \text{if } b \text{ then } w = u \text{ else } w = v \Rightarrow \left\{ \begin{array}{l} b \xrightarrow{h_b \wedge (h_u \vee h_v)} h_w \\ h_u \xrightarrow{b \wedge h_u} h_w \\ h_v \xrightarrow{\bar{b} \wedge h_v} h_w \\ u \xrightarrow{b \wedge h_u} w \\ v \xrightarrow{\bar{b} \wedge h_v} w \end{array} \right. \quad (21) \\
\text{(R-3)} \quad & u \xrightarrow{b} w \Rightarrow b \longrightarrow h_w \\
\text{(R-4)} \quad & \left. \begin{array}{l} w = f(u_1, \dots, u_k) \\ h_w = h_{u_1} = \dots = h_{u_k} \end{array} \right\} \Rightarrow u_i \xrightarrow{h_w} w
\end{aligned}$$

Note that there is no rule involving variables of the form ξ_z^- , as previous state variables are available prior to starting the current reaction and thus do not participate to the causality calculus. Rules (R-1,...,R-4) are formally justified in section 5. We briefly report the corresponding results. For \mathbf{P} an STS, first apply Rules (R-1,...,R-4) until fixpoint is reached: this yields an STS we call **sched**(\mathbf{P}). Then, a *sufficient* condition for \mathbf{P} to have a unique deterministic run is:

1. **sched**(\mathbf{P}) is circuitfree at each instant, meaning that *it is never true that*

$$\begin{aligned}
x_1 \xrightarrow{b_1} x_2 \xrightarrow{b_2} x_1 \\
\text{and} \\
(b_1 \wedge b_2 = \text{T})
\end{aligned}$$

where x_1 and x_2 are distinct variables.

2. **sched**(**P**) has no multiple definition of variables at any instant, meaning that, whenever

$$\begin{aligned} & \mathbf{if } b_1 \mathbf{ then } x = \text{exp}_1 \\ \wedge & \mathbf{ if } b_2 \mathbf{ then } x = \text{exp}_2 \end{aligned}$$

holds in **P** and the exp_1 and exp_2 are different expressions, then

$$b_1 \wedge b_2 = \top$$

never holds in **P**.

Then **P** is said to be *executable*, and **sched**(**P**) provides (dynamic) scheduling specifications for this run. Note that proof obligations resulting from the above two conditions are generally not automatically provable, therefore abstractions may have to be considered.

Summary. What do we have at this stage?

1. STS composition is just the conjunction of constraints.
2. Scheduling specifications do compose as well.
3. Since causality analysis is based on an abstraction, the rules (R-1,...,R-4) for inferring scheduling from causality are bound to the *syntax* of the STS conjuncts. Hence, in order to maximize the chance of effectively recognizing that an STS **P** is executable, **P** is generally rewritten in a different but semantically equivalent syntax (runs remain the same) while causality analysis is performed⁴. But this latter operation is global and not compositional: here we reach the limits of ideal compositionality.

3.3 Causality analysis : examples


We show here some STS statements and their associated scheduling as derived from causality analysis. In the following figures, vertices in boldface denote input clocks, vertices in bold-italic denote input data, and vertices in courier

⁴This is part of the job performed by the SIGNAL compiler's "clock calculus".

denote other variables. It is of interest to split between these two different types of inputs, as input reading for an STS can occur with any combination of data- and demand-driven mode. Note that, for each vertex of the graph, the labels sitting on the incoming branches are evaluated prior to the considered vertex. Thus, when this vertex is to be evaluated, the other variables needed for its evaluation are already known. Resulting directed graphs (which are labelled with booleans) specify the set of all legal schedulings for the execution of the considered STS; this is formalized in section 5.


A reactive STS :

if b **then** $z = u$ **else** $z = v$



In the above example, input data are associated with their corresponding input clocks: this STS reads its inputs on a purely data-driven mode, input patterns (u, v, b) are free to be present or absent, and, when they are present, their value is free also. We call it a “reactive” STS.

The full example, a proactive STS :



if b **then** $z = u$ **else** $z = v$
 \wedge **if** h_z **then** $v = \xi_z^- - 1$ **else** $v = \perp$
 \wedge **if** h_v **then** $b = (v \leq 0)$ **else** $b = \perp$
 \wedge $h_v = h_z = h_b$
 \wedge $(b = \top) = (h_u = \top)$

Applying scheduling rules (R-1,...,R-4) and then performing some straightforward simplifications, we get the result shown in figure 1. Note the change in control: $\{ \text{input clock, input data} \}$ have been drastically modified from the “**if** b **then** $z = u$ **else** $z = v$ ” statement to the complete STS: inputs now consist of the pair $\{h, v_u\}$, where v_u refers to the value carried by u




Figure 1: *Scheduling from causality analysis for the example.*

when present. Reading of u occurs on demand, when condition b is true. We propose to call such an STS “proactive”.

3.4 Generating scheduling for separate modules

Relevant target architectures for embedded applications are typically 1/ purely sequential code (such as C-code), 2/ code using a threading or tasking mechanism provided by some kind of a real-time OS (here the threading mechanism offers some degree of concurrency), or 3/ DSP-type multiprocessor architectures with associated communication media.


On the other hand, the scheduling specifications we derive from causality rules (R-1,...,R-4) still exhibit maximal concurrency. Actual implementations will have to conform to these scheduling specifications. In general, they will exhibit less (and even sometimes no) concurrency, meaning that further sequentialization has been performed to generate code.

Of course, this additional sequentialization can be the source of potential, otherwise unjustified, deadlock when the considered module is reused in the form of object code in some environment, this was illustrated in subsection 3.1. The traditional answer to this problem by the synchronous programming school has been to refuse considering separate compilation: modules for further reuse should be stored as source code, and combined as such before code generation.

We shall however see that this does not need to be the case, however. Instead, a careful use of the scheduling specifications of an STS will allow us to decompose it into modules that can be stored as object code for further reuse, whatever the actual environment and implementation architecture will be.


For the sake of clarity, we restrict our discussion to the case of single-clocked STS, i.e., an STS in which all declared variables have the same clock.

The issue is illustrated in the following picture, in which the directed graph defining the circuitfree scheduling specification of some single-clocked STS is depicted :



In the above picture, the gray zones group all variables which depend on the same subset of inputs, let us call them “*tasks*”. Tasks are not subject to the risk of creating fake deadlocks from implementation, unlike the example from subsection 3.1. In fact, as all variables belonging to the same task depend on the same inputs, each task can be executed safely according to the following scheme: 1/ collect inputs, 2/ execute task.


In the next picture, we show how the actual implementation is prepared :



The thick arrows inside the task depicted on the right show one possible fully sequential scheduling of this task. Then, what should be really stored as source code for further reuse is only *the abstraction consisting of the tasks viewed as black-boxes, together with their associated interface scheduling specifications*. In particular, if the supporting execution architecture involves a real-time tasking system implementing some preemption mechanism in order to dynamically optimize scheduling for best response time, tasks can be freely suspended/resumed by the real-time kernel, without impairing conformity of the object code to its specification. Using our notion of scheduling specification, the above approach easily extends to general STS, in which several different clocks are involved.

3.5 Relaxing synchrony


Loosening synchrony. The major problem is that of testing for absence in an asynchronous environment. This is illustrated in the following picture in which the information about presence of variables in the considered instant is lost when passing from left- to right-hand side, since explicit definition of the “instant” is not available any more :



The question mark indicates that it is generally not possible, in an asynchronous environment, to decide upon presence/absence of a signal relatively to another one. While testing for absence is perfectly sound in a synchronous paradigm, it is meaningless in an asynchronous one.

The solution consists in restricting ourselves to so-called *endochronous* STS. Endochronous STS are those for which the control depends only on 1/ the past state, and 2/ the values possibly carried by environment signals, but not on the presence/absence status of these signals. For an endochronous STS, loosing the synchronization barriers that define the successive reactions will not result in changing its semantics ; this is formalized in subsection 4.2.


An example of an STS which is “exochronous” is the “reactive” STS given on the left-hand side of the following picture, whereas the “proactive” STS shown on the right-hand side is endochronous :



In the diagram on the left-hand side, three different clocks are source nodes of the directed graph. This means that the first decision in executing a


reaction consists in deciding upon relative presence/absence of these clocks. In contrast, in the diagram on the right-hand side, only one clock, the activation clock h , is a source node of the graph. Hence no test for relative presence/absence is needed, and the control only depends on the value of the internally computed boolean variable b .

How endochrony allows us to desynchronize an STS is illustrated in an intuitive way on the following diagram, which depicts the scheduling specification associated with the (endochronous) pseudo-statement “**if** b **then** get_u ” :



In the diagram on the left, a history of this statement is depicted, showing the successive instants (or reactions) separated by thick dashed lines. In the right-hand side diagram, thick dashed lines have been removed. Clearly, no information has been lost: we know that u should happen exactly when $b = \top$, and thus awaiting for the value of b is enough for deciding whether u is to be waited for. A formal study of desynchronization and endochrony is presented in section 4.

Moving from exochronous programs to endochronous programs can be performed, we only show one typical but simple example :



The idea is to add to the considered STS a monitor which delivers the presence/absence information via two boolean variables b, b' with identical clocks h , and such that $[k = \top] = [b = \top]$, and similarly for k', b' . The resulting STS is endochronous, since boolean variables b, b' are scrutinized at the pace of activation clock h . Other schemes are also possible, this is discussed in subsection 4.5.

Loosening synchronous composition. The second question is that of preserving the semantics of synchronous composition when an asynchronous communication medium is used. In the synchronous programming paradigm, communication occurs via instantaneous broadcast, meaning that all components must agree on 1/ which variable is present/absent in the considered reaction, and then 2/ what is the value carried by each present variable. Again this protocol is meaningless in an asynchronous communication medium. In subsection 4.3, it is shown that the condition for semantics preserving desynchronization of the communication is that the considered pair of STS should be *isochronous*.

Isochrony is a property of the synchronous composition $P \parallel Q$ of two STS. Roughly speaking, a pair of STS is isochronous if every pair of reactions, of P and Q respectively, which agree on *present* common variables, also agree on *all* common variables. Thus, again, common agreement for composition of reactions can disregard absence.

Endochrony and isochrony are the basic concepts for our theory of desynchronization. For this theory to hold, requirements for the communication medium are: 1/ it should not lose messages, and, 2/ it should not change the order of messages associated with each given variable.

3.6 Modular design, GALS architectures

From the theory informally presented in the previous subsections, the following approach results for modular design and distributed implementations of reactive systems. The target architecture is Globally Asynchronous, Locally Synchronous (GALS) by nature. The whole approach is summarized in the diagram of figure 2, where the considered STS is assumed to possess a unique, deterministic execution, i.e., it satisfies the correctness criteria stated in section 3.2. In this diagram, gray rectangles denote three modules P_1, P_2, P_3 of the source STS specification, hence given by $P = P_1 \parallel P_2 \parallel P_3$. We assume here that this partitioning has been given by the designer, based on functional and architectural considerations.

White bubbles inside the gray rectangles depict the structuration into tasks as discussed in subsection 3.4. The black half-ellipses denote the monitors. Monitors are in charge of 1/ providing the additional protocols if asynchronous communication media are to be used, and 2/ specifying the scheduling of the abstract tasks.




Figure 2: *Implementation architecture.*

In principle, communication media and real-time kernels do not need to be specified here, as they can be used freely provided they respect the send-receive abstract communication model and conform to the scheduling constraints set by the monitors.

4 Formal study of desynchronization

How far/close is indeed synchrony from asynchrony has already been discussed in the literature, thus questioning the oversimplified vision of “zero time” computation and instantaneous broadcast communication. Early paper [Benveniste and Berry, 1991] informally discussed the link between perfect synchrony and token-based asynchronous dataflow networks, see in particular section V therein. The first formal and deep study is [Caspi 1992]: a precise relation is established between so-called well-clocked synchronous functional programs and the subset of Kahn networks amenable to “bufferless” evaluation.

Distributed code generation from synchronous programs, requires to address the issue of the relationship between synchrony and asynchrony in some way or another. Mapping synchronous programs to a network of automata, communicating asynchronously via unbounded fifos, has been proposed in [Caillaud et al., 1997]. Mapping SIGNAL programs to distributed architectures was proposed in [Maffeis and LeGuernic, 1994, Aubry 1997], based on an early version of the theory we present in this paper. The SYNDEx

tool [Sorel and Lavarenne, Sorel 1996] also implements a similar approach. Recent work [Berry and Sentovich 1998] on the POLIS system proposes to reuse the “constructive semantics” approach for the ESTEREL synchronous language, with CFSM (Codesign Finite State Machines) as a model of synchronous machines which can be desynchronized.

Independently, another route to relate synchrony and asynchrony has been followed. In [Benveniste and LeGuernic 1990, LeGuernic et al., 1991] it was shown how *nondeterministic* SIGNAL programs can be used to model asynchronous communication media such as queues, buffers, etc. *Reactive Modules* were proposed [Alur and Henzinger 1996] as a synchronous language for hardware modelling, in which asynchrony is emulated by the way of non-determinism. Although this is of interest, we believe this approach is not suited to analyze true asynchrony, in which no notion of a global state is available, unlike for synchrony.

We first informally discuss the essentials of asynchrony. Synchronous Transition Systems were defined in section 2.2, and their asynchronous counterpart is defined in subsection 4.1, where desynchronization is also formally defined. The rest of this section is devoted to the analysis of desynchronization and its inverse, namely resynchronization.

4.1 Desynchronizing STS, and two fundamental problems

We first start with an informal discussion, following the discussion of subsection 2.1. Keeping in mind the essentials of the synchronous paradigm, we are now ready to discuss informally how asynchrony relates to synchrony. Referring to points 1, 2, and 3 of the discussion of subsection 2.1, the following can be stated about asynchrony :

1. Reactions cannot be observed any more : as no global clock exists, the global synchronization barriers which indicate the transition from one reaction to the next one are no more available. Instead, we only assume a reliable distributed communication medium, in which messages are not lost, and messages within each individual channel are sent and delivered in the same order. We call a *flow* such a totally ordered sequence of messages.
2. Absence cannot be sensed, and thus cannot be used to exercise control.

3. Composition occurs by means of separately unifying each common flow of the two components. This models in particular the communications via asynchronous unbounded fifos, such as used, say, in Kahn networks. Rendez-vous type of communication can also be abstracted in this way.

From the definition (1) of a run of an STS, we can say that a run is a sequence of tuples of values in domains extended with the extra symbol \perp . Desynchronizing a run amounts to discarding the synchronization barriers defining the successive reactions. Hence, for each variable $v \in V$, we only know the ordered sequence of *present* values. Thus desynchronizing a run amounts to mapping a *sequence of tuples* of values in domains extended with the extra symbol \perp , into a *tuple of sequences* of present values, one sequence per each variable. This is formalized next.

For $\sigma : s_0, s_1, s_2, \dots$ a run for Φ , we decompose state s_k as

$$s_k = (s_k[v])_{v \in V}$$

Thus we can rewrite run σ as follows :

$$\begin{aligned} \sigma &= (\sigma[v])_{v \in V}, \text{ where} \\ \sigma[v] &= s_0[v], s_1[v], \dots, s_k[v], \dots \end{aligned}$$

Now, compress each $\sigma[v]$ by deleting those $s_k[v]$ that are equal to \perp . Formally, we denote by k_0, k_1, k_2, \dots the subsequence of $k = 0, 1, 2, \dots$ such that $s_k[v] \neq \perp$. Then we set

$$\begin{aligned} \sigma^a &= (\sigma^a[v])_{v \in V}, \text{ where} \\ \sigma^a[v] &= s_{k_0}[v], s_{k_1}[v], s_{k_2}[v], \dots \end{aligned}$$

This defines the *desynchronization mapping*

$$\sigma \longmapsto \sigma^a, \tag{22}$$

where each

$$\sigma^a[v] = s_{k_0}[v], s_{k_1}[v], s_{k_2}[v], \dots$$

is called a *flow* in the sequel.

For $\Phi = \langle V, \Theta, \rho \rangle$ an STS, we define

$$\Phi^a =_{\text{def}} \langle V, \Sigma^a \rangle, \tag{23}$$

where Σ^a is the family of all σ^a , for σ ranging over the set of runs of Φ . For $\Phi_i = \langle V_i, \Theta_i, \rho_i \rangle$, $i = 1, 2$, we define

$$\Phi_1^a \parallel^a \Phi_2^a =_{\text{def}} \langle V, \Sigma^a \rangle, \text{ where } \begin{cases} V &= V_1 \cup V_2 \\ \Sigma^a &= \Sigma_1^a \wedge^a \Sigma_2^a \end{cases} \quad (24)$$

and \wedge^a denotes the conjunction of sets of asynchronous runs, which we define now. For $\sigma_i^a \in \Sigma_i^a$, $i = 1, 2$, we say that σ_1^a and σ_2^a are *unifiable*, written

$$\sigma_1^a \bowtie^a \sigma_2^a, \quad (25)$$

if the following condition holds:

$$\forall v \in V_1 \cap V_2 : \sigma_1^a[v] = \sigma_2^a[v] \text{ holds.}$$

If condition (25) holds, then we define $\sigma^a =_{\text{def}} \sigma_1^a \wedge^a \sigma_2^a$ as

$$\begin{aligned} \forall v \in V_1 \cap V_2 & : \sigma^a[v] = \sigma_1^a[v] = \sigma_2^a[v] \\ \forall v \in V_1 \setminus V_2 & : \sigma^a[v] = \sigma_1^a[v] \\ \forall v \in V_2 \setminus V_1 & : \sigma^a[v] = \sigma_2^a[v] \end{aligned}$$

Finally, Σ^a is the set of the so defined σ^a . Thus asynchronous composition proceeds via unification of shared flows.

Synchrony vs. Asynchrony? At this point two natural questions arise, namely:

Question 1 (desynchronizing a single STS) *Is resynchronization feasible and uniquely defined? More precisely, is it possible to uniquely reconstruct the original run σ for our STS from its desynchronised version σ^a as defined in (22)?*

Question 2 (desynchronizing a communication) *Does communication behave equivalently for both the synchronous and asynchronous compositions? More precisely, does the following property hold:*

$$\Phi_1^a \parallel^a \Phi_2^a = (\Phi_1 \parallel \Phi_2)^a ? \quad (26)$$

If question 1 had a positive answer, then we could desynchronize a run of the considered STS, and then still recover the original synchronous run. Thus a positive answer to question 1 would guarantee the preserving of the synchronous semantics when performing desynchronization, for a single STS.

On the other hand, if question (26) had a positive answer, then we could interpret our STS composition equivalently as synchronous or asynchronous.

Unfortunately, neither 1 nor 2 have positive answers in general, due to the possibility to exercise control by the way of absence in synchronous composition \parallel . In the following section, we show that questions 1 and 2 have positive answers under certain sufficient conditions, in which the two notions of *endochrony* (for point 1) and *isochrony* (for point 2) play a central role⁵.

4.2 Endochrony and re-synchronization

4.2.1 Formal results

In this section, we use notations from section 2.2. For $\Phi = \langle V, \Theta, \rho \rangle$ an STS, and s a reachable state of Φ , we denote by s^h the clock-abstraction of s , defined by

$$\forall v \in V : s^h[v] \in \{\perp, \top\}, \text{ and } s^h[v] = \perp \Leftrightarrow s[v] = \perp \quad (27)$$

For $\Phi = \langle V, \Theta, \rho \rangle$ an STS, s^- a reachable previous state for Φ , and $W' \subseteq W \subseteq V$, we say that W' is a *clock inference of W given s^-* , written

$$W' \hookrightarrow_{s^-} W, \quad (28)$$

if, for each state s reachable from s^- for Φ , knowing the presence/absence and actual value carried by each variable belonging to W' , allows us to determine exactly the presence/absence for each variable belonging to W . In other words,

$$s[W'] \text{ determines } s^h[W]. \quad (29)$$

If $W' \hookrightarrow_{s^-} W_1$ and $W' \hookrightarrow_{s^-} W_2$ hold, then $W' \hookrightarrow_{s^-} (W_1 \cup W_2)$ follows, thus there exists a greatest W such that $W' \hookrightarrow_{s^-} W$ holds. Hence we can consider the unique increasing chain, for s^- given,

$$\emptyset = V(0) \hookrightarrow_{s^-} V(1) \hookrightarrow_{s^-} V(2) \hookrightarrow_{s^-} \dots \quad (30)$$

⁵ Endochronous, from ancient greek *ενδο*–inside and *χρονος*–time; Isochronous, from ancient greek *ισο*–identical and *χρονος*–time. It's sometimes nice to remember that ancient greeks used to be great scientists, and thus honor them by reusing their words in our context.

of subsets of V such that, for each k , $V(k)$ is the greatest set of variables such that $V(k-1) \hookrightarrow_{s^-} V(k)$ holds. As $\emptyset = V(0)$, $V(1)$ consists of the subset of variables that are present as soon as the considered STS gets activated⁶. Of course chain (30) must become stationary at some finite k_{\max} : $V(k_{\max}+1) = V(k_{\max})$. In general, we only know that $V(k_{\max}) \subseteq V$. Chain (30) is called the *synchronization chain* of Φ .

Definition 1 (endochrony) STS Φ is said to be endochronous if, for each state s^- reachable for Φ , $V(k_{\max}) = V$, i.e., if the following condition is satisfied: the synchronization chain

$$(E) \quad \emptyset = V(0) \hookrightarrow_{s^-} V(1) \hookrightarrow_{s^-} V(2) \hookrightarrow_{s^-} \dots \text{ converges to } V. \quad (31)$$

Condition (31) expresses that presence/absence of all variables can be inferred *incrementally* from already known values carried by present variables and state variables of the STS in consideration. Hence no test for presence/absence on the environment is needed. The following theorem justifies our approach:

Theorem 1 Consider an STS $\Phi = \langle V, \Theta, \rho \rangle$.

1. Conditions (a) and (b) are equivalent, where:

(a) Φ is endochronous.

(b) For each $\delta \in \Sigma^a$, we can reconstruct the corresponding synchronous run σ such that $\sigma^a = \delta$, in a unique way up to silent reactions.

2. Assume Φ is endochronous and stuttering invariant. If $\Phi' = \langle V, \Theta, \rho' \rangle$ is another endochronous and stuttering invariant STS then

$$(\Phi')^a = \Phi^a \Rightarrow \Phi' = \Phi \quad (32)$$

Proof: We prove successively points 1 and 2.

1. We fix the previous state s^- and prove the result by induction. Pick a $\delta \in \Sigma^a$, and assume for the moment that we were able to decompose it as:

$$\underbrace{s_1, s_2, \dots, s_n}_{n\text{-initial segment of } \sigma}, \delta_n \quad (33)$$

⁶Of course we assume here that no variable is absent in every reachable state.

i.e., into a finite sequence of length n composed of non-silent states s_i (the head of the synchronous run σ we wish to reconstruct), followed by the tail of the asynchronous run δ , which we denote by δ_n , and we assume that such a decomposition is unique. Then we claim that

$$(33) \text{ is also valid with } n \text{ substituted by } n + 1. \quad (34)$$

To prove (34), we note that, when STS Φ gets activated, then we know that variables belonging to $V(1)$ will be present in the considered state. By assumption, the clock-abstracted state $s_{n+1}^h[V(1)]$, having $V(1)$ as variables, is uniquely determined. In the sequel we write $s_{n+1}^h(1)$ for short instead of $s_{n+1}^h[V(1)]$. Thus, presence/absence of variables for state $s_{n+1}(1)$ is known, it remains to determine the values carried by present variables.

For $v \in V_1$, we simply pick the value carried by the minimal element of the sequence associated with variable v in δ_n . Values carried by corresponding state variables are updated accordingly. Thus we know all of $s_{n+1}(1)$.

Next we move on constructing $s_{n+1}(2)$. From $s_{n+1}(1)$ we know $s_{n+1}^h(2)$. Thus we know how to split V_2 into present and absent variables for the considered state. Pick the present ones, and repeat the same argument as before to get $s_{n+1}(2)$.

Repeat this argument until $V(k) = V$ for some finite k (by endochrony assumption). This proves claim (34).

Given the initial condition for δ , we get from (34), by induction, the desired proof that $(a) \Rightarrow (b)$.

Next, we prove $(b) \Rightarrow (a)$. We assume that Φ is not endochronous, and show that condition (b) cannot be satisfied. If Φ is not endochronous, there must be some reachable state s^- for which chain (31) does not converge to V . Thus again we pick a $\delta \in \Sigma^a$, decomposed as for case 1, cf. formula (33):

$$\underbrace{s_1, s_2, \dots, s_n}_{n\text{-initial segment of } \sigma}, \delta_n$$

and we assume in addition that $s_n = s^-$, the given state for which endochrony is violated. We now show that (34) is disproved. Let

$k_* \geq 0$ be the smallest index such that $V(k) = V(k + 1)$, we know $V_{k_*} \neq V$. Thus we can apply the algorithm of case 1 for reconstructing the reaction, until variables of V_{k_*} . Then presence/absence for variables belonging to $V \setminus V_{k_*}$ cannot be determined based on the knowledge of variables belonging to V_{k_*} . Thus there are several possible extensions for $s_{n+1}^h(k_* + 1)$ and thus $(n + 1)$ -st reaction is not determined in a unique way. Hence condition (b) is falsified.

2. Assume Φ is endochronous, and consider Φ' as in point 2 of the theorem. As both Φ and Φ' are stuttering invariant, point 2 is an immediate consequence of point 1. \diamond

COMMENTS.

1. For an STS, endochrony is not decidable in general. It is decidable for STS involving, say, only finite domains for their variables, and model checking can be used for that. For general STS, model checking can be used, in combination with abstraction techniques. The case of interest is when the chain $V(0), V(1), \dots$ does not depend upon the particular state s^- , and we write simply $V(k) \leftrightarrow V(k + 1)$ in this case.
2. The proof of this theorem in fact provides an effective algorithm for the on-the-fly reconstruction of the successive reactions, for a desynchronized run of an endochronous program.

(COUNTER)EXAMPLES.


examples :

- a single-clocked STS.
- STS “**if** $b = \top$ **then** get u ”, where b, u are the two inputs, and b is boolean. The clock of b coincides with the activation clock for this STS, and thus $V(1) = \{b\}$. Then, knowing the value for b indicates whether or not u is present, thus $V(2) = \{b, u\} = V$.

counterexample : STS “**if** ($[\mathbf{present} \ a] \parallel [\mathbf{present} \ b]$) **then**...” is not endochronous, as the environment is free to offer any combination of presence/absence for the two inputs a, b . Thus $\emptyset = V(0) = V(1) = V(2) = \dots \stackrel{\subset}{\neq} V$, and endochrony does not hold.

4.2.2 Practical consequences

A first use of endochrony is shown in the following figure :



In this figure, a pair (Φ_1, Φ_2) of STS is depicted, with W as set of shared variables. Rewrite their composition as follows :


$$\Phi_1 \parallel \Phi_2 = \Phi_1 \parallel \Psi_{1,2} \parallel \Phi_2$$

where $\Psi_{1,2}$ is the restriction of $\Phi_1 \parallel \Phi_2$ to W , hence $\Psi_{1,2}$ models the synchronous communication channel. Using the property $\Phi \parallel \Phi = \Phi$ for every STS Φ , we get

$$\Phi_1 \parallel \Phi_2 = \underbrace{(\Phi_1 \parallel \Psi_{1,2})}_{\tilde{\Phi}_1} \parallel \underbrace{(\Psi_{1,2} \parallel \Phi_2)}_{\tilde{\Phi}_2} = \tilde{\Phi}_1 \parallel \tilde{\Phi}_2 \quad (35)$$

Assume now that channel model $\Psi_{1,2}$ is endochronous, and composition $\Phi_1 \parallel \Phi_2$ is implemented as the (equivalent) composition $\tilde{\Phi}_1 \parallel \tilde{\Phi}_2$. Then, as $\tilde{\Phi}_1$ knows channel $\Psi_{1,2}$ and the latter is endochronous, then communication can be equivalently implemented according to perfect synchrony or full asynchrony.

This is fine, but it does not extend to networks of STS involving more than two nodes. The following figure shows an example :



Assume Ψ_1, Ψ_2 are both endochronous. Then communication between Φ_1 and Φ on the one hand, and Φ and Φ_2 on the other hand, can be desynchronized. Unfortunately, communication between Φ_1 and Φ_2 via Φ can't, as it is not true in general that $\Psi_1 \parallel \Phi \parallel \Psi_2$ is endochronous. The problem is that endochrony is not compositional, hence even ensuring in addition that Φ itself is endochronous would not do. Thus we would need to ensure that Ψ_1, Ψ_2 as well as $\Psi_1 \parallel \Phi \parallel \Psi_2$ are all endochronous, not an elegant solution

when networks are considered! Thus we move on introducing the alternative notion of *isochrony*, which focusses on communication, and is compositional.

4.3 Isochrony, and synchronous and asynchronous compositions

The next result addresses the question of when property (26) holds true. We are given two STS $\Phi_i = \langle V_i, \Theta_i, \rho_i \rangle, i = 1, 2$. Denote by $W = V_1 \cap V_2$ the set of their common variables, and by $\Phi = \Phi_1 \parallel \Phi_2$ their synchronous composition. For s a reachable state in Φ , we denote by $s_1 =_{\text{def}} s[V_1]$ and $s_2 =_{\text{def}} s[V_2]$ the restrictions of state s to Φ_1 and Φ_2 , respectively. Note that, for $i = 1, 2$, s_i is a reachable state for Φ_i . Corresponding notations s^-, s_1^-, s_2^- for past states will be used accordingly.

Definition 2 (isochrony) *Consider a pair (Φ_1, Φ_2) of STS. Transitions of $\Phi_i, i = 1, 2$, are written (s_i^-, s_i) . Consider the following conditions on pairs $((s_1^-, s_1), (s_2^-, s_2))$ of transitions for (Φ_1, Φ_2) :*

- (i)
 1. $s_1^- = s^-[V_1]$ and $s_2^- = s^-[V_2]$ holds for some reachable state s^- for Φ , in particular s_1^- and s_2^- are unifiable;
 2. none of the states $s_i, i = 1, 2$ are silent on the common variables, i.e., it is not the case that, for some $i = 1, 2$: $s_i[v] = \perp$ holds $\forall v \in W$;
 3. s_1 and s_2 coincide over the set of present common variables⁷, i.e.:

$$\forall v \in W : (s_1[v] \neq \perp \text{ and } s_2[v] \neq \perp) \Rightarrow s_1[v] = s_2[v] ;$$


- (ii) *States s_1 and s_2 coincide over the whole set of common variables, i.e., states s_1 and s_2 are unifiable:*

$$s_1 = s[V_1] \text{ and } s_2 = s[V_2] \text{ holds for some state } s \text{ for } \Phi .$$

The pair (Φ_1, Φ_2) is called isochronous if condition (i) implies condition (ii), for each pair $((s_1^-, s_1), (s_2^-, s_2))$ of transitions for (Φ_1, Φ_2) .

⁷By convention this is satisfied if the set of present common variables is empty.

COMMENT. Roughly speaking, condition of isochrony expresses that unifying over *present* common variables is enough to guarantee the unification of the two considered states s_1 and s_2 . Condition of isochrony is illustrated on the following figure :



The figure depicts, for unifiable previous states s_1^-, s_2^- , corresponding states s_1, s_2 where (s_i^-, s_i) is a valid transition for Φ_i . It shows the interpretation of s_1 (circle on the left) and s_2 (circle on the right) over shared variables W . White and dashed areas represent absent and present values, respectively. The two left and right circles are superimposed in the mid circle. In general, vertically and horizontally dashed areas do not coincide, even if s_1 and s_2 unify over the subset of shared variables that are present for both transitions (double-dashed area). Pictorially, unification over double-dashed area does not imply in general that dashed areas coincide. Isochrony indeed requires that unification over double-dashed area does imply that dashed areas coincide, hence unification of s_1 and s_2 follows. It is interesting to reformulate isochrony in a different way.

Define the *desynchronized conjunction* of two transition relations $\rho_1 \wedge_a \rho_2$ as follows. For t_1 and t_2 two transitions, we define *asynchronous unifiability* $t_1 \bowtie^a t_2$ by :

$$t_1 \bowtie^a t_2 \quad \text{iff} \quad \left(\begin{array}{l} v \in V_1 \cap V_2, \text{ and} \\ t_1[v] \neq \perp \text{ and } t_2[v] \neq \perp \end{array} \right) \Rightarrow (t_1[v] = t_2[v]) \quad (36)$$

Note that $t_1 \bowtie^a t_2$ means that transitions t_1 and t_2 are unifiable on their common *present* ports, regardless of absence (this is just the restriction to transitions of the definition of \bowtie^a which was formulated for flows). Definition (36) is in contrast to *synchronous unifiability*, or unifiability for short, $t_1 \bowtie t_2$ defined by :

$$t_1 \bowtie t_2 \quad \text{iff} \quad (v \in V_1 \cap V_2) \Rightarrow (t_1[v] = t_2[v]) \quad (37)$$

which means that transitions t_1 and t_2 are unifiable on their common ports, including presence/absence. Condition (37) corresponds to the conjunction of transition relations introduced in the definition of STS composition. If $t_1 \bowtie^a t_2$, we can define $t_1 \sqcup^a t_2$ by

$$(t_1 \sqcup^a t_2)[v] \\ =_{\text{def}} \text{ if } \exists i = 1, 2 : (v \in V_i \text{ and } t_i[v] \neq \perp) \text{ then } t_i[v] \text{ else } \perp$$

With this in mind, we define $\rho_1 \wedge_a \rho_2$ as follows :

$$\rho_1 \wedge_a \rho_2 = \{t_1 \sqcup^a t_2 : t_i \models \rho_i \forall i = 1, 2 \wedge t_1 \bowtie^a t_2\},$$

and isochrony is equivalently reformulated as follows :

Definition 3 (Isochrony, reformulation) *Let (Φ_1, Φ_2) be a pair of STS and $\Phi = \Phi_1 \parallel \Phi_2$ be their parallel composition. The pair (Φ_1, Φ_2) is called **isochronous** if*

$$\rho_1 \wedge \rho_2 = \rho_1 \wedge_a \rho_2 \tag{38}$$

holds, restricted to the set of reachable states for Φ .

The following theorem justifies introducing this notion of isochrony.

Theorem 2

1. *If the pair (Φ_1, Φ_2) is isochronous, then it satisfies property (26).*
2. *Conversely, assume in addition that Φ_1 and Φ_2 are both endochronous. If the pair (Φ_1, Φ_2) satisfies property (26), then it is isochronous.*

Thus, isochrony is sufficient for (26) to hold, and it is also in fact necessary when the components are endochronous.

COMMENTS :

1. We already discussed the importance of guaranteeing property (26). Now, why is this theorem interesting? Mainly because it replaces condition (26), which involves infinite runs, by condition **(I)** of isochrony, which only involves a single reaction for the considered pair of STS.
2. Comment 1 for endochrony also applies here.

Proof: We successively prove points 1 and 2.

1. Isochrony implies property (26). We proceed into two steps.

1. The desynchronization of Φ , defined by (23), is denoted by Φ^a , and we denote by δ a run of Φ^a . For each $\delta \in \Sigma^a$, there is at least one corresponding synchronous run σ for Φ such that $\delta = \sigma^a$. Any such σ is clearly the synchronous composition of two unifiable runs σ_1 and σ_2 for Φ_1 and Φ_2 , respectively. Hence associated asynchronous runs σ_1^a and σ_2^a are also unifiable, and their asynchronous composition $\sigma_1^a \wedge^a \sigma_2^a$ belongs to $\Sigma_1^a \wedge^a \Sigma_2^a$. Thus we always have the inclusion

$$\Phi_1^a \parallel^a \Phi_2^a \supseteq (\Phi_1 \parallel \Phi_2)^a, \quad (39)$$

which proves the first part of (26). So far we have only used the definition of desynchronization and asynchronous composition, isochrony has not yet been used.

2. To prove the opposite inclusion, we need to prove that, when moving from asynchronous composition to synchronous one, the additional need for a reaction-per-reaction matching of unifiable runs will not result in rejecting pairs of runs that otherwise would be unifiable in the asynchronous sense. This is where condition **(I)** of isochrony enters the game.

Pick a pair (δ_1, δ_2) such that $\delta_1 \bowtie^a \delta_2$ (cf. (25)): they can be combined while performing the asynchronous composition $\Phi_1^a \parallel^a \Phi_2^a$ to form some δ (cf. (24)), this is denoted by $\delta_1 \wedge^a \delta_2 = \delta$. By definition of desynchronization (cf. subsection 4.1), there exist a (synchronous) run σ_1 for Φ_1 , and a (synchronous) run σ_2 for Φ_2 , such that δ_i is obtained by desynchronizing σ_i , $i = 1, 2$ (as we do not assume endochrony at this point, run σ_i is not uniquely determined). Thus each run σ_i is a succession of states. Clearly, inserting finitely many silent states between successive states of σ_i would also provide valid candidates for recovering δ_i after desynchronization. We shall show, by induction over successive states, that:

$$\begin{aligned} &\text{properly inserting such a silent state in the appropriate} \\ &\text{component will provide two runs which are} \quad (40) \\ &\textit{unifiable} \text{ in the synchronous sense.} \end{aligned}$$

This will show that, from a pair (δ_1, δ_2) such that $\delta_1 \bowtie^a \delta_2$, we can reconstruct (at least) one pair (σ_1, σ_2) of runs for Φ_1 and Φ_2 that are unifiable in the synchronous sense, and thus will prove the alternative inclusion

$$\Phi_1^a \parallel^a \Phi_2^a \subseteq (\Phi_1 \parallel \Phi_2)^a . \quad (41)$$

From (39) and (41) we then deduce property (26). We prove (40) now, by induction over successive states.

We are given a pair (δ_1, δ_2) such that $\delta_1 \bowtie^a \delta_2$. Pick a σ_1 such that $\sigma_1^a = \delta_1$, and similarly for σ_2 . For s_1, s_2, \dots, s_n a finite run, we say that another run s'_1, s'_2, \dots, s'_m is a *stretching* of s_1, s_2, \dots, s_n , written

$$s'_1, s'_2, \dots, s'_m = (s_1, s_2, \dots, s_n)^\uparrow \quad (42)$$

if there is a strictly increasing subsequence k_1, \dots, k_n of $1, \dots, m$ such that $s'_{k_j} = s_j, j = 1, \dots, n$, and $s'_k = \perp$ for $k \neq k_1, \dots, k_n$. Note that (42) implies $m \geq n$. Using notation (42) we introduce the following hypothesis, for use in our inductive reasoning: for $i = 1, 2$, run σ_i decomposes as

$$\sigma_i = \underbrace{s_{i,1}, s_{i,2}, \dots, s_{i,n_i}}_{\text{initial segment of length } n_i}, \sigma_{i,n_i} \quad (43)$$

and there are stretchings such that

$$\begin{aligned} s'_{i,1}, s'_{i,2}, \dots, s'_{i,n} &= (s_{i,1}, s_{i,2}, \dots, s_{i,n_i})^\uparrow \text{ for } i = 1, 2 \\ s'_{1,m} &\bowtie s'_{2,m} \text{ for } m = 1, \dots, n \end{aligned} \quad (44)$$

Note that (44) implies $\sigma_{1,n_1}^a \bowtie^a \sigma_{2,n_2}^a$. Define index

$$\zeta(n) = \min\{n_1, n_2\}$$

where n_i is defined in (43). To perform the proof by induction, we need to extend (43,44) in such a way that index $\zeta(n)$ grows to infinity.

To this end, decompose the tail σ_{i,n_i} into

$$\sigma_{i,n_i} = s_{i,n_i+1}, \sigma_{i,n_i+1} .$$

The following cases can occur:

CASE 1 : none of the two states s_{1,n_1+1} and s_{2,n_2+1} is silent over the common W variables. Concentrate on those $v \in W$ variables that are present in both states s_{1,n_1+1} and s_{2,n_2+1} . As $\delta_1 \bowtie^a \delta_2$ holds, then we must have $s_{1,n_1+1}[v] = s_{2,n_2+1}[v]$ for any such v . Thus points 1,2,3 of condition **(I)** of isochrony are satisfied. Hence s_{1,n_1+1} and s_{2,n_2+1} are indeed unifiable in this case, *by isochrony*. Therefore, in this case, hypothesis (43,44) extends in such a way that $\zeta(n+1) = \min\{n_1+1, n_2+1\} = \zeta(n)+1$ holds.

CASE 2 : both states s_{1,n_1+1} and s_{2,n_2+1} are silent over the common W variables. They are unifiable. Again, hypothesis (43,44) extends in such a way that $\zeta(n+1) = \zeta(n)+1$ holds.

CASE 3 : one and only one of the two states s_{1,n_1+1} and s_{2,n_2+1} is silent over the common W variables, say $\forall v \in W : s_{1,n_1+1}[v] = \perp$. In this case we unify state s_{1,n_1+1} with the silent state \perp for Φ_2 . Thus the matching hypothesis (44) is extended as :

$$\begin{aligned} s'_{1,1}, s'_{1,2}, \dots, s'_{1,n}, s'_{1,n+1} &= (s_{1,1}, s_{1,2}, \dots, s_{1,n_1}, s_{1,n_1+1})^\uparrow \\ s'_{2,1}, s'_{2,2}, \dots, s'_{2,n}, \underbrace{\perp}_{s'_{2,n+1}} &= (s_{2,1}, s_{2,2}, \dots, s_{2,n_2})^\uparrow \\ s'_{1,m} \bowtie s'_{2,m} &\text{ for } m = 1, \dots, n+1. \end{aligned} \quad (45)$$

Therefore $\zeta(n+1) = \min\{n_1+1, n_2\}$ and we cannot infer that $\zeta(n+1) > \zeta(n)$ holds in this case.

Given the analysis above, we only need to show that

$$\begin{aligned} \text{CASE 3 cannot occur for infinitely} \\ \text{many successive induction steps.} \end{aligned} \quad (46)$$

Assume (46) does not hold. Then this implies that the whole tail σ_{1,n_1} is silent over the common W variables, while σ_{2,n_2} is not. But on the other hand we should have $\sigma_{1,n_1}^a \bowtie^a \sigma_{2,n_2}^a$, see(44), whence a contradiction. This finishes the induction proof, hence (41) follows.

2. Under endochrony of the components, property (26) implies isochrony. This is easy. From Theorem 1 we know that, in our argument for proving point 1 of theorem 2, the synchronous runs σ_i are uniquely defined, up to silent states, from their desynchronized respective versions σ_i^a .

Now, focus on CASE 1 of this argument. If isochrony is not satisfied, then, for some pair $\sigma_1^a \bowtie^a \sigma_2^a$ of unifiable asynchronous runs, and some decomposition (43) of them, it follows that points 1,2,3 of condition **(I)** of isochrony are satisfied, but states $s_{1,n+1}$ and $s_{2,n+1}$ are *not* unifiable. As our only possibility is to try to insert silent states for one of the two components – not feasible in CASE 1 – our process of incremental unification on a per reaction basis fails. Thus (41) is violated, and so is property (26). This finishes the proof of the theorem. \diamond

The following result is instrumental in proving compositionality of isochrony.

Lemma 1 *If pairs (Ψ, Φ_1) and (Ψ, Φ_2) are isochronous, then so is pair $(\Psi, \Phi_1 \parallel \Phi_2)$.*

Proof: Let (s^-, s) and (t^-, t) be pairs of successive states, for Ψ and $\Phi_1 \parallel \Phi_2$ respectively, satisfying condition **(I)** for isochrony, see definition 2 or 3. Let t be the unification of the two states s_1 and s_2 for Φ_1 and Φ_2 , respectively. By point 2 of **(I)**, at least one of these two states is not silent, assume s_1 is not silent. From point 3 of **(I)**, s and s_1 coincide over the set of *present* common variables, and thus, since pair (Ψ, Φ_1) is isochronous, states s and s_1 coincide over the *whole* set of common variables for Ψ and Φ_1 . Thus s and s_1 are unifiable. But, on the other hand, s_1 and s_2 are also unifiable since they are just restrictions of the same global state t for $\Phi_1 \parallel \Phi_2$. Thus states s and t are unifiable, and thus pair $(\Psi, \Phi_1 \parallel \Phi_2)$ is isochronous. This proves lemma 1. \diamond

An interesting immediate byproduct is the extension of the results on desynchronization, to networks of communicating synchronous components :

Corollary 1 (desynchronizing a network of components) *We are given a finite family $(\Phi_k)_{k=1,\dots,K}$ of STS. Assume that each pair $(\Phi_k, \Phi_{k'})$ is isochronous. Then*

1. *For each disjoint subsets I and J of set $\{1, \dots, K\}$, the pair*

$$\left(\parallel_{k \in I} \Phi_k \quad , \quad \parallel_{k' \in J} \Phi_{k'} \right) \quad (47)$$

is isochronous. Thus isochrony is compositional.

2. *Also, desynchronization extends to the network :*

$$(\Phi_1 \parallel \dots \parallel \Phi_K)^a = \Phi_1^a \parallel^a \dots \parallel^a \Phi_K^a . \quad (48)$$

Proof:

1. Property (47) follows from lemma 1 via obvious induction on the cardinal of sets I, J .
2. The second statement is proved via induction on the cardinal of the number of components:

$$\begin{aligned} (\Phi_1 \parallel \dots \parallel \Phi_K)^a &= ((\Phi_1 \parallel \dots \parallel \Phi_{K-1}) \parallel \Phi_K)^a \\ &= (\Phi_1 \parallel \dots \parallel \Phi_{K-1})^a \parallel^a \Phi_K^a, \end{aligned}$$

and the induction step follows from (47). \diamond

The next corollary expresses that isochrony is a “local” property.

Corollary 2 (locality of isochrony) *Assume pair (Φ_1, Φ_2) is isochronous, and pair (Ψ_1, Ψ_2) is such that Ψ_1 has no common variable with $\Phi_2 \parallel \Psi_2$ and Ψ_2 has no common variable with $\Phi_1 \parallel \Psi_1$. Then pair $(\Psi_1 \parallel \Phi_1, \Phi_2 \parallel \Psi_2)$ is also isochronous.*

Proof: This follows directly from lemma 1. \diamond

This is a useful result, it says that, in order for a pair $(\parallel_{k \in I} \Phi_k, \parallel_{k' \in J} \Phi_{k'})$ to be isochronous, it is enough to check isochrony for pairs $(\Phi_k, \Phi_{k'})$ of *interacting components*.

Note however that, in order for a pair $(\Psi_1 \parallel \Phi_1, \Phi_2 \parallel \Psi_2)$ to be isochronous, it is not necessary, but only sufficient, that the pair (Φ_1, Φ_2) is isochronous.

(COUNTER)EXAMPLES.

examples :

- a single-clocked communication between two STS.
- the pair $(\tilde{\Phi}_1, \tilde{\Phi}_2)$ of formula (35).

counterexample : assume an STS communicates with another one according to the synchronous protocol “**await** $x \parallel$ **await** y ”, the resulting pair of STS is not isochronous.

4.4 Getting GALS architectures

In practice, only partial desynchronization of networks of communicating STS may be considered. This means that we really want to have *Locally Synchronous* components communicating via a *Globally Asynchronous* communication medium — this is referred to as GALS architectures.

In fact, theorems 1 and 2 provide the adequate solution. Let us assume we have a finite collection Φ_i of STS such that :

1. each Φ_i is *endochronous*, and
2. each pair (Φ_i, Φ_j) is *isochronous*.

Then, from corollary 1 and theorem 1, we know that

$$(\Phi_1 \parallel \dots \parallel \Phi_K)^a = \Phi_1^a \parallel^a \dots \parallel^a \Phi_K^a$$

and each Φ_k^a is in one-to-one correspondence with its synchronous counterpart Φ_k . Here is the resulting running mode for this GALS architecture :

- For communications involving a pair (Φ_i, Φ_j) of STS, each flow is preserved individually, but global synchronization is lost.
- Each STS Φ_i reconstructs its own successive reactions by just observing its (desynchronized) environment, and then locally behaves as a synchronous STS.
- Note that it is allowed, for each Φ_i , to have an internal activation clock which is *faster* than communication clocks. Resulting local activation clocks evolve asynchronously from one another.

4.5 Handling endo/isochrony in practice

While we have given criteria for endochrony and isochrony, we did not propose a practical algorithm for checking these criteria. We do this now. Our aim is to prepare for GALS architectures such as discussed in subsection 4.4. In particular, throughout this subsection, a network of STS satisfying conditions 1 and 2 of subsection 4.4 will be called *endo/isochronous*.

In this subsection, we shall indicate 1/ how a (tight) sufficient condition for endo/isochrony can be actually tested, and 2/ how making an STS endo/isochronous can be performed. As both the DC₊ format and the SIGNAL




Figure 3: *The clock hierarchy computed by the DC₊ or SIGNAL compiler.*

language can be considered as concrete instances of our STS model, we shall rely for our explanation on tools and algorithms already developed in these environments.

4.5.1 Checking endo/isochrony

As one of the modules of the existing DC₊ or SIGNAL compiler, the data structure shown in Figure 3 is computed, for a given program P : In this figure, b, c denote boolean variables, $[b], [c]$ denote clocks composed of the instants at which $b, c = \top$ holds, respectively. Finally, h, k are also clocks. The down-arrows $h_0 \rightarrow b_1$, $[b_1] \rightarrow b_2$, $[b_2] \rightarrow b_3$, etc, indicate that boolean variable b_1 has a clock equal to h_0 and only needs variables with clock h_0 for its evaluation, and so on. Roots of the trees are related by clock equations, depicted for instance by the bidirectional arrow relating h_0 and k_0 . This defines a tree under each clock h_0, k_0, \dots , and yields the so-called *clock hierarchy* in the form of a “forest”, i.e., a collection of trees related by clock equations. This structure is detailed in [Amagbegnon et al., 1994] [Amagbegnon et al., 1995], where it is shown to be a canonical representation of the combination of clock equations and scheduling specifications of a program. Now, considering this clock hierarchy, one easily proves the following :

Theorem 3 *Assume program P has a clock hierarchy consisting of a single tree. Also assume it is decomposed as $P = P_1 \parallel \dots \parallel P_K$, and, for each k , the clock hierarchy of component P_k is a subtree of the clock tree of P . Then the corresponding network of STS is endo/isochronous.*

Theorem 3 is an immediate corollary of Theorem 1 of section 4, it only states a sufficient condition. In computing a clock hierarchy, the abstractions performed are twofold: 1/ inferring dependencies from causality analysis, and 2/ abstracting boolean variables which result from the evaluation of a predicate involving a non-boolean expression. In practice, we shall use the clock hierarchy as the practical criterion for checking endo/isochrony.

4.5.2 Enforcing endo/isochrony

Assume we have an STS P having a clock hierarchy which is not a tree, and we still *want* it to be a tree. What can we do? As revealed by inspecting the previous figure, it is sufficient to make the roots h_0, k_0, \dots of the clock hierarchy belonging to some single clock tree. In other words, we can concentrate on the roots of the clock hierarchy. Thus the problem can be restated as follows:

We are given a set h_1, \dots, h_k of clocks, which are related by a set of clock equations of the form:

$$\begin{aligned} p_1(h_1, \dots, h_k) &\neq \text{F} \\ &\dots \\ p_q(h_1, \dots, h_k) &\neq \text{F} \end{aligned} \tag{49}$$

This corresponds to having a collection p_1, \dots, p_q of predicates on clocks, which are boolean-valued expressions that are either true or absent. Note that being always true is the case for predicates in classical boolean logic, while in our case, due to the requirement for stuttering robustness, we must accept the possibility for a “clock predicate” to be absent. Systems of equations of the form (49) can be solved for their variables h_1, \dots, h_k , meaning that we can find a set h_1^o, \dots, h_k^o of clocks, and a set p_1^o, \dots, p_k^o of clock expressions, such that equation system:

$$\begin{aligned} h_1 &= p_1^o(h_1^o, \dots, h_k^o) \\ &\dots \\ h_k &= p_k^o(h_1^o, \dots, h_k^o) \end{aligned} \tag{50}$$

has the same set of solutions for h_1, \dots, h_k as the original system (49), and new clocks h_1^o, \dots, h_l^o are free, i.e., unconstrained by the system of equations (50). Finally, we introduce boolean variables b_1^o, \dots, b_l^o , and a “master clock” h^o , such that

$$\begin{aligned} h_1^o &= [b_1^o], \dots, h_l^o = [b_l^o] \\ h_{b_1^o} &= \dots = h_{b_l^o} = h \end{aligned} \tag{51}$$

The bottom line is :

1. System of clock equations (49) is equivalent to (50,51) after hiding auxiliary variables h, b_1^o, \dots, b_l^o .
2. System (50,51) is a clock tree.

Discussion. Basically, building (51,50) from (49) intuitively corresponds to equipping the original P program with a suitable *communication protocol* Q in such a way that the compound program $P \parallel Q$ is endo/isochronous. This is not surprising indeed, for it is known in the area of distributed systems that components in a distributed system must be equipped with suitable protocols for their communications.

Finally, the way we moved from (49) to (50) reveals one unpleasant feature of this technique, namely: this part of the process is not unique, and thus there are possibly many different correct protocols.

5 Formal study of causality

In this section we develop a formal theory of causality for STS. Our basic tool is that of scheduling specifications and labelled preorders. We first formalize this, by adding the value *unknown* to our domains, like in the Constructive Boolean logic used in [Berry and Sentovich 1998]. Using this extended domain, we are able to formally state and prove our criterion that circuit-freedom implies executability. Then we formalize the rules (R-1,2,3,4) of (21), and we finally show how correct deterministic execution results from a successful causality analysis.

5.1 Encoding scheduling specifications using an algebraic domain

In this section, we consider the following domain \mathcal{D} and its two orderings \prec and $<$ as an abstraction of arbitrary domains of values :

$$\mathcal{D} = \{?, \underbrace{\perp, \mathbf{F}, \mathbf{T}}_{\top}\} \quad (52)$$

$$? \prec \perp, \mathbf{F}, \mathbf{T} \quad \perp < \mathbf{F} < \mathbf{T} \quad (53)$$

In these formulas , symbols ? (resp. $\mathfrak{!}$) indicate that the value is “unknown” (resp. “known”). The “unknown” status should not be confused with absence (\perp): absence is a perfectly known status, while “unknown” is intended to model that a variable has not been produced yet in the current reaction. Non-boolean types are abstracted as the single distinguished element \top , hence, for booleans, the pair $\{\mathbf{F}, \mathbf{T}\}$ can be seen as a refinement of the symbol \top , this is shown by the underbrackets. And $\{\perp, \top\}$ is a refinement of $\mathfrak{!}$, this is shown by the overbrackets. Ordering $<$ has already been introduced, and the additional partial order \prec is the Scott information ordering: $? \prec \perp, \mathbf{F}, \mathbf{T}$, the three values $\perp, \mathbf{F}, \mathbf{T}$ being incomparable with respect to \prec .

Definition: Relation $x \xrightarrow{b} y$ is defined in table 1, where it is specified in the form of a multivalued function. Its main feature is that it forbids, whenever $b = \top$, that y gets known while x is not.

Properties of scheduling specifications. The following properties hold :

$$\begin{aligned} & x \xrightarrow{b} y \wedge y \xrightarrow{c} z \Rightarrow x \xrightarrow{b \wedge c} z \\ \text{if } b, c \neq ?, \text{ then :} & \\ & x \xrightarrow{b} y \wedge x \xrightarrow{c} y \Rightarrow x \xrightarrow{b \vee c} y \end{aligned} \quad (54)$$

In these equations, $b \wedge c$ and $b \vee c$ are respectively defined as the infimum (resp. supremum) w.r.t. relation “ $<$ ” defined in (53) when both values belong to the subdomain $\{\perp, \top, \mathbf{F}\}$. In fact, we do not need formulas (54) in case b or c are unknown, because the label of a branch is known prior to its extremity, in executable programs equipped with their scheduling specifications as inferred from rules (R-1, . . . , R-4).

x	$?$	\perp	\top
b			
$?$	$?$	$?$	$?$
\perp			
\mathbf{F}			
\mathbf{T}	$?$		

Table 1: *Definition of the dependency* $x \xrightarrow{b} y$. This table gives the result of this multivalued function for its output y . When nothing is written, this means that any value is accepted. If x is boolean, then \top is to be refined as any of the two values $\{\mathbf{F}, \mathbf{T}\}$.

5.2 Circuitfree schedulings

We are given a set of variables x_1, \dots, x_n . Some of them are boolean; for the sake of readability, boolean variables used as labels in scheduling specifications, will be generically denoted by b_1, b_2, \dots . Then we are given 1/ a set of constraints of the form $C(b_1, \dots, b_k)$ on boolean variables restricted to subdomain $\{\perp, \mathbf{T}, \mathbf{F}\}$ of known values; and 2/ a set of scheduling specifications defined on x_1, \dots, x_n . Constraints $C(b_1, \dots, b_k)$ are extended to the “unknown” value by simply assuming $C(b_1, \dots, b_k)$ is satisfied as soon as at least one of the variables b_1, \dots, b_k is “unknown”.

Each dependency is interpreted as specified in Table 1. Thus, together with the boolean constraints of the form $C(b_1, \dots, b_k)$, they specify a subdomain of the product domain \mathcal{D}^n of all possible states. The set of states satisfying these constraints is denoted by \mathcal{S} , and we call it a *scheduling* of x_1, \dots, x_n . States in \mathcal{S} are written s, t, \dots and corresponding interpretations are denoted by s_1, \dots, s_n for short instead of $s[x_1], \dots, s[x_n]$, and similarly for t . The “totally unknown state” :

$$\forall i, s_i = ?, \text{ is denoted by } s?. \quad (55)$$

Two states of \mathcal{S} are said to be *neighbours* if they differ exactly in one variable, we call it their *discriminating* variable. We call a *path* in \mathcal{S} any finite sequence $s(1), s(2), \dots, s(K)$ of neighbouring states belonging to \mathcal{S} .

For s and t two neighbouring states of \mathcal{S} , we write $s \prec t$ if their respective values for their discriminating variable x_i satisfy the relation $s_i \prec t_i$ defined

in (52). A path $s(1), s(2), \dots, s(K)$ such that $s(k) \prec s(k+1)$ is called *increasing*.

A scheduling \mathcal{S} is called *circuitfree* if it is never true in \mathcal{S} that

$$\begin{aligned} x_{i_1} \xrightarrow{b_1} x_{i_2} \xrightarrow{b_2} x_{i_3} \dots x_{i_p} \xrightarrow{b_p} x_{i_1} \\ \text{and} \\ (b_1 \wedge \dots \wedge b_p = \top) \end{aligned} \tag{56}$$

Theorem 4 (circuitfree schedulings) *A scheduling is circuitfree iff, for every state $s \in \mathcal{S}$ satisfying $\forall i : s_i \neq ?$, there is an increasing path linking $s_?$ to s .*

The intuitive interpretation of this theorem is that, for an STS with a circuitfree scheduling, it is possible to compute sequentially without deadlock all variables, starting from the inputs. Each increasing path mentioned in theorem 4 corresponds to one possible sequential execution.

Proof: We first prove the “if” part by contradiction. Assume (56) is violated for some circuit $x_{i_1} \xrightarrow{b_1} x_{i_2} \xrightarrow{b_2} x_{i_3} \dots x_{i_p} \xrightarrow{b_p} x_{i_1}$, i.e., $b_1 \wedge \dots \wedge b_p = \top$ is possible for this circuit in \mathcal{S} . We want to deduce from this assumption that there are states for which all variables are known, but there is no increasing path originating from $s_?$ and terminating at the states in consideration. Without loss of generality, we can restrict \mathcal{S} to those states for which

$$\begin{aligned} \forall i = 1, \dots, p : [b_i = ? \text{ or } b_i = \top] \text{ holds,} \\ \text{the set of such states is called } \mathcal{S}_{(b_1 \wedge \dots \wedge b_p = \top)}. \end{aligned} \tag{57}$$

By table 1, condition $x_{i_1} \xrightarrow{b_1} x_{i_2} \xrightarrow{b_2} x_{i_3} \dots x_{i_p} \xrightarrow{b_p} x_{i_1}$ implies that, on $\mathcal{S}_{(b_1 \wedge \dots \wedge b_p = \top)}$, the following holds :

$$x_{i_1} \succeq x_{i_2} \succeq \dots \succeq x_{i_p} \succeq x_{i_1} ,$$

and thus the x_{i_j} 's are either all unknown, or alternatively all known. Thus there is no increasing path originating from $s_?$ and leading to any known state belonging to $\mathcal{S}_{(b_1 \wedge \dots \wedge b_p = \top)}$. This proves the “if” part.

Next, we prove the “only if” part, also by contradiction. Beforehand, we need a lemma. Two states s and s' are said *complementary* if, for each variable x ,

$$\text{either } s[x] = ? \text{ or } s'[x] = ? .$$

Two states s and s' are said *compatible* if, for each variable x ,

$$\text{either } s[x] = ? \text{ or } s'[x] = ? \text{ or } s'[x] = s[x] .$$

Complementary states are also compatible. For two compatible states s and s' , we define their sum $s \uplus s'$ by:

$$(s \uplus s')[x] = \text{if } s[x] \neq ? \text{ then } s[x] \text{ else } s'[x]$$

Lemma 2 (monotonicity) *Let t_0 and t_1 be two neighbouring states belonging to \mathcal{S} , such that $t_0 \prec t_1$. Let t be a state such that*

1. t_1 and t are complementary,
2. $t_0 \uplus t \in \mathcal{S}$,
3. there is an increasing path contained in \mathcal{S} originating from t_0 and terminating in $t_0 \uplus t$, and
4. $t_1 \uplus t$ satisfies the boolean constraints $C(b_1, \dots, b_k)$ which contribute to the definition of \mathcal{S} .

Then, $t_1 \uplus t \in \mathcal{S}$ and there is an increasing path contained in \mathcal{S} originating from t_1 and terminating in $t_1 \uplus t$.

Proof: Note that $t_0 \uplus t$ is well defined, since t_0 and t are also complementary. Let $t_0 \rightarrow t_0 \uplus t$ denote the path referred to in item 3. Denote by \tilde{t} the state such that $1/\tilde{t}$ and t_0 are complementary, and $2/t_1 = t_0 \uplus \tilde{t}$, such a state exists and is unique. Denote by $t_0 \uplus \tilde{t} \rightarrow t_0 \uplus t \uplus \tilde{t}$ the increasing path obtained by complementing each state belonging to path $t_0 \rightarrow t_0 \uplus t$ by \tilde{t} . This is possible since each intermediate state of path $t_0 \rightarrow t_0 \uplus t$ and \tilde{t} are complementary. We claim that

$$\text{path } t_0 \uplus \tilde{t} \rightarrow t_0 \uplus t \uplus \tilde{t} \text{ is contained in } \mathcal{S}. \quad (58)$$

Clearly, claim (58) is equivalent to the conclusion of the lemma. To prove (58), using item 4, we first note that each state belonging to path $t_0 \uplus \tilde{t} \rightarrow t_0 \uplus t \uplus \tilde{t}$ satisfies the boolean constraints $C(b_1, \dots, b_k)$ which contribute to the definition of \mathcal{S} . We thus only need to check that they also satisfy the dependencies contributing to the definition of \mathcal{S} ; but the latter results from an inspection of table 1. This proves the lemma. \diamond

We now return to the proof of theorem 4 and proceed by steps.

1. Assume $\exists s^* \in \mathcal{S}$ satisfying $\forall i : s_i^* \neq ?$, such that there is no increasing path linking $s_?$ to s^* . Denote by b_1, \dots, b_p the boolean variables such that $b_1 \wedge \dots \wedge b_p = \top$ holds at state s^* . Denote by S the set of states $s \in \mathcal{S}$ such that $s \preceq s^*$. We have $s_? \in S$ and $s^* \in S$. States belonging to S are all compatible.
2. Let $s, s' \in S$ be two states such that increasing paths $s_? \rightarrow s$ and $s_? \rightarrow s'$ are both contained in S . Then we claim that

$$\begin{aligned}
s'' = s \uplus s' \in S, \quad \text{and} \quad \text{there exists an increasing path contained in } S, \\
\text{originating from } s_?, \\
\text{and terminating in } s''
\end{aligned} \tag{59}$$

As all $s \in \mathcal{S}$ satisfy $s \preceq s^*$, they satisfy in particular the boolean constraints $b_1 \wedge \dots \wedge b_p = \top$. Thus we only need to verify the dependencies. There is a unique state $s_0 \in S$ such that $1/ s_0 \in [s_? \rightarrow s] \cap [s_? \rightarrow s']$, and $2/ [s_0 \rightarrow s] \cap [s_0 \rightarrow s'] = \{s_0\}$, meaning that s_0 is the latest point at which the two considered path deviate from each other. Let s_1 be the neighbour state of s_0 belonging to path $[s_0 \rightarrow s']$. Apply lemma 2 with the following substitutions: $t_0/s_0, t_1/s_1, t/\tilde{s}$ such that $s = s_0 \uplus \tilde{s}$. We deduce that path $[s_? \rightarrow s \uplus s_1] \subseteq S$. Then, let s_2 be the neighbour state of s_1 belonging to path $[s_1 \rightarrow s']$, we can repeat the same argument. And we proceed repeatedly in the same way until we prove the claim (59).

3. Consider the set of $s \in S$ for which there exists an increasing path $[s_? \rightarrow s] \subseteq S$. From (59) we know that this set has a unique maximal element s_{\max} for partial order \prec . By hypothesis we have $s_{\max} \prec s^*, s_{\max} \neq s^*$. Thus there are at least two variables, denote them by x and x' , such that

$s_{\max}[x] = s_{\max}[x'] = ?$, but $s[x] = s[x'] \neq ?$ for every $s \in S \setminus [s? \rightarrow s_{\max}]$. Hence, the following holds at each state belonging to S :

$$x \xrightarrow{b} x' \xrightarrow{b} x \quad \text{where} \quad b = b_1 \wedge \dots \wedge b_p = \top$$

Hence condition of circuit freedom is violated on S , and thus it can be violated on \mathcal{S} . This finishes the proof of theorem 4. \diamond

In the sequel, for Φ an STS with scheduling specifications, we shall consider its associated scheduling

$$\mathcal{S}_\Phi \tag{60}$$

which is obtained by keeping, from the set of predicates defining the transition relation of Φ ,

1. the scheduling specifications, and
2. the assertions involving only boolean variables and clocks,

and discarding the other ones.

5.3 Deriving scheduling specifications as causality constraints

In this section, we formally justify rules (21). The principles we follow for our abstraction mechanism are given next:

- (P-1) For x not a boolean variable, we abstract its domain \mathcal{D}_x as the singleton $\{\top\}$, and then extend $\{\top\}$ with the additional values $\{?, \perp\}$.
- (P-2) Within equations of the form “ $y = exp$ ” or “**if** b **then** $y = exp_1$ **else** $y = exp_2$ ” we shall further abstract y by mapping the set $\{\perp, \text{F}, \top\}$ to the single value \mathbf{i} (known). Note the asymmetry of this abstraction principle: for the statement “**if** b **then** $y = x$ ” where x, y are booleans, we abstract y but not x .

(P-3) Since we are interested in causality constraints, we only need to keep track of configurations for which y cannot be known, i.e., $y = ?$ is the only allowed possibility. For other configurations, we weaken the constraint on y to “ y unconstrained”, which is depicted in the tables by an empty box.

We now proceed on deriving the scheduling associated to each primitive s-statement, using **(P-1,2,3)**. We use the notation: $?, \perp$ to indicate that, for the considered configuration, either $y = ?$ or $y = \perp$ holds, and similarly for other cases.

Lemma 3 *The following holds :*

$$x \xrightarrow{b} y \Rightarrow b \longrightarrow h_y$$

Proof: by inspection of table 1.

Lemma 4 *The following holds :*

$$h_x \longrightarrow x$$

Proof: by inspection of the following tables (the first table relates x to h_x , as extended to unknown values):

h_x	$?$	\perp	\top
x	$?$	$?, \perp$	$?, \top$

abstracted as (using **P-2**) :

h_x	$?$	\perp	\top
x	$?$	$?, \perp$	$?, \perp$

which is equal to :

h_x	$?$	\perp	\top
x	$?$		

which turns out to be equivalent to $h_x \longrightarrow x$ by table 1.

Lemma 5 *The following holds :*

$$(f) : \begin{cases} y = f(u, v) \\ h_u = h_v = h_y \end{cases} \Rightarrow (u, v) \xrightarrow{h_y} y$$

Proof: by inspection of table 1 and of the following tables (# denotes a prohibited value):

abstraction of (f) , using **(P-1)** :

u	?	\perp	\top
v			
?	?	$?, \perp$?
\perp	$?, \perp$	\perp	#
\top	?	#	\top

using **(P-2)** :

u	?	\perp	\top
v			
?	?	$?, i$?
\perp	$?, i$	i	#
\top	?	#	i

using **(P-3)** :

u	?	\perp	\top
v			
?	?		?
\perp			
\top	?		

which is equivalent to the formulas of the conclusion of the rule of lemma 5.

Lemma 6 *The following holds :*

$$[\text{if } b \text{ then } x = u] \wedge [\text{if } b \text{ then } h_x = h_u] \Rightarrow \begin{cases} u \xrightarrow{b \wedge h_u} x \\ b \xrightarrow{h_b \wedge h_u} h_x \\ h_u \xrightarrow{b \wedge h_u} h_x \end{cases}$$

Proof: by inspection of table 1 and of the following two tables. These tables define the possible values, of x and h_x respectively, for $[\text{if } b \text{ then } x = u] \wedge [\text{if } b \text{ then } h_x = h_u]$:

	u			
b		?	\perp	\top
?		?	?, \perp	?
\perp		?, \perp	?, \perp	?, \perp
\top		?	?, \perp	?, \top
F		?, \perp	?, \perp	?, \perp

	h_u			
b		?	\perp	\top
?		?	?, \perp	?
\perp		?, \perp	?, \perp	?, \perp
\top		?	?, \perp	?, \top
F		?, \perp	?, \perp	?, \perp

Applying principles **(P-2,3)** then yields the formulas corresponding to the conclusion of the rule of lemma 6. Note the asymmetry between x and u , while statements $x = u$ and $u = x$ are clearly identical. This asymmetry is due to principle **(P-2)** for STS abstraction.

5.4 Correct programs

In this subsection, we formally state and prove the result establishing the link between circuit freedom and executable STS.

Theorem 5 (correct programs) *Let \mathbf{P} be an STS satisfying the following conditions :*

1. *For each statement of \mathbf{P} , the scheduling specifications derived from applying the rules of lemmas 3, 4, 5, 6 are also statements of \mathbf{P} .*
2. *The scheduling $\mathcal{S}_{\mathbf{P}}$ (cf. (60)) defined by \mathbf{P} is circuitfree.*
3. *There is no multiple definition of a variable, meaning that, whenever*

$$\begin{aligned} & \mathbf{if} \ b_1 \ \mathbf{then} \ x = \text{exp}_1 \\ \wedge \quad & \mathbf{if} \ b_2 \ \mathbf{then} \ x = \text{exp}_2 \end{aligned}$$

is part of \mathbf{P} , then :

$$b_1 \wedge b_2 = \top \ \text{never holds.}$$


Then :

1. *As far as control is concerned, the inputs of \mathbf{P} are the source nodes of the dependency graph.*
2. *Input values are those variables which never occur on the left-hand side of statements of the form “ $x = \text{exp}$ ”.*
3. *For each given input control history of \mathbf{P} and compatible input value history, there is exactly one run of \mathbf{P} , i.e., \mathbf{P} is deterministic.*

NOTA : Clearly, theorem 5 provides us with a sufficient condition, this condition is not necessary. Furthermore, the rules for inferring scheduling specifications as causality constraints is bound to the syntax, not to the semantics of the program. In particular, from statement “ **if** b **then** $x = u$ ”, we choose to infer dependency $u \xrightarrow{b \wedge h_u} x$ but not the symmetric one in which x and u are exchanged. This means that, while P may not satisfy the assumptions of theorem 5 for a given syntactic form of P , it may satisfy them after a proper rewriting into a semantically equivalent form. Here, semantic equivalence means identical runs when scheduling specifications are discarded.

Proof: It is organized into several steps.

1. With the formula $x \xrightarrow{b} y$ we associate the following automaton :



Transitions are labelled with actions. Label “*set x*” indicates that variable x is set to an arbitrary value of its (extended) domain $\mathcal{D}_x \cup \{\perp\}$. States are labelled with those variables that are ?, i.e., have not been set. This automaton is the most permissive one with the following properties :

- (a) states are valued with configurations of the triple (x, b, y) that are compatible with the scheduling constraint $x \xrightarrow{b} y$.
- (b) Variables are set sequentially.
- (c) All variables are eventually set.

Thus each path of this automaton specifies an evaluation scheme for the triple (x, b, y) which is compatible with the considered scheduling specification. Conversely, any correct evaluation scheme for triple (x, b, y) can be specified in this way. We call this automaton the *execution automaton* associated to scheduling specification $x \xrightarrow{b} y$.


2. To each primitive statement we associate the conjunction of its causality constraints and possible constraints involving clocks and boolean variables, and we take the product of associated execution automata. The paths of the resulting automaton specify all correct schedulings to evaluate the involved variables. We call the resulting product automaton the *execution automaton* associated to the considered primitive.
3. Then we take the product of the execution automata associated to each statement. By theorem 4 we know that, for each tuple of variables which satisfies the specification, there is a path of the product automaton which originates from its initial state and terminates at the final state in which all variables are set, meaning that all variables of the considered tuple are sequentially set.
4. Finally, we refine the transition labels of the form “*set x*” etc., by assigning to x etc the value specified by the program. As source nodes of the dependency graph are set first, they appear as inputs of \mathbf{P} for its control. Also, variables u that are set and do not occur on the left-hand side of any statement $u = \textit{expression}$ must be read from the environment: their values are inputs of the considered program \mathbf{P} . Finally, thanks to condition 3 of theorem 5, actions of the form “*set x*” etc., are refined into single writings. This finishes the proof of the theorem. \diamond

We illustrate this technique on the following simple STS :

$$y = f(u, v) \quad \wedge \quad h_u = h_v = h_y =_{\text{def}} h .$$

The causality constraint and associated execution automaton are :

$$\begin{aligned}
& h \longrightarrow (u, v, y) \\
\wedge & (u, v) \xrightarrow{h} y \\
\wedge & h_u = h_v = h_y =_{\text{def}} h
\end{aligned}$$



Clock h is the activation clock. The refined execution automaton is obtained by replacing $set\ u$ and $set\ v$ by $read\ u$ and $read\ v$, and $set\ y$ by the assignment $y := f(u, v)$.

6 Conclusion

Our contribution can be summarized as follows :

- We have proposed STS with scheduling specifications as a paradigm for causality analysis, STS abstraction, separate compilation and reuse.
- We have characterized those STS for which asynchronous and synchronous semantics are equivalent in some precise meaning.

We advocate system design methodology based on the synchronous paradigm, possibly followed by a provably correct desynchronization. Advantages of this approach are numerous, they are listed below according to the different phases of the design :

Specification : designing within the synchronous paradigm allows the designer to exploit the simplicity and elegance of compositionality of synchronous specifications. In addition, specification can be performed in-

dependently from the execution architecture ; therefore, upgrading an execution architecture does not require redesigning the specifications.

Verification :

- In the synchronous paradigm, composition of specifications and composition of properties are both performed by using the composition “ \parallel ” of STS. This facilitates reasoning in general, and in particular compositional reasoning.
- For endo/isochronous STS, proofs based on the synchronous semantics carry over without modifications to asynchrony. For such systems, verifications can be performed within the synchronous framework. This allows to avoid state explosion resulting from the use of the asynchronous interleaving semantics.

Abstraction, modularity, and reuse :

- Scheduling specifications provide the adequate notion of abstraction for separate compilation. It allows the designer to check the correctness of component encapsulation at systems integration phase.
- STS with scheduling specifications can be composed using a proper generalization of the composition “ \parallel ” of STS. Thus advantages of compositionality naturally extend to STS with scheduling specifications.
- The structuration of specifications into scheduler and tasks allows us to define proper reusable modules. Of course, if assumptions are available on the possible behaviours of the environment, then larger modules can be stored as object code for further reuse.

GALS networks : the elegant feature is that isochrony is a *local* property within a network of components. As isochrony is compositional, adding a new component Φ_{new} to a pre-existing GALS network $(\Phi_i)_{i=1,\dots,n}$ while preserving its GALS nature, only requires to check whether pairs $(\Phi_{\text{new}}, \Phi_i)$ are isochronous, for each Φ_i having direct communication with Φ_{new} in the extended network. Thus GALS designs can be built *compositionally*, it is not needed to desynchronize at once the whole synchronous design.

Thanks to the outcomes of the SACRES project, the above approach is supported by the SIGNAL-V4 language⁸, and by the DC₊ common format for synchronous languages [DC+ Sacres 1996]. SIGNAL-V4 and the DC₊ format both are concrete implementations of our STS model. This includes scheduling specifications, which are available as primitive statements in both formalisms.

In particular, the 1999 release of SILDEX [Sildex]⁹ implements distributed code generation based on the approach presented in this paper. The target architectures above all else are POSIX compliant real-time OS.

The new SIGNAL-V4 compiler developed at Inria implements the whole methodology, including separate compilation. Services for architecture generation are also provided, using our notion of abstraction.

Research perspectives. Further work is needed to show that the above principles are viable for generating architectures built up from pre-existing C/C++/Java/... modules. Then, not all communication media or operating systems provide services satisfying the requirements of our theory of desynchronization, namely: no loss of messages, first-in/first-out semantics for each individual channel. Additional work is needed for getting a full implementation on each different type of distributed architecture; this can be very easy (writing a few generic drivers, e.g., for POSIX), or can be more demanding when adequate services are not provided by the architecture, and thus need to be emulated.

ACKNOWLEDGEMENT. *The authors are gratefully indebted to Michael Siegel for a thorough reading and detailed comments, and in particular the discovery of several inconsistencies in an earlier version of this manuscript in the formal study of causality.*

References

[Aabelsberg and Rozenberg, 1988] I.J. Aabelsberg, and G. Rozenberg, “Theory of traces”, *Theoretical Computer Science*, 60, 1–82, 1988.

⁸*Loic Besnard and other members of the “EpAtr” team at IRISA are gratefully acknowledged for the development of this environment.*

⁹the SILDEX tool is a commercial tool for reactive systems design based on the SIGNAL language. It is marketed by TNI, Brest, France.

- [Alur and Henzinger 1996] R. Alur and T. A. Henzinger, “Reactive Modules”, *Proceedings of the 11th IEEE Symposium on Logic in Computer Science* (LICS), 207–218, 1996, extended version submitted for publication.
- [Amagbegnou et al., 1994] T.P. Amagbegnou, L. Besnard and P. Le Guernic, “Arborescent canonical form of boolean expressions”, Inria Research Report n°2290, June 1994.
- [Amagbegnou et al., 1995] T.P. Amagbegnou, L. Besnard and P. Le Guernic, “Implementation of the dataflow language SIGNAL”, in *Programming Languages Design and Implementation*, ACM, 163–173, 1995.
- [Aubry 1997] P. Aubry, “Mises en œuvre distribuées de programmes synchrones”, PhD Thesis, Univ. Rennes I, 1997.
- [Benveniste and LeGuernic 1990] A. Benveniste and P. Le Guernic, “Hybrid dynamical systems theory and the Signal language”, *IEEE Transactions on Autom. Control*, 35 No 5, 535–546, may 1990.
- [Benveniste and Berry, 1991] A. Benveniste and G. Berry, “Real-Time systems design and programming”, *Another look at real-time programming*, special section of *Proc. of the IEEE*, vol. 9 n° 9, September 1991, 1270–1282.
- [Benveniste LeGuernic and Jacquemot, 1991] A. Benveniste, P. Le Guernic, and C. Jacquemot. “Synchronous programming with events and relations: the SIGNAL languages and its semantics”, *Sci. Comp. Prog.*, 16:103–149, 1991.
- [Benveniste et al., 1992] A. Benveniste, P. Le Guernic, Y. Sorel, and M. Sorine, “A denotational theory of synchronous communicating systems”, *Information and Computation*, Vol. 99 No 2, 192–230, August 1992.
- [Benveniste Caspi et al., 1994] A. Benveniste, P. Caspi, N. Halbwachs, and P. Le Guernic, “Data-flow synchronous languages”, In *A Decade of Concurrency, reflexions and perspectives, REX School/Symposium*, pages 1–45, LNCS Vol. 803, Springer Verlag, 1994.

- [Berry, 1989] G. Berry, “Real time programming: Special purpose or general purpose languages”, In *IFIP World Computer Congress*, San Francisco, 1989.
- [Berry, 1995] G. Berry, *The Constructive Semantics of Esterel*, Draft book, <http://www.inria.fr/meije/esterel>, December 1995.
- [Berry and Sentovich 1998] G. Berry and E.M. Sentovich, “An implementation of constructive synchronous programs in polis”, manuscript, November 1998.
- [Caillaud et al., 1997] B. Caillaud, P. Caspi, A. Giraud, and C. Jard, “Distributing automata for asynchronous networks of processors”, *European Journal on Automated Systems (JESA)*, Hermes, 31(3), 503–524, May 1997.
- [Caspi 1992] P. Caspi, “Clocks in Dataflow languages”, *Theoretical Computer Science*, vol. 94:125–140, 1992.
- [Clerbout and Latteux, 1987] M. Clerbout, and M. Latteux, “Semi-commutations”, *Information and Computation*, 73, 59–74, 1987.
- [de Roever et al., Eds, 1998] W-P. de Roever, H. Langmaack, and A. Pnueli, Eds. *Compositionality: the significant difference*, Proc. of the International Symposium COMPOS’97, Bad Malente, Germany, Sept. 1997, LNCS vol 1536, Springer Verlag, 1998.
- [DC+ Sacres 1996] Sacres consortium, The Declarative Code DC+, Version 1.2, May 1996; Esprit project EP 20897: Sacres, see <http://www.tni.fr/sacres/>
- [LeGuernic and Gautier, 1991] P. Le Guernic and T. Gautier, “Dataflow to von Neumann: the SIGNAL approach”, in *Advanced topics in dataflow computing*, L. Biv and J-L. Gaudiot Eds., Prentice Hall, 413-438, 1991.
- [Halbwachs, 1993] N. Halbwachs, *Synchronous programming of reactive systems*,. Kluwer Academic Pub., 1993.
- [Lamport, 1983a] L. Lamport, “Specifying concurrent program modules”, *ACM Trans. on Prog. Lang. and Sys.*, 5(2):190-222, 1983.

- [Lamport, 1983b] L. Lamport, “What good is temporal logic?”, In *Proc. IFIP 9th World Congress*, R.E.A. Mason (Ed.), North Holland, 657-668, 1983.
- [LeGuernic et al., 1991] P. Le Guernic, T. Gautier, M. Le Borgne, C. Le Maire, “Programming real-time applications with SIGNAL”, *Another look at real-time programming*, special section of *Proc. of the IEEE*, vol. 9 n° 9, September 1991, 1321–1336.
- [Maffeis and LeGuernic, 1994] O. Maffeis and P. Le Guernic, “Distributed implementation of Signal: scheduling and graph clustering”, in: *3rd International School and Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*, Lecture Notes in Computer Science 863, Springer Verlag, 149–169, Sept. 1994.
- [Manna and Pnueli 1992] Z. Manna and A. Pnueli, *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, New York, 1992.
- [Manna and Pnueli 1995] Z. Manna and A. Pnueli, *The Temporal Logic of Reactive and Concurrent Systems: Safety*. Springer-Verlag, New York, 1995.
- [Sorel and Lavarenne] Y. Sorel and C. Lavarenne, “SynDEx v4.2 User Guide”,
<http://www-rocq.inria.fr/syndex/.articles/doc/doc/SynDEx42.html>
- [Sorel 1996] Y. Sorel, “Sorel: Real-time embedded image processing applications using the A3 methodology”, *Proc. IEEE International Conf. on Image Processing*, (Lausanne, September 1996).
- [Sildex] TNI, SILDEX tool, see <http://www.tni.fr/indexgb.html>