



HAL
open science

Compositionality in Dataflow Synchronous Languages: Specification & Code Generation

Albert Benveniste, Paul Le Guernic, Pascal Aubry

► **To cite this version:**

Albert Benveniste, Paul Le Guernic, Pascal Aubry. Compositionality in Dataflow Synchronous Languages: Specification & Code Generation. [Research Report] RR-3310, INRIA. 1997. inria-00073379

HAL Id: inria-00073379

<https://hal.inria.fr/inria-00073379>

Submitted on 24 May 2006

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 & code generation

Albert Benveniste , Paul Le Guernic , Pascal Aubry

N° 3310

November 1997

_____ THÈME 1 _____



*Rapport
de recherche*



Compositionality in dataflow synchronous languages : specification & code generation

Albert Benveniste , Paul Le Guernic * , Pascal Aubry †

Thème 1 — Réseaux et systèmes
Projet EpAtr

Rapport de recherche n3310 — November 1997 — 40 pages

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 the issues of compositionality aiming at modular code generation, for dataflow synchronous languages. As careless storing of object code for further reuse in systems design 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. Then we study separate compilation for synchronous programs, and we discuss the issue of distributed implementation using an asynchronous medium of communication ; for both topics we provide a complete formal study. Corresponding algorithms are currently under development in the framework of the *DC+ common format* for synchronous languages.

Key-words: synchronous languages, modularity, compositionality, code generation, distributed code generation, causality analysis, separate compilation, desynchronisation.

(Résumé : *tsvp*)

This paper is an extended 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, and the background on labelled partial orders is mostly acknowledged to Paul Caspi.

* A.B. and P.L.G. are with Irisa/Inria, Campus de Beaulieu, 35042 Rennes cedex, France; email: name@irisa.fr

† Irisa/Ifsic, Campus de Beaulieu, 35042 Rennes cedex, France; email: name@ifsic.fr

Unité de recherche INRIA Rennes
IRISA, Campus universitaire de Beaulieu, 35042 RENNES Cedex (France)
Téléphone : 02 99 84 71 00 - International : +33 2 99 84 71 00
Télécopie : 02 99 84 71 71 - International : +33 2 99 84 71 71

Compositionnalité dans les langages synchrones flots de données : spécification et génération de code

Résumé : La compositionnalité est une bonne traduction mathématique du concept de modularité, concept important pour la conception de systèmes embarqués. Nous étudions la compositionnalité pour les langages flots de données synchrones, avec comme objectif la spécification et la génération de code distribué. Nous exhibons un formalisme uniforme pour traiter des questions de causalité, ordonnancement, et communication : la notion de *spécification d'ordonnancement* ou de *dépendance*. Nous appuyant sur cette notion, nous conduisons une étude formelle complète de la compilation séparée et de la désynchronisation en vue de la génération de code distribué. Les algorithmes correspondants sont en cours de développement autour du format commun DC_+ des langages synchrones.

Mots-clé : programmation synchrone, dataflow, spécification, génération de code, génération de code distribué, compositionnalité, modularité, ordres partiels, causalité, désynchronisation.

Contents

1	Motivations	4
2	The essentials of the synchronous paradigm	5
3	Specification : Symbolic Transition Systems (STS) [Pnu97]	6
4	Compositional reasoning on causality and scheduling specifications	9
4.1	What is the problem?	9
4.2	Scheduling specifications	10
5	Modular and distributed code generation	16
5.1	Relaxing synchrony	16
5.2	Generating scheduling for separate modules	18
5.3	The bottom line: modular and distributed code generation	20
6	Conclusion	21
A	Appendix : formal study of causality and scheduling specifications	22
A.1	Dependencies	22
A.2	Circuitfree schedulings	23
A.3	Deriving dependencies as causality constraints	25
A.4	Correct programs	27
B	Appendix : formal study of desynchronization (co-authored with Benoît Caillaud¹)	30
B.1	Desynchronizing STS: from STS to infinite partial orders, and their asynchronous composition	30
B.2	Endochrony	32
B.3	Handling endochrony in practice	36
B.3.1	Checking endochrony	36
B.3.2	Enforcing endochrony	38

¹IRISA/INRIA, name@irisa.fr

1 Motivations

Modularity is advocated as the ultimate solution for the design of large systems, and this holds in particular for embedded systems, and their software and architecture. Modularity allows the designer to scale down design problems, and facilitate reuse of predefined modules.

The mathematical translation of the concept of modularity is often that of *compositional*ity. Paying attention to the composition of specifications [MP92] is central to any system model involving concurrency or parallelism. More recently, significant effort has been devoted toward introducing compositionality in verification with aiming at deriving proofs of large programs from partial proofs involving (abstractions of) components [MP95].

Compilation and code generation has been given less attention from the same point of view, however. 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 only using these interface descriptions, while still guaranteeing that final encapsulation of actual code within these black-boxes together with their composition, will maintain correctness of the design w.r.t. specification.

This paper is devoted the issues of compositionality aiming at modular code generation, for dataflow synchronous languages. As dataflow synchronous is rather a paradigm more than a few concrete languages or visual formalisms [BB91], it was desirable to abstract from such and such particular language. Thus we have chosen to work with a formalism proposed by Amir Pnueli, that of Symbolic Transition Systems (STS [Pnu97]), which is at the same time very lightweight, and fully general to capture the essence of synchronous paradigm.

Using this formalism, we study composition of specifications, a very trivial topic indeed. Most of our effort is then devoted to issues of compositionality that are critical to code generation. As careless storing of object code for further reuse in systems design 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. We pay strong attention to this notion, a whole appendix is devoted to its mathematical analysis, using a technique not unlike the one for the causality analysis of Esterel [Ber95]. 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, and again, a full appendix is devoted to its formal study.

Besides this, let us mention that Amir Pnueli & coworkers [Pnu97] have introduced a more elaborated version of our STS formalism, for the purpose of investigating issues of compositionality in proofs involving liveness properties.

This work was initiated within the context of the SIGNAL synchronous language by P. Le Guernic and its students B. Le Goff, O. Maffeis [MLG94], and P. Aubry [Aub97] who finally implemented these ideas.

2 The essentials of the synchronous paradigm

There has been several attempts to characterize the essentials of the synchronous paradigm [BB91] [Halb93]. With some experience and after attempts to address the issue of moving from synchrony to asynchrony (and back), we feel the following features are indeed essential for characterizing this paradigm :

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

$$P = R^\omega$$

where R denotes the family of possible reactions².

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 selfexplanatory. 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. When it is defined, parallel composition is always given by taking the conjunction of associated reactions :

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

Typically, if specifying is the intention, then the above formula is a perfect definition of parallel composition. In contrast, if programming was the intention, then the need for this definition to be compatible with an operational semantics very much complicates the “when it is defined” prerequisite³.

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

Clearly, these remarks call for a common format implementing this paradigm, the DC/DC₊ format [DC96] has been proposed with this objective. Also, this calls for a simplest possible formalism with the above features, on which fundamental questions should be investigated (the purpose of this basic synchronous formalism would not be to allow better specification or programming, however): the STS formalism we describe next has this in its objectives.

²In 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.

³For instance, most of the effort related to the semantics of ESTEREL has been directed toward solving this issue satisfactorily [Ber95].

3 Specification : Symbolic Transition Systems (STS) [Pnu97]

Symbolic 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 $u \in \mathcal{V}$ a value $s[u]$ 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 .

Following [Pnu97]⁴, we define a *Symbolic Transition System* (STS) to be a system

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

consisting of the following components :

- V is a finite set of typed *variables*,
- $\Theta(V)$ is an assertion characterizing *initial states*.
- $\rho = \rho(V^-, V)$ is the *transition relation* relating previous and current states s^- and s , by referring to both past⁵ and current versions of variables V^- and V . 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 $\rho(s^-[V], s[V]) = \mathsf{T}$, we say that state s^- is a ρ -*predecessor* of state s .

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

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

The *composition* of two STS $\Phi = \Phi_1 \wedge \Phi_2$ 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 of initial and transition relations.

⁴The talented reader will also notice some close relation to the TLA model of L. Lamport.

⁵Usually, variables and *primed* variables 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 formulæ 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 formulæ.

Notations for STS: we shall use the following generic notations in the sequel :

- c, v, w, \dots denote STS *variables*.
- for v a variable, $h_v \in \{\top, \perp\}$ denotes its *clock*:

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

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

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

with the convention that $s_0[\xi_v^-] = \perp$, i.e., ξ_v is absent before the 1st occurrence of v , and is always present after the 1st occurrence of v . State variable ξ_v is assumed to be *local*, i.e., is never shared among different STS, and thus state variables play no role for STS composition.

Examples of Transition Relations :

- a selector :

$$\text{if } b \text{ then } z = u \text{ else } z = v .$$

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

- decrementing a register :

$$\text{if } h_z \text{ then } v = \xi_z^- - 1 \text{ else } v = \perp ,$$

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

$$\text{if } h_z \text{ then } v = \xi_z^- \text{ else } v = \perp .$$

Note that both statements imply the equality of clocks :

$$h_z \equiv h_v .$$

- testing for a property :

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

Note that a consequence of this definition is, again,


$$h_z \equiv h_v .$$

- a synchronization constraint :

$$(b = \top) \equiv (h_u = \top) ,$$

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

Putting things together yields the STS :



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

A run of this STS for the z variable is depicted on the figure above. Each time u is received, z is reset and gets 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 latter instant, 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 these primitives, dataflow synchronous languages such as LUSTRE [Halb93] and SIGNAL [LG91] are easily encoded.

Open environments and modularity. As modularity is wanted, it is desirable that the pace of an STS is local to it rather than global. Since any STS is subject to further composition in some yet unknown environment, this makes the requirement of having a global pace quite inconvenient. This is why *we prohibit the use of clocks that are always present*. This has several consequences. First, it is not possible to consider the “complement of a clock” or the “negation of a clock”: this would require referring to the always present clock. Thus, as will be revealed by our examples, clocks will always be variables, and we shall be able to relate clocks only using \wedge (intersection of instants of presence), \vee (union of instants of presence), and \setminus (set difference of instants of presence).

Stuttering. In the same vein, it should be permitted, for an STS, to do nothing while the environment is possibly working. This feature has been yet identified in the litterature and is known as *stuttering* invariance or robustness [Lam83a, Lam83b]. It is central to TLA, where it is understood that a transition with no event at all and no change of states is always legal.

For an STS Φ , stuttering invariance is defined as follows : if

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

is a run of Φ , so is

$$\sigma' : s_0, \underbrace{\perp_0, \dots, \perp_0}_{0 \leq \#\{\perp_0\} < \infty}, s_1, \perp_1, \dots, \perp_1, s_2, \perp_2, \dots, \perp_2, \dots, \quad (3)$$

where symbol \perp_i denotes a *silent* state in which all state variables keep the value they had at state s_i , while other variables take the value \perp . The number of inserted silent states is ≥ 0 but finite. This models that the considered STS can do nothing for any arbitrary but finite “duration”.

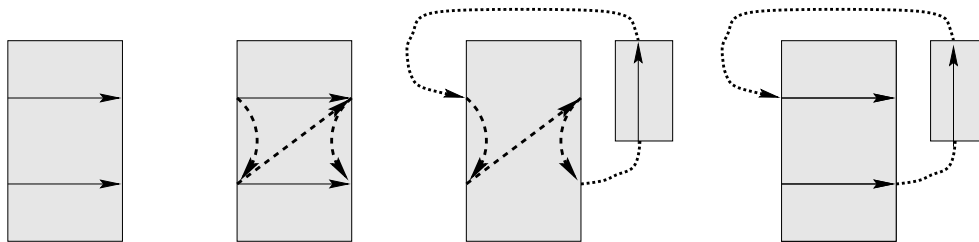
Stuttering invariance is not hardwired into our STS formalism, but a quick inspection of all the statements we have introduced in our example reveals that any composition of them is stuttering invariant. More generally, any STS not involving the always present clock (be it directly or indirectly) is stuttering invariant.

4 Compositional reasoning on causality and scheduling specifications

4.1 What is the problem ?

Basically, the problem is twofold : 1/ bruteforce separate compilation can cause deadlocks, 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 :

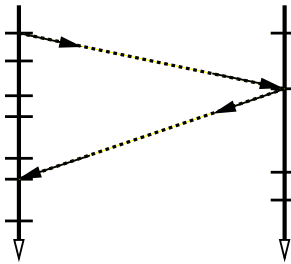


The first diagram depicts the “dependencies” associated with some STS specification : the 1st output needs the 1st input for its computation, and the 2nd output needs the 2nd input for its computation. The second diagram shows a possible scheduling, corresponding to the standard scheduling : 1/ read 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 however later see that this does not need to be the case, however.

Desynchronisation. 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*.

4.2 Scheduling specifications

Preorders and partial orders to model causality relations, schedulings, and communications. 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 [AR88] has been extended to so-called “semi-commutations” [CL87], and this technique has been recently

applied to the implementation of reactive automata on distributed architectures [CCGJ97]. Causality preorder relations have also been used in a different way in [BCHIG94], 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.

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*.

The *conjunction* of two preorders is the minimal preorder which is an extension of the two considered conjuncts.

STS with scheduling specifications. Now we consider STS $\Phi = \langle V, \Theta, \rho \rangle$ as before, but with the following additional feature: as part of

- $\Theta(V)$ (relation defining initial states), and
- $\rho(V^-, V)$ (transition relation),

we have *preorders* denoted by

$$x \preceq y \quad \text{for } x, y \in V ,$$

specified via (*possibly cyclic*) directed graphs:

$$x \rightarrow y \quad \text{for } x, y \in V .$$

STS involving such type of preorder relation shall be called in the sequel *STS with scheduling specifications*. As preorders are just like any other relation, STS with scheduling specifications are just like any other STS, hence they inherit their properties, in particular *they can be composed*.

Notations for scheduling specifications : for b a variable of type $bool \cup \{\perp\}$, and u, v variables of any type,

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

is denoted by

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

Note the following:

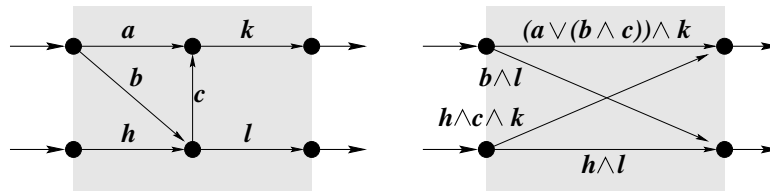
$$\underbrace{u \xrightarrow{-b} v}_{b=\text{false}} \quad \neq \quad \underbrace{u \xrightarrow{\bar{b}} v}_{b=\text{false} \vee b=\perp} \quad !!!$$

In appendix A, it is shown that scheduling specifications have the following properties :

$$x \xrightarrow{b} y \wedge y \xrightarrow{c} z \Rightarrow x \xrightarrow{b \wedge c} z \quad (4)$$

$$x \xrightarrow{b} y \wedge x \xrightarrow{c} y \Rightarrow x \xrightarrow{b \vee c} y \quad (5)$$

Properties (4,5) 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 (4,5).

Inferring scheduling specifications from causality analysis. The idea supporting causality analysis of an STS specification is quite simple. On the one hand, a transition relation involving only the types “pure” and “boolean” can be solved by unification and thus made executable. On the other hand, a transition relation involving arbitrary types is abstracted as term rewriting, encoded via directed graphs. For instance, relation $y = 2uv^2$ (involving, say, real types) is abstracted as $(u, v) \longrightarrow y$, since y can be substituted by expression $2uv^2$. On the other hand, the relation $w + uv^2 \geq 0$ (again involving real types) is abstracted as the full directed graph with vertices (u, v, w) , as no term rewriting is possible. A difficulty arises from the hybrid nature of general STS, in which boolean variables can be computed from the evaluation of non-boolean expressions (e.g., $b = (x \geq 0)$), and then used as a control variable.

We now provide a technique for inferring schedulings from causality analysis for STS specified as conjunctions of the particular set of primitive statements we have introduced so far. In formulæ (6), 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 formulæ (6), labels of schedulings are expressions involving variables in the domain $\{\perp, F, T\}$ ordered by $\{\perp < F < T\}$; with this in mind, expressions involving the symbols “ \wedge ” (min) and “ \vee ” (max) have a clear meaning.

$$\begin{aligned}
 \text{(R-1)} \quad & \forall u \quad h_u \longrightarrow u \\
 \text{(R-2)} \quad & \text{if } b \text{ then } w = u \quad \Rightarrow \quad \left\{ \begin{array}{l} b \xrightarrow{h_b \wedge (h_u \vee h_v)} h_w \\ h_u \xrightarrow{b \wedge h_u} h_w \\ h_v \xrightarrow{\bar{b} \wedge h_v} h_w \\ u \xrightarrow{b \wedge h_u} w \\ v \xrightarrow{\bar{b} \wedge h_v} w \end{array} \right. \quad \text{(6)} \\
 \text{else } w = v \\
 \text{(R-3)} \quad & u \xrightarrow{b} w \Rightarrow b \longrightarrow h_w \\
 \text{(R-4)} \quad & \left. \begin{array}{l} w = f(u_1, \dots, u_k) \\ h_w = h_{u_1} = \dots = h_{u_k} \end{array} \right\} \Rightarrow u_i \xrightarrow{h_w} w
 \end{aligned}$$

Rules (R-1,...,R-4), as well as the theorem to follow, are formally justified in appendix A. In this appendix, a precise definition of “*deterministic*” is also given.

Theorem 1 (executable STS) For P an STS,

1. Apply Rules (R-1,...,R-4) until fixpoint is reached: this yields an STS we call **sched**(P).

2. A sufficient condition for P to have a unique deterministic run is:

(a) **sched**(P) is provably circuitfree at each instant, meaning that it is never true that

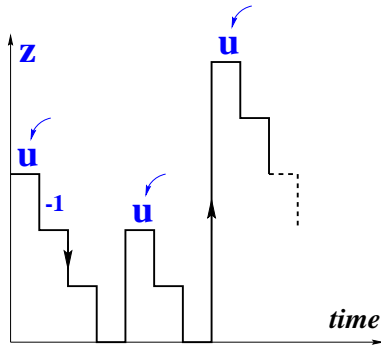
$$\begin{array}{c}
 x_1 \xrightarrow{b_1} x_2 \xrightarrow{b_2} x_1 \\
 \text{and} \\
 (b_1 \wedge b_2 = \text{T})
 \end{array}$$

unless $x_1 = x_2$ holds.

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

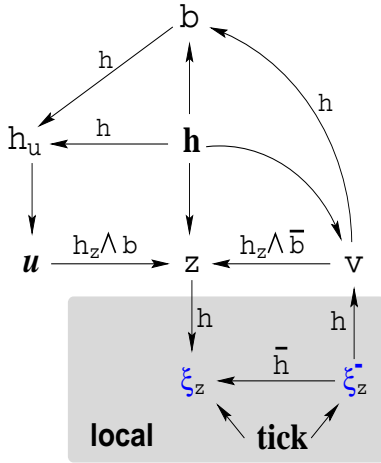
$$\begin{array}{l}
 \text{if } b_1 \text{ then } x = \text{exp}_1 \\
 \wedge \text{ if } b_2 \text{ then } x = \text{exp}_2
 \end{array}$$

The full example:



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

Applying the rules (R-1,...,R-4) for inferring schedulings from causality, we get :



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

Note the change in control: { 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 when present. Reading of the u occurs on demand, when condition b is true. Thus we call such an STS “proactive”.

Summary. What do we get at this stage?

1. STS composition is just the conjunction of constraints.
2. Since preorders are just relations, scheduling specifications do compose as well.
3. As 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 performed⁶. But this latter operation is global and not compositional: here we reach the limits of brute-force compositionality.

5 Modular and distributed code generation


Two major issues need to be considered:

- Relaxing synchrony* is needed if distribution over possibly asynchronous media is desired without paying the price for maintaining the strong synchrony hypothesis via costly protocols.
- Designing modules equipped with proper interfaces for subsequent reuse*, and generating a correct scheduling and communication protocol for these modules, is the key to modularity.

We consider these two issues next.

5.1 Relaxing 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 of “which are the variables present 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 questionmark indicates that it is generally not possible, in an asynchronous environment, to decide upon presence/absence of a signal relatively to another one. The major problem is that of “testing for absence” as being part of the control. While this is perfectly sound in a synchronous paradigm, this 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 previous state, and 2/ the values possibly carried by environment signals, but not on the presence/absence status

⁶This is part of the job performed by the SIGNAL compiler’s “clock calculus”.

of these signals. An endochronous STS can work as a (synchronous) module working in a distributed execution using an asynchronous communication medium, provided that this medium satisfies the two requirements of 1/ not losing messages, and 2/ not changing the order of messages.


An example of an STS which is “exochronous” (i.e., not endochronous) 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 an endochronous STS:



In the diagram on the left hand side, three different clocks are source nodes of the directed graph. This means that the primary 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 boolean variable b , which is computed internally.

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

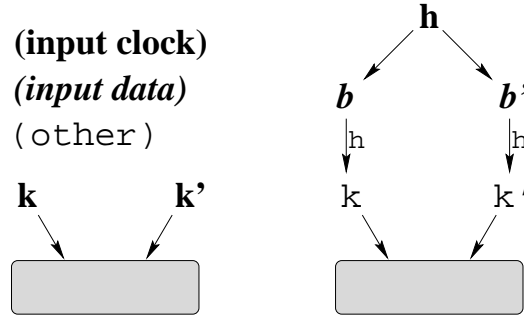
“ if b then get_u ”:



In the left diagram, a history of this statement is depicted, showing the successive instants (or reactions) separated by thick dashed lines. In the middle, an instant has been twisted, and in the last one, thick dashed lines have been removed. Clearly, no information has been lost: we know that u should be got exactly when $b = T$, and thus it is only needed to wait for

b in order to know whether u is to be waited for also. A formal study of desynchronization and endochrony is presented in Appendix B.

Moving from exochronous to endochronous is easily performed, we only show one typical but simple example :



The idea is to add to the considered STS a monitor which delivers the information of presence/absence via the b, b' boolean variables with identical clock h , i.e., $\{k = \tau\} \equiv \{b = \tau\}$, 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.

5.2 Generating scheduling for separate modules


Relevant target architectures for embedded applications typically are 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 rules (R-1,...,R-4) of causality analysis still exhibit maximum 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 4.1. 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.

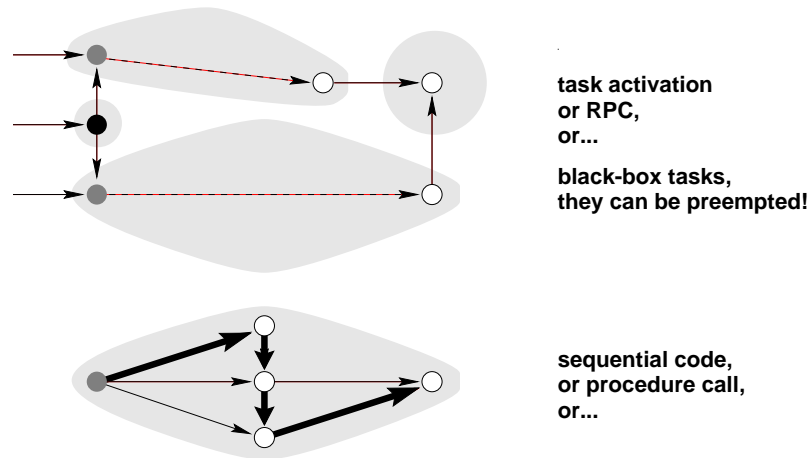
We shall however see that this does not need to be the case, however. Instead, a careful use of the scheduling specifications of an STS will allow us to decompose it into modules that can be stored as object code for further reuse, whatever the actual environment and implementation architecture will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

The case of single-clocked STS. We first discuss the case of single-clocked STS, in which all 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. In fact, as all variables belonging to the same task depend on the same inputs, each task can be executed according to the following scheme: 1/ collect inputs, 2/ execute task. The actual way the task is executed is arbitrary, provided it conforms the scheduling specification.


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



The thick arrows inside the task depicted on the bottom 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 task 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.

The general case of multiple-clocked STS. The generalization is illustrated in the following picture :

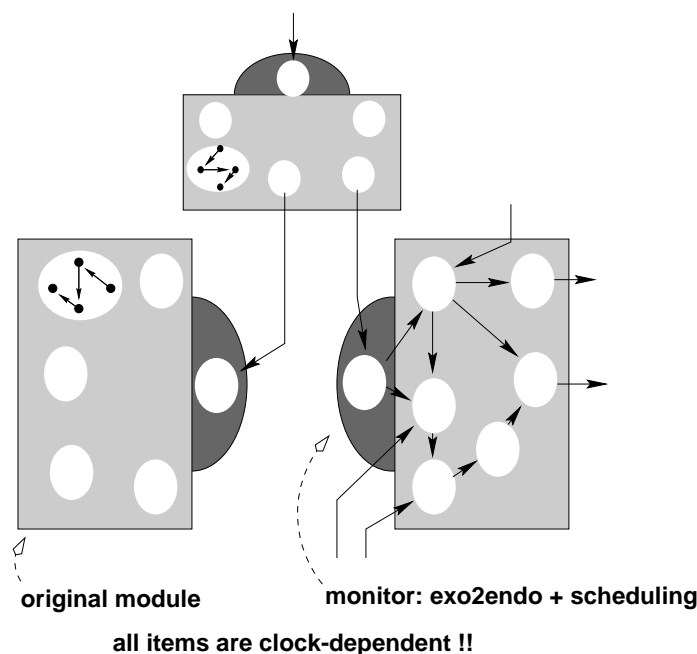


The only new point is that all items discussed before are now clock-dependent. Thus the computations of the “tasks”, their internal scheduling, and their abstraction, must be performed using scheduling specifications labelled by booleans. In the above example, the considered STS is activated by some clock h , which in turns defines the activation clock k of the depicted task. The bottom line is :

1. Different tasks can have different activation clocks.
2. For each task, the internal scheduling generally depends on some internal clocks, i.e., on some predicates involving internal states of the task. Thus the internal scheduling is dynamic, but can be precomputed at compile time.
3. Also, the task abstraction has a clock-dependent scheduling specification. This is another source of dynamic scheduling, also computed at compile time : it is implemented in the form of a “software monitor” which opens/forbids the possibility of executing a given task at a given instant, depending on the current status of the STS clocks.
4. This software monitor can be combined with the interruption mechanism of the supporting real-time kernel intended to optimize response time of the execution, without impairing conformity of the object code to its specification.

5.3 The bottom line: modular and distributed code generation

The whole approach is summarized in the following diagram :



In this diagram, gray rectangles denote three modules P_1, P_2, P_3 of the source STS specification, hence given by $P = P_1 \wedge P_2 \wedge P_3$. We assume here that this partitioning has been given by the designer, based on functional and architectural considerations. Note that the idempotence of STS composition ($Q \wedge Q = Q$) allows us to duplicate source code at convenience while partitioning the specification.

Then, white bubbles inside the gray rectangles depict the structuration into tasks as discussed before.

Finally the black half-ellipses denote the monitors. Monitors are in charge of 1/ providing the additional control to make the considered module endochronous if asynchronous communication media are to be used, and 2/ specifying the scheduling of the abstract tasks.

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.

6 Conclusion

The above approach is supported by the DC_+ common format for synchronous languages [DC96]. The DC_+ format is one possible concrete implementation of our STS model, including scheduling specifications. It serves as a basis for the code generation suite in the Esprit

SACRES project. S. Machard and E. Rutten are currently working on generic code generation based on this work, with different real-time O.S. and architectures as targets.

ACKNOWLEDGEMENT: *The authors are indebted to Benoît Caillaud for careful reading of and remarks on a draft version of the manuscript.*

A Appendix : formal study of causality and scheduling specifications

A.1 Dependencies

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

$$? \prec \overbrace{\{\perp, F, T\}}^{\mathfrak{I}} \quad (7)$$

$$\perp < F < T \quad (8)$$

In these formulæ, symbols $?$ (resp. \mathfrak{I}) indicate “undefined” (resp. “defined”). “Undefined” should not be confused with “absent” (\perp): “absent” is a perfectly defined status, while “undefined” is intended to model that a variable has not been produced yet in the current reaction. Finally, non-boolean types are abstracted as the single distinguished element T . Thus, for booleans, the pair $\{F, T\}$ can be seen as a refinement of the symbol T , this is shown by the underbrackets. And $\{\perp, T\}$ is a refinement of \mathfrak{I} , this is shown by the overbrackets.

Definition of dependencies: 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, when $b = T$, that y gets defined while x is not.

Properties of dependencies. The following properties hold :

$$x \xrightarrow{b} y \wedge y \xrightarrow{c} z \Rightarrow x \xrightarrow{b \wedge c} z \quad (9)$$

$$x \xrightarrow{b} y \wedge x \xrightarrow{c} y \Rightarrow x \xrightarrow{b \vee c} y \quad (10)$$

In these equations, $b \wedge c$ and $b \vee c$ are respectively defined as the infimum (resp. supremum) w.r.t. relation “ $<$ ” defined in (8) when both values belong to the subdomain $\{\perp, T, F\}$. Otherwise, they are both equal by definition to $?$.

x	$?$	\perp	\top
b			
$?$	$?$	$?$	$?$
\perp			
F			
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 $\{F, T\}$.

A.2 Circuitfree schedulings

We are given a set x_1, \dots, x_n of variables with some of them being boolean. The boolean variables will also be denoted by b_1, b_2, \dots , and will be used as labels in dependencies. Then we are given 1/ a set of constraints on the boolean variables of the form $C(b_1, \dots, b_k)$ as restricted to the subdomain $\{\perp, T, F\}$ of defined values, and 2/ a set of dependencies defined on the set x_1, \dots, x_n . We call the whole a *scheduling* of x_1, \dots, x_n .

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, denote by \mathcal{S} the set of states satisfying these constraints. States in \mathcal{S} are written s, t, \dots and corresponding interpretations for the variables x_1, \dots, x_n are denoted by s_1, \dots, s_n for short instead of $s[x_1], \dots, s[x_n]$, and similarly for t . In \mathcal{S} , some values s_i are set to $?$, and the other ones are free, or alternatively constrained by boolean equations for boolean variables.

Lemma 1 *Dependency $x \xrightarrow{b} y$ is part of the considered scheduling iff, for variables (x, y, b) , none of the triples of values $(?, \cdot, \top)$, $(\cdot, \cdot, ?)$ belongs to \mathcal{S} , where ‘ \cdot ’ denotes an arbitrary value for x .*

Proof: The “only if” part follows by inspection of table 1. As for the “if” part, we note that boolean constraints of the form $C(b_1, \dots, b_k)$ do not involve the $?$ value, so that the condition that the considered triples do not belong to \mathcal{S} can only result from the conjunction of dependencies, and we see that $x \xrightarrow{b} y$ is implied by this conjunction, and can thus be added without changing \mathcal{S} . \diamond

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

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

in (7). A path $s(1), s(2), \dots, s(K)$ such that $s(k) \prec s(k+1)$ (resp. $s(k) \preceq s(k+1)$) is called *increasing* (resp. *nondecreasing*).

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

$$\begin{aligned} x_{i_1} \xrightarrow{b_1} x_{i_2} \xrightarrow{b_2} x_{i_3} \dots x_{i_p} \xrightarrow{b_p} x_{i_1} \\ \text{and} \\ (b_1 \wedge \dots \wedge b_p = \top) \end{aligned} \quad (11)$$

Lemma 2 (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 $(?, \dots, ?)$ to s .*

Proof: We first prove the “if” part. Assume (11) 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} . Restrict \mathcal{S} to those states for which

$$(b_1 \wedge \dots \wedge b_p = \top) \text{ or } (b_1 \wedge \dots \wedge b_p = ?) \text{ holds,}$$

call $\mathcal{S}_{(b_1 \wedge \dots \wedge b_p = \top)}$ this subset.

Then, on $\mathcal{S}_{(b_1 \wedge \dots \wedge b_p = \top)}$, 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 :

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

and thus the x_{i_j} 's are, either all = ?, or alternatively all defined (equal to one of the values \perp, \top, F). Thus there is no increasing path leading to any defined state belonging to $\mathcal{S}_{(b_1 \wedge \dots \wedge b_p = \top)}$. This proves the “if” part.

We now prove the “only if” part by contradiction. Assume $\exists s \in \mathcal{S}$ satisfying $\forall i : s_i \neq ?$, such that there is no increasing path linking $(?, \dots, ?)$ 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 . And denote by S the maximal (for set theoretic union), connected⁷ subset of \mathcal{S} , such that 1/ $(b_1 \wedge \dots \wedge b_p = \top)$, 2/ every $t \in S$ satisfies : $\forall i, t_i \neq ?$ and there is no increasing path linking $(?, \dots, ?)$ to t , and finally 3/ $S \ni s$. By assumption, S is not empty. By construction of S , there are some variables, say $x_{j_1}, x_{j_2}, \dots, x_{j_q}$, with $q > 1$, such that :

$$\begin{aligned} \text{on } \mathcal{S}_{(b_1 \wedge \dots \wedge b_p = \top)} \setminus S & : x_{j_1} = x_{j_2} = \dots = x_{j_q} = ? , \\ \text{on } S & : x_{j_1} = x_{j_2} = \dots = x_{j_q} = \mathbf{i} . \end{aligned} \quad (12)$$

Using lemma 1 we conclude that

$$\begin{aligned} x_{j_1} \xrightarrow{b} x_{j_2} \xrightarrow{b} x_{j_3} \dots x_{j_q} \xrightarrow{b} x_{j_1} \text{ holds, where} \\ b = (b_1 \wedge \dots \wedge b_p) , \text{ and } b = \top \text{ is possible in } \mathcal{S}. \end{aligned} \quad (13)$$

⁷with respect to our notion of neighbouring states.

Since $q > 1$, we have a circuit, and this finishes the proof. \diamond

The intuitive interpretation of this lemma is that, for an STS with a circuitfree scheduling, it is possible to sequentially compute all variables from the inputs with no deadlock. Each increasing path mentioned in Lemma 2 corresponds to one possible sequential execution.

A.3 Deriving dependencies as causality constraints

In this section, we infer dependencies as causality constraints from STS statements. 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 it to 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 $\{\perp, F, T\}$ to the single value \mathbf{i} (defined). This is where asymmetry between left- and right hand side of equations occurs in our abstraction technique.
- (P-3) Since we are interested in causality constraints, we only need to keep track of configurations for which y cannot be defined. Hence, when $y = \mathbf{i}$ results from applying (P-2), we weaken it to “ y unconstrained”, which is depicted in the tables by an empty box.

We now proceed on deriving the scheduling associated to each primitive statement.

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 table 1 and of the following tables (the first table expresses that it is possible, for a variable, to know that it is present in the considered reaction, without knowing its actual value yet) :

h_x	?	\perp	\top
x	?	\perp	?

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

h_x	?	\perp	\top
x	?	\perp	?

,

weakened as (using **P-3**) :

h_x	?	\perp	\top
x	?		

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

Lemma 5 *The following holds :*

$$y = f(u, v) \Rightarrow (u, v) \xrightarrow{h_y} y$$

Proof: by inspection of table 1 and of the following table for the abstracted relation derived from $y = f(u, v)$ using principles (**P-1,P-2,P-3**) (# denotes a prohibited value) :

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

, weakened as :

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

which is equivalent to the formulæ 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 yielding respectively x and h_x for $[\text{if } b \text{ then } x = u] \wedge [\text{if } b \text{ then } h_x = h_u]$:

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

,

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

which is equivalent to the formulæ of lemma 6 after weakening by removing the \perp 's. 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.

A.4 Correct programs

In this section, we prove the following refinement of Theorem 1 :

Theorem 2 (correct programs) *Let P be an STS satisfying the following conditions :*

1. *Each statement in P has its associated set of derived dependencies as derived from lemmas 3, 4, 5, 6 as part of P .*
2. *Resulting scheduling is circuitfree.*
3. *There is provably no multiple definition of a variable, meaning that, whenever*

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

is part of P , then

$$b_1 \wedge b_2 \wedge \dots \wedge b_p = \text{T}$$

has to be provably impossible.

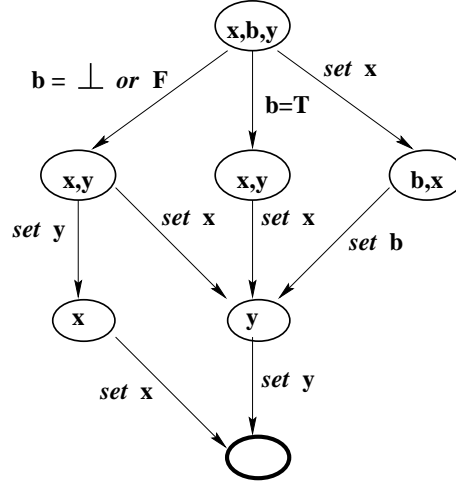
Then :

1. *As far as control is concerned, the inputs of 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 such as $x = \text{expression}$.*
3. *For each given input control history of P and compatible input value history, there is exactly one run of P , i.e., P is deterministic.*

NOTA: Clearly, theorem 2 provides us with a sufficient condition, this condition is not necessary. Furthermore, the rules for inferring dependencies 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 2 for a given syntactic form of P , it may satisfy them after proper rewriting into a syntactically different, but semantically equivalent form. Here, semantically equivalent means that the two programs in consideration have 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 with 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 with the considered primitive.
3. Then we take the product of the execution automata associated with each statement. By Lemma 2 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. Thanks to condition 3 of theorem 2, 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 :

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

The causality constraint and associated execution automaton are :

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



The refined execution automaton is obtained by replacing *set u* and *set v* by *read u* and *read v*, and *set y* by $y := f(u, v)$.

B Appendix : formal study of desynchronization (co-authored with Benoît Caillaud⁸)

Throughout this appendix, we only consider executable STS denoted by $\Phi = \langle V, \Theta, \rho \rangle$. For each considered STS, we assume that it has been equipped with scheduling specifications associated with each of its statements, following rules (R-1,2,3,4).

B.1 Desynchronizing STS: from STS to infinite partial orders, and their asynchronous composition

From the definition of scheduling specifications in section 4.2, we know that ρ specifies in particular a finite family of partial orders on V , consisting of a collection of circuitfree directed graphs, with elements of V as vertices, parameterized by the possible configurations of the set of boolean variables and clocks of Φ . Denote by \mathcal{G}^Φ this finite family of partial orders. Thus each $G \in \mathcal{G}^\Phi$ is a circuitfree directed graph having a subset of V as its vertices; by convention, each variable v which is not a vertex of G is added to G as a singleton, with no branch linked to it. This convention will hold in the sequel.

Each run $\sigma : s_0, s_1, s_2, \dots$ (cf. (1)) of Φ defines in particular a sequence $G_0^\Phi, G_1^\Phi, G_2^\Phi, \dots$ of elements of \mathcal{G}^Φ , where G_k^Φ is the partial order part of transition relation ρ at state s_k . Now, we turn this sequence $G_0^\Phi, G_1^\Phi, G_2^\Phi, \dots$ into a single infinite directed graph

$$G = G(\sigma) = G_0^\Phi \circ G_1^\Phi \circ G_2^\Phi \circ \dots, \quad (14)$$

where “ \circ ” is a *concatenation* operation which we define next :

1. For each variable $v \in V$, the occurrences of v in $G_0^\Phi, G_1^\Phi, G_2^\Phi, \dots$ are totally ordered as a chain $v_0 \prec v_1 \prec v_2 \prec \dots$.
2. Consider the disjoint union $\uplus_k G_k^\Phi$ of the graphs $G_0^\Phi, G_1^\Phi, G_2^\Phi, \dots$. Then concatenation $G_0^\Phi \circ G_1^\Phi \circ G_2^\Phi \circ \dots$ is defined as the minimal partial order greater than both $\uplus_k G_k^\Phi$ and the chains $v_0 \prec v_1 \prec v_2 \prec \dots$ for each $v \in V$.

Next, we label each vertex of $G_0^\Phi \circ G_1^\Phi \circ G_2^\Phi \circ \dots$ with the pair {variable, carried value} attached with this vertex in the considered run. This makes run $\sigma : s_0, s_1, s_2, \dots$ an infinite *labelled partial order*, which we denote by

$$\sigma = s_0 \circ s_1 \circ s_2 \circ \dots$$

Finally, we define

$$\sigma^a \stackrel{\text{def}}{=} \sigma / \text{absent variables}, \quad (15)$$

(where symbol “/” denotes hiding) meaning that 1/ we remove, in partial order $G_0^\Phi \circ G_1^\Phi \circ G_2^\Phi \circ \dots$, those vertices labelled with variables carrying \perp as a value in run σ , and then 2/ take the resulting labelled partial order and call it σ^a , where superscript “ $.^a$ ” refers to “asynchronous”.

⁸IRISA/INRIA, name@irisa.fr

Informally speaking, successive scheduling graphs are concatenated on a variable-by-variable basis, then absent variables are deleted while transitive closure is taken, and finally we take transitive reduction of the so obtained partial order to get its associated directed graph. This procedure is illustrated in the diagram of the pseudo-statement “**if** b **then** get_u ” in section 5.1. The key point is that, in σ^a , all vertices correspond to present occurrences of variables. Thus no test for absence can be performed. Infinite graph σ^a is our model of an *asynchronous run* of STS Φ .

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

$$\Phi^a =_{\text{def}} \langle V, \Sigma^a \rangle, \quad (16)$$

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

$$\Phi_1^a \bigwedge^a \Phi_2^a =_{\text{def}} \langle V, \Sigma^a \rangle, \text{ where } \begin{cases} V &= V_1 \cup V_2 \\ \Sigma^a &= \Sigma_1^a \wedge^a \Sigma_2^a \end{cases} \quad (17)$$

and \wedge^a denotes 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 *compatible*, written

$$\sigma_1^a \bowtie^a \sigma_2^a, \quad (18)$$

if the following conditions hold :

C₁: $\forall v \in V_1 \cap V_2$, labelled partial orders σ_1^a and σ_2^a have identical restrictions to the chain $v_0 \prec v_1 \prec v_2 \prec \dots$, meaning that, for v a shared variable, the values carried by v_0, v_1, v_2, \dots are identical for σ_1^a and σ_2^a . If this condition holds, then, for v a shared variable, we can unify chains $v_0 \prec v_1 \prec v_2 \prec \dots$ in labelled partial orders σ_1^a and σ_2^a . We assume this holds in discussing condition **C₂** to follow.

C₂: assuming condition **C₁** holds, then we require in addition that the least upper bound of labelled partial orders σ_1^a and σ_2^a exists. This amounts to assuming that there is no pair (x, y) of shared vertices of σ_1^a and σ_2^a such that $x \prec_1 y$ holds in σ_1^a , while $y \prec_2 x$ holds in σ_2^a . Intuitively speaking, no short circuit (i.e., no “deadlock”) is created in combining the two considered asynchronous runs.

If conditions (**C₁**, **C₂**) hold, then we define $\sigma_1^a \wedge^a \sigma_2^a$ as being the above upperbound, and $\Sigma_1^a \wedge^a \Sigma_2^a$ as being the set of $\sigma_1^a \wedge^a \sigma_2^a$, where (σ_1^a, σ_2^a) ranges over the set of pairs of compatible labelled partial orders. This formally defines (17).

Synchrony vs. Asynchrony ? At this point a natural question arises, namely, does the following property hold :

$$\Phi_1^a \bigwedge^a \Phi_2^a = \left(\Phi_1 \bigwedge \Phi_2 \right)^a, \quad (19)$$

also illustrated by the following commutative diagram :

$$\begin{array}{ccc} (\Phi_1, \Phi_2) & \xrightarrow{a} & (\Phi_1^a, \Phi_2^a) \\ \downarrow & & \downarrow \\ \Phi_1 \wedge \Phi_2 & \xrightarrow{a} & \Phi_1^a \wedge^a \Phi_2^a \end{array}$$

Property (19) intuitively expresses that communication behaves equivalently for both synchronous and asynchronous composition. Roughly speaking, if (19) was true, then we could interpret our STS model equivalently as synchronous or asynchronous.

Unfortunately, property (19) does not hold in general, due to the possibility to exercise control by the way of absence in synchronous composition \wedge . In the following section, we show that (19) indeed holds under certain sufficient conditions, in which the notion of *endochrony* plays a central role.

B.2 Endochrony

In this section, we use notations from section 3. For each STS $\Phi = \langle V, \Theta, \rho \rangle$, we denote by

$$\Phi^h \tag{20}$$

the STS obtained by abstracting variables into their clocks. Also, for $W \subseteq V$, we denote by Φ_W the restriction of STS Φ to W , and by Φ_W^h the abstraction of Φ_W into its clocks.

For $W' \subseteq W \subseteq V$ we say that W' is a *clock inference* of W , written

$$W' \hookrightarrow W, \tag{21}$$

if the following diagram defines a surjective function associating, for each run of $\Phi_{W'}$, a unique run of Φ_W^h :

$$\begin{array}{ccc} \Phi_W & \xrightarrow{\text{proj}_{W'}} & \Phi_{W'} \\ \text{clock abstraction } \downarrow & & \swarrow \\ & & \Phi_W^h \end{array}$$

This means that presence/absence of variables belonging to $W \setminus W'$ in Φ_W can be inferred from the values of variables in $\Phi_{W'}$. If $W' \hookrightarrow W_1$ and $W' \hookrightarrow W_2$ hold, then $W' \hookrightarrow (W_1 \cup W_2)$ follows, thus there exists a greatest W such that $W' \hookrightarrow W$ holds. Hence we can consider the unique increasing chain

$$\emptyset = V(0) \hookrightarrow V(1) \hookrightarrow V(2) \hookrightarrow \dots \tag{22}$$

of subsets of V such that, for each k , $V(k)$ is the greatest set of variables such that $V(k-1) \hookrightarrow V(k)$ holds. As $\emptyset = V(0)$, $V(1)$ consists of the subset of variables that are present as soon as the considered STS gets activated. Of course chain (22) must become stationary at some finite k_o : $V(k_o + 1) = V(k_o)$. In general, we only know that $V(k_o) \subseteq V$.

Definition 1 (endochrony) STS Φ is said to be endochronous if $V(k_o) = V$, i.e., if the following condition is satisfied:

$$\text{the chain } \emptyset = V(0) \hookrightarrow V(1) \hookrightarrow V(2) \hookrightarrow \dots \text{ converges to } V. \quad (23)$$

Condition (23) 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 3 Consider an STS $\Phi = \langle V, \Theta, \rho \rangle$. Then, condition (a) implies condition (b):

- (a) Φ is endochronous.
- (b) For each $\sigma^a \in \Sigma^a$, we can reconstruct the corresponding synchronous run σ for Φ , in a unique way up to silent reactions.

Proof: we prove the result by induction. Pick a $\sigma^a \in \Sigma^a$, and assume for the moment that we were able to decompose it via concatenation as:

$$\underbrace{s_1 \circ s_2 \circ \dots \circ s_n}_{n\text{-initial segment of } \sigma} \circ \sigma_n^a \quad (24)$$

into a finite n -length concatenation of non-silent reactions s_i , followed by the tail of the infinite graph σ^a , which we denote by σ_n^a , and we assume that such a decomposition is unique. Then we claim that

$$(24) \text{ is also valid with } n \text{ substituted by } n + 1. \quad (25)$$

To prove (25), we note that, when STS Φ gets activated, then we know that variables belonging to $V(1)$ will be present in the considered reaction. By assumption, there is a unique $s_{n+1}^h(1)$ having $V(1)$ as variables. Thus, present vertices and branches of graph $s_{n+1}(1)$ are 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 chain named v in σ_n^a . 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 reaction. 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 (25).

Given the initial condition for σ^a , we get from (25), by induction, the desired proof that (a) \Rightarrow (b). \diamond

The next result addresses the question of when property (19) holds true.

Theorem 4 Consider two STS $\Phi_i = \langle V_i, \Theta_i, \rho_i \rangle, i = 1, 2$. Denote by $V = V_1 \cap V_2$ the set of their common variables, and by $\Phi = \Phi_1 \wedge \Phi_2$ their synchronous composition. Consider the following conditions :

(A₁) Φ_1, Φ_2 , and Φ are endochronous.

(A₂) Denote by $\Phi_{i,V}^h$ and Φ_V^h the restriction of Φ_i^h and Φ^h to V , respectively. Then $\Phi_{1,V}^h = \Phi_{2,V}^h = \Phi_V^h$ holds.

Condition (A₂) expresses that no clock constraint results from composing the two considered STS, i.e., their composition only results in constraints involving present values of variables.

If conditions (A₁) and (A₂) are satisfied, then property (19) holds true.

COMMENTS : We already discussed the importance of guaranteeing property (19). Now, why is this theorem interesting? Mainly because it replaces condition (19), which involves infinite runs, by conditions (A₁) and (A₂), which only involve a single reaction. While conditions (A₁) and (A₂) are generally undecidable, some decidable approximation of them (by abstracting nonboolean domains) is actually checked in the SIGNAL compiler, see section B.3.

Proof: We proceed into several steps.

1. Denote by Φ^a the desynchronisation of Φ , defined by (16), and denote by σ^a a run of Φ^a . For each $\sigma^a \in \Sigma^a$, there is at least one corresponding synchronous run σ for Φ . Any such σ is clearly the synchronous composition of two compatible 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 compatible, and their asynchronous composition $\sigma_1^a \wedge^a \sigma_2^a$ belongs to $\Sigma_1^a \wedge^a \Sigma_2^a$. Thus we have the inclusion

$$\Phi_1^a \wedge^a \Phi_2^a \supseteq \left(\Phi_1 \wedge \Phi_2 \right)^a, \quad (26)$$

which proves the first part of (19). So far we have only used the definition of desynchronisation and asynchronous composition, none of the conditions (A₁), (A₂) was requested.

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 compatible runs will not result in rejecting pairs of runs that otherwise would be compatible in the asynchronous sense. This is where conditions (A₁), (A₂) enter the game.

Let $V_1(k)$, $V_2(k)$, and $V(k)$, denote the chains (22) associated with Φ_1 , Φ_2 , and $\Phi_1 \wedge \Phi_2 = \Phi$ respectively. Consider the set $V(1)$ of variables that are present as soon as Φ gets activated. Due to condition (A₂), we know that each such variable

must belong to at least one of the sets $V_i(1)$ for $i = 1, 2$. Hence one of the following cases must occur:

$$\begin{aligned} \text{case 1} & : V(1) = V_1(1) \\ \text{case 2} & : V(1) = V_2(1) \\ \text{case 3} & : V(1) = V_1(1) \cup V_2(1) \end{aligned}$$

Clearly, cases 1 and 2 are equivalent up to exchanging the names of the two considered STS. Thus we assume either case 1 or case 3.

Case 3 is the easiest one. In this case, again using condition (\mathbf{A}_2) , we get that $V(k) = V_1(k) \cup V_2(k)$ holds. Thus focus on case 1, and consider next $V(2)$. Two subcases occur: 1/ $V(2)$ contains the activation clock of Φ_2 , and then $V(2) \supseteq V_2(1)$ and $V(k) = V_1(k) \cup V_2(k-1)$ follows; or alternatively 2/ $V(2)$ does not contain the activation clock of Φ_2 , and then we continue our argument. Arguing in the same way for $V(3)$, etc, we get that $\exists k_1 \geq 0$ such that:

$$\begin{aligned} \text{for } 0 \leq k \leq k_1 & : V_1(k) = V(k) \cap V_1, \\ \text{whereas for } k > k_1 & : V(k) = V_1(k) \cup V_2(k - k_1). \end{aligned} \quad (27)$$

In particular, the activation of Φ_2 is determined by Φ_1 .

Hence, pick a pair (σ_1^a, σ_2^a) such that $\sigma_1^a \bowtie^a \sigma_2^a$ (cf. (18)): they can be combined while performing the asynchronous composition $\Phi_1^a \wedge^a \Phi_2^a$ to form some σ^a (cf. (17)), this is denoted by $\sigma_1^a \wedge^a \sigma_2^a = \sigma^a$. Then:

- (a) Apply scheme (b) of Theorem 3 to σ_1^a , this yields a unique decomposition of σ_1^a into a sequence of reactions

$$\sigma_1^a = s_{1,1} \circ s_{1,2} \circ \dots \circ s_{1,n} \circ \dots$$

- (b) Using (27), we know the subsequence n_k of reactions of Φ_1 at which σ_{1,n_k} activates Φ_2 .

- (c) Apply scheme (b) of Theorem 3 to σ_2^a , this yields a unique decomposition of σ_2^a into a sequence of reactions

$$\sigma_2^a = s_{2,1} \circ s_{2,2} \circ \dots \circ s_{2,k} \circ \dots \quad (28)$$

- (d) Relabel sequence $1, 2, \dots, k, \dots$ in (28) as $n_1, n_2, \dots, n_k, \dots$ (defined at step 2b) and fill the holes in so relabelled decomposition (28) with silent reactions⁹ for Φ_2 , denoted by \emptyset_2 , we get another decomposition of σ_2^a :

$$\begin{aligned} \sigma_2^a & = \emptyset_2 \circ \dots \circ \emptyset_2 \circ s_{2,n_1} \circ \emptyset_2 \circ \dots \circ \emptyset_2 \circ s_{2,n_2} \circ \dots \\ & \quad \dots \circ s_{2,n_k} \circ \dots \\ & \stackrel{\text{def}}{=} s_{2,1} \circ s_{2,2} \circ \dots \circ s_{2,n} \circ \dots \end{aligned}$$

⁹Cf. formula (3) and neighbouring explanations for the definition of a silent reaction.

- (e) Using again condition (\mathbf{A}_2) , we can assert that, for each n , $s_{1,n} \bowtie s_{2,n}$, where \bowtie denotes compatibility of considered states as needed for synchronous composition. Finally, we get

$$\begin{aligned} \sigma_1^a \wedge^a \sigma_2^a &= (s_{1,1} \circ s_{1,2} \circ \dots)^a \wedge^a (s_{2,1} \circ s_{2,2} \circ \dots)^a \\ &= (s_1 \circ s_2 \circ \dots)^a \\ &= \sigma^a \end{aligned}$$

for some run σ of STS Φ , where s_n is the state obtained by composing the two compatible states $s_{1,n}$ and $s_{2,n}$. This proves the inclusion symmetric to (26) and thus proves (19).

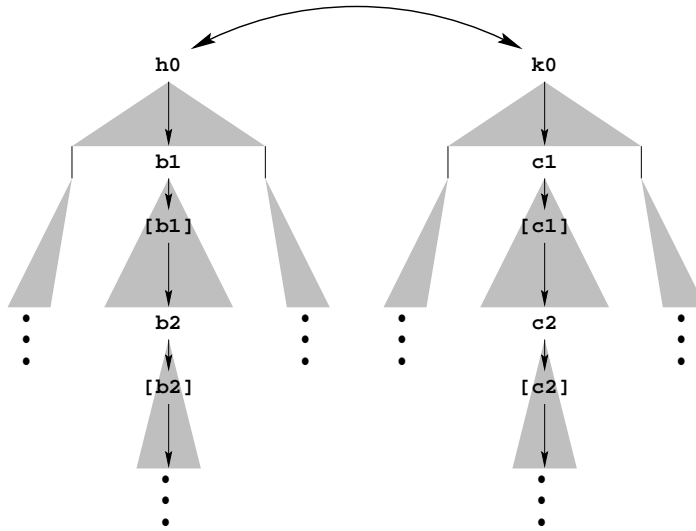
This finishes the proof of the theorem. \diamond

B.3 Handling endochrony in practice

While we have given a finitary criterion for endochrony, we did not propose a practical algorithm for checking this criterion. We do this now. In this subsection, we shall indicate 1/ how a (tight) sufficient condition for endochrony can be actually tested, and 2/ how making an STS endochronous can be performed. As both the DC_+ format and the `SIGNAL` 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.

B.3.1 Checking endochrony

As one of the modules of the existing DC_+ or `SIGNAL` compiler, the following structure is computed, for a given program P :



In this picture, b, c denote boolean variables, $[b], [c]$ denote corresponding clocks composed of the instants at which $b, c = \top$ holds, respectively. Finally, h, k are also clocks (i.e., variables of type `event`). 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. In doing so, a tree is built under each of the clocks h_0, k_0, \dots , this yields the so-called *clock hierarchy* in the form of a “forest”, i.e., a collection of trees. Roots of the trees are related by some clock equation, this is depicted as the bidirectional arrow relating h_0, k_0, \dots . This structure is detailed in [ABIG94] [ABIG95], where it is shown to be a canonical form for representing the combination of clock equations and scheduling specifications of a program. Now, considering this clock hierarchy, one easily proves the following :

Theorem 5 *If program P has a clock hierarchy consisting of a single tree, then it is endochronous. Actually, the chain $\emptyset = V(0) \leftrightarrow V(1) \leftrightarrow V(2) \leftrightarrow \dots$ consists of the successive sets of variables that have as their clock the subsets of clocks figuring at levels $\leq 0, \leq 1, \leq 2, \dots$, respectively.*

Theorem 5 is an immediate corollary of Theorem 3 of Appendix B. It does not state that P is endochronous if and only if its clock hierarchy is a tree, but nearly so. In fact, such a result would be true for STS restricted to types `{event, boolean}`. For general STS, abstractions performed while computing the clock hierarchy prevents from obtaining a necessary and sufficient condition for endochrony — actually, property (23) of definition 1 is undecidable for general programs. 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.

Nevertheless, in practice, we shall use the clock hierarchy as the practical criterion for checking endochrony.

B.3.2 Enforcing endochrony

Assume we have an STS P which is not endochronous, and we *want* it to be so. What can we do? As revealed by inspecting the previous picture, if we can make the roots h_0, k_0, \dots of the clock hierarchy belonging to some single clock tree, then we are done. In other words, we can concentrate on the roots of the clock hierarchy. Thus the situation for study is the following.

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

$$\begin{aligned} p_1(h_1, \dots, h_k) &\neq \text{F} \\ &\dots \\ p_q(h_1, \dots, h_k) &\neq \text{F} \end{aligned} \tag{29}$$

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 (29) can be solved for their variables h_1, \dots, h_k , meaning that we can find a set h_1^o, \dots, h_l^o of clocks, and a set p_1^o, \dots, p_k^o of clock expressions, such that

$$\begin{aligned} h_1 &= p_1^o(h_1^o, \dots, h_l^o) \\ &\dots \\ h_k &= p_k^o(h_1^o, \dots, h_l^o) \end{aligned} \tag{30}$$

has the same set of solutions for h_1, \dots, h_k as the original system (29), and new clocks h_1^o, \dots, h_l^o are free, i.e., unconstrained by the system of equations (30). Finally, we introduce boolean variables b_1^o, \dots, b_l^o , and a “master clock” h^o , such that

$$\begin{aligned} h_1^o &= [b_1^o], \dots, h_l^o = [b_l^o] \\ h_{b_1^o} &= \dots = h_{b_l^o} = h \end{aligned} \tag{31}$$

The bottom line is :

1. System of clock equations (29) is equivalent to (30,31) after hiding auxiliary variables h, b_1^o, \dots, b_l^o .
2. System (30,31) is endochronous.
3. Further combining (30,31) with another STS processed in the same way, will result in the satisfaction of additional condition \mathbf{A}_2 of theorem 4.

Discussion. Basically, building (31,30) from (29) intuitively corresponds to equipping the original P program with a suitable *communication protocol* Q in such a way that the compound program $P \wedge Q$ becomes endochronous. This is not a surprise, 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.

Also, the way we moved from (29) to (30) reveals one unpleasant feature of this technique, namely: this part of the process is not unique, and thus there are actually many possible different protocols that would work. For simple cases (e.g., root clocks h_1, \dots, h_k are actually themselves free), some obvious choice prevails, and corresponding solutions are used to generate (endochronous) code by the DC_+ compiler. For general cases, there is no obvious choice and it is expected that other considerations would enter the game¹⁰, thus no code is currently generated for corresponding cases by the DC_+ compiler.

References

- [AR88] I.J. AABELSBERG, AND G. ROZENBERG, *Theory of traces*, Theoretical Computer Science, 60, 1–82, 1988.
- [ABIG94] T.P. AMAGBEGNON, L. BESNARD AND P. LE GUERNIC, “Arborescent canonical form of boolean expressions”, Inria Research Report n°2290, June 1994.
- [ABIG95] T.P. AMAGBEGNON, L. BESNARD AND P. LE GUERNIC, “Implementation of the dataflow language SIGNAL”, in *Programming Languages Design and Implementation*, ACM, 163–173, 1995.
- [Aub97] P. AUBRY, “Mises en œuvre distribuées de programmes synchrones”, PhD Thesis, Univ. Rennes I, 1997.
- [BB91] A. BENVENISTE, 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.
- [BIGJ91] 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.
- [BCHIG94] 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.
- [Ber95] GÉRARD BERRY, *The Constructive Semantics of Esterel*, Draft book, <http://www.inria.fr/meije/esterel>, December 1995.

¹⁰such as: what is the communication layer in the supporting distributed target architecture ?

- [CCGJ97] 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.
- [CL87] M. CLERBOUT, AND M. LATTEUX, *Semi-commutations*, Information and Computation, 73, 59–74, 1987.
- [DC96] SACRES CONSORTIUM, The Declarative Code DC+, Version 1.2, May 1996; Esprit project EP 20897: Sacres.
- [Halb93] N. HALBWACHS, *Synchronous programming of reactive systems*,. Kluwer Academic Pub., 1993.
- [KP96] Y. KESTEN AND A. PNUELI, An α STS-based common semantics for SIGNAL and STATECHARTS, March 1996. Sacres Manuscript.
- [Lam83a] L. LAMPORT, Specifying concurrent program modules, *ACM Trans. on Prog. Lang. and Sys.*, 5(2):190-222, 1983.
- [Lam83b] L. LAMPORT, What good is temporal logic? In *Proc. IFIP 9th World Congress*, R.E.A. Mason (Ed.), North Holland, 657-668, 1983.
- [LG91] 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.
- [MLG94] 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.
- [MP92] Z. MANNA AND A. PNUELI, *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, New York, 1992.
- [MP95] Z. MANNA AND A. PNUELI, *The Temporal Logic of Reactive and Concurrent Systems: Safety*. Springer-Verlag, New York, 1995.
- [Pnu97] A. PNUELI, personal communication.



Unit e de recherche INRIA Lorraine, Technop le de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS L ES NANCY
Unit e de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unit e de recherche INRIA Rh ne-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unit e de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unit e de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

 diteur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399