



HAL
open science

Factorization of Unfoldings for Distributed Tile Systems Part 2: General Case

Eric Fabre

► **To cite this version:**

Eric Fabre. Factorization of Unfoldings for Distributed Tile Systems Part 2: General Case. [Research Report] RR-5186, INRIA. 2004. inria-00071402

HAL Id: inria-00071402

<https://inria.hal.science/inria-00071402>

Submitted on 23 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.

***Factorization of Unfoldings
for Distributed Tile Systems
Part 2 : General Case***

Eric Fabre

N°5186

May 2004

————— Systèmes communicants —————



***R**apport
de recherche*

Factorization of Unfoldings for Distributed Tile Systems Part 2: General Case

Eric Fabre

Systemes communicants
Projet DistribCom

Rapport de recherche n° 5186 — May 2004 — 37 pages

Abstract: We consider large distributed discrete event systems, *i.e.* systems obtained by connecting a possibly large number of elementary components. By “large” we mean that the size of the system prevents its study at a global level. We propose instead a framework to analyse such systems by parts, taking advantage of their modular definition. The key is twofold. We first represent their runs in a true concurrency semantic, which already reduces the state explosion phenomenon: trajectories which only differ by the ordering of concurrent events are not distinguished. The convenient data structures to represent sets of trajectories in this semantic are called branching processes of the system, and their maximal representative corresponds to the unfolding of the system. The second idea lies in the following result: the unfolding of a modular system factorizes as the product of unfoldings of its components, which gives an even more compact representation of runs of a distributed system. Therefore one should rather perform all computations on this factorized form. We propose an algebraic setting to do so: computations take place in the category of “augmented branching processes,” and rely on two operations: projection and product. We illustrate this approach on a simple problem: find the minimal factorization for the unfolding of the global system, which amounts to determining, for each component, runs that remain possible once this component is inserted in the global system.

Key-words: discrete event system, distributed system, concurrency, unfolding, branching process, factorization, turbo algorithm, graphical model of interactions

(Résumé : *tsvp*)

This work was supported by RNRT projects MAGDA and MAGDA 2, funded by the French Ministry of Research.

Factorisation des dépliages pour des systèmes de pièces

Partie 2 : cas général

Résumé : On s'intéresse à de grands systèmes à événements discrets, construits en connectant un nombre élevé de composants. Par "grand", on signifie que l'analyse de ces systèmes ne peut se faire à un niveau global, à cause d'un phénomène d'explosion combinatoire. On propose une alternative, consistant à étudier ces systèmes par morceaux, en profitant justement de leur nature modulaire. Cela repose sur deux idées. En premier lieu, on suppose une sémantique de concurrence vraie sur les trajectoires de ces systèmes. Cela signifie que des trajectoires ne différant que par l'entrelacement d'événements concurrents ne sont pas distinguées, ce qui permet déjà de réduire l'explosion combinatoire. Dans cette sémantique, les ensembles de trajectoires se décrivent par des processus de branchements, le plus grand d'entre eux correspondant au dépliage du système. Le second résultat clef est le suivant : le dépliage d'un système distribué se factorise en un produit de dépliages de ses composants. Cette forme factorisée permet de réduire encore la description des trajectoires d'un système distribué. Nous proposons justement de fonder l'analyse d'un système distribué sur cette forme factorisée de son dépliage. Les objets manipulés sont des processus de branchements augmentés des composants du système, et les opérations de base le produit et la projection. Nous donnons des algorithmes pour mener à bien les calculs sur cette forme factorisée, algorithmes inspirés des méthodes "turbo" en communications numériques. Cette technique d'analyse modulaire d'un système distribué est illustrée sur un système simple : le calcul des facteurs minimaux du dépliage du système global. Cela revient à calculer, pour chaque composant, les comportements qui demeurent possibles une fois que ce composant est inséré dans le système global.

Mots-clé : système à événements discrets, système réparti, concurrence, dépliage, processus de branchement, factorisation, algorithme turbo, graphe d'interaction

Contents

1	Introduction	4
2	Augmented branching processes (ABP)	5
2.1	Tile systems, occurrence nets and branching processes	5
2.2	Augmentation of occurrence nets and of branching processes	7
3	Properties and basic operations on ABPs	10
3.1	Basic properties	10
3.2	Intersection	12
3.3	Union	13
3.4	Trimming	14
4	Product of ABPs	15
4.1	Definition	16
4.2	Categorical product	17
4.3	Relations to trimming	18
4.4	Other properties	19
5	Projection of ABPs	20
5.1	Standard projection of BP : its weaknesses	20
5.2	Projection of ABPs	21
6	ABP calculus on trees	22
6.1	A key property	22
6.2	Modular computations on tree-shaped systems	25
7	Equivalence of ABPs and involutivity	27
7.1	Sub-involutivity of GABPs	27
7.2	Pre-order on GABPs	27
7.3	Equivalence of GABPs	28
7.4	Product, projection and equivalence	29
7.5	Involutivity of minimal ABPs	31
7.6	Minimal product covering of an ABP	31
8	Local computations on systems with cycles	32
8.1	Summary of results obtained on trees	32
8.2	Weak convergence for cyclic systems	33
9	Discussion and conclusion	34

1 Introduction

Distributed and concurrent systems have been studied for a long time in computer science, with emphasis on problems like election, consensus, mutual exclusion, protocol design and verification, etc. [18, 19]. Relying on this sound basis, the powerful concept of distributed computing has deeply penetrated the field of computer science. Applications range from telecommunication and computer networks on one end, to distributed databases, and more generally distributed programming on the other end. In conjunction with object oriented programming (and remote object invocation routines), distributed processing is now becoming a design paradigm, as in web-services for example. This approach offers very appealing aspects, such as an easy access to distributed resources/information, the reusability of software, some flexibility in the choice of components, and most importantly modularity, which allows the design of large distributed systems from elementary components. As a counterpart of this powerful trend, some actual systems have now reached a critical level in terms of size and heterogeneity, which makes them hardly tractable by standard analysis methods operating at a global level. The purpose of this paper, extending results of [1], is to propose a framework for analyzing/monitoring large distributed systems. Specifically, we propose to take advantage of the modular nature of a system to design modular analysis/monitoring procedures.

What we call a *distributed* system in this paper is first of all a *modular* system, *i.e.* a system made of several interconnected components. Each component is a dynamic system, interacting with neighboring components, and having some degree of autonomy in its dynamics. In particular, there is no global clock giving the pace to all components, which justifies the term “distributed” by contrast with “modular.” To make this assumption more concrete, we adopt a true concurrency semantics on runs (or trajectories) of a distributed system: runs are described as event structures, or partial orders of events, rather than sequences of events. This encodes the fact that two components can run concurrently (*i.e.* in parallel) as far as they have no interaction. In this semantic, a set of runs of a system is represented under the form of a *branching process*, a compact data structure, and the set of all possible runs corresponds to the maximal branching process, named the *unfolding* of that system [5, 6, 7].

In the first part of this work, we proved that the unfolding of a distributed system can be expressed as the product of unfoldings of its components, which gives an even more compact and modular (thus scalable) representation of a possibly huge set of runs. (After developing this work, we came aware of early results of Winskel deriving exactly this property, from simple standard arguments in category theory, see [3], p.258.) More generally, any branching process of the global system can be embedded into a minimal product of branching processes of its components. Each factor of this product covering contains runs of a component that are part of at least one run of the global system. We believe this product representation is the key to efficient modular processings for distributed systems. As a typical example of such modular processings, we address the problem of computing the minimal product covering of the unfolding of a distributed system. This amounts to determining local runs of each component that remain possible once this component is inserted into the global system. It turns out that this problem can be solved in an efficient manner by an appropriate combination of local computations on branching processes, based on two operations: product and projection. While the product of branching processes, or of event structures, is actually a quite standard operation [3, 14, 13, 15], the projection we define seems less known: we use it to abstract the behaviors of a given component on its interface with another component. The originality of our contribution probably lies in the emphasis we put on algebraic relations between product and projection, that allow the orchestration of modular computations. As a matter of fact, the computations we describe are based on a message passing algorithm between “computation modules” associated to each component. Since the exchange of messages is asynchronous, our modular computation schemes can actually be easily *distributed*. An application of this framework to distributed failure diagnosis, with applications to telecommunication networks, can be found in [24, 25, 26].

The product representation of the unfolding of a distributed system, and the orchestration of local computations based on product and projection, were already presented in [1]. This previous contribution required that all objects appearing in the course of computations were branching processes of restrictions of the global system. So computations remained valid as long as projections in this category were not misleading, *i.e.* did not unduly transform causality or conflict relations between events into an apparent concurrency. In this paper, we remove this strong limitation by embedding computations into the larger and new category of *augmented branching processes* (ABP). The latter allow to define a projection operation that preserves all necessary causality and conflict relations. ABPs are defined in section 2, while section 3 studies their properties and defines elementary operations on them, like intersection and union and the new trimming operation. The product of ABPs is defined in section 4, and extends the usual product of branching processes. Unfortunately, it is not stable in the category of ABPs and requires the introduction of a larger category: generalized ABPs. Projection is defined in section 5, and used jointly with the product (section 6) to define an “ABP calculus” for distributed systems where component interactions have a tree structure. The extension of this ABP calculus to systems with arbitrary interaction structure is done in section 8, and takes the form of an iterative algorithm. This algorithm doesn’t yield the exact minimal product covering for the unfolding of a distributed system, but only an approximation of it, that nevertheless has interesting properties. The convergence of this iterative procedure relies on an involutivity property of ABPs, studied in section 7.

2 Augmented branching processes (ABP)

2.1 Tile systems, occurrence nets and branching processes

This subsection essentially recalls notations and notions introduced in [1] (we refer the reader to that paper for more details).

Tile systems. A *tile system* \mathcal{S} is a tuple $(\mathcal{V}, \mathcal{T}, \mathbf{v}^0, \alpha, \beta, \gamma)$ where \mathcal{V} and \mathcal{T} are finite sets of variables and tiles respectively, and \mathbf{v}^0 is the initial state of the system (α, β, γ are defined below). Variables are denoted by capital letters A, B, V_1, V_2, \dots , and take values a, b, v_1, v_2, \dots in the finite domain \mathcal{D} . A *state* of \mathcal{S} is a function $\mathbf{v} : \mathcal{V} \rightarrow \mathcal{D}$. Tiles are generically denoted by $\mathbf{t}, \mathbf{t}', \mathbf{t}_1, \mathbf{t}_2$, etc. Function $\gamma : \mathcal{T} \rightarrow 2^{\mathcal{V}}$ associates to each tile the variable set on which it operates; the variable set $\gamma(\mathbf{t})$ impacted by tile \mathbf{t} is also denoted by $\mathcal{V}_{\mathbf{t}}$ for short. In the same way, there exist partial functions $\alpha, \beta : \mathcal{T} \times \mathcal{V} \rightarrow \mathcal{D}$ defining for each tile \mathbf{t} and each variable $V \in \mathcal{V}_{\mathbf{t}}$ the value of V respectively before and after \mathbf{t} fires. \mathbf{t} is enabled by state \mathbf{v} , denoted by $\mathbf{v}[\mathbf{t}]$, iff $\forall V \in \mathcal{V}_{\mathbf{t}}, \mathbf{v}(V) = \alpha(\mathbf{t}, V)$. \mathbf{t} can thus fire, which yields the new state \mathbf{v}' defined by $\forall V \in \mathcal{V}_{\mathbf{t}}, \mathbf{v}'(V) = \beta(\mathbf{t}, V)$, and $\forall V \in \mathcal{V} \setminus \mathcal{V}_{\mathbf{t}}, \mathbf{v}'(V) = \mathbf{v}(V)$. We adopt the standard notation $\mathbf{v}[\mathbf{t}]\mathbf{v}'$.

A *distributed* tile system is defined as the composition of several components. Consider tile systems $\mathcal{S}_i = (\mathcal{V}_i, \mathcal{T}_i, \mathbf{v}_i^0, \alpha_i, \beta_i, \gamma_i)$, $1 \leq i \leq N$. These systems are *coherent* iff, $\forall 1 \leq i, j \leq N, i \neq j$,

1. $\forall V \in \mathcal{V}_i \cap \mathcal{V}_j, \mathbf{v}_i^0(V) = \mathbf{v}_j^0(V)$
2. $\forall \mathbf{t} \in \mathcal{T}_i \cap \mathcal{T}_j, \gamma_i(\mathbf{t}) \cap \mathcal{V}_j = \gamma_j(\mathbf{t}) \cap \mathcal{V}_i$
3. $\forall \mathbf{t} \in \mathcal{T}_i \cap \mathcal{T}_j, \forall V \in \mathcal{V}_i \cap \mathcal{V}_j, \alpha_i(\mathbf{t}, V) = \alpha_j(\mathbf{t}, V)$ and $\beta_i(\mathbf{t}, V) = \beta_j(\mathbf{t}, V)$ on pairs (\mathbf{t}, V) where these functions are defined.

Point 1 requires that initial states coincide on shared variables, so there exists a state $\mathbf{v}^0 \triangleq \bigwedge_i \mathbf{v}_i^0$ defined on $\cup_i \mathcal{V}_i$ such that $\mathbf{v}_i^0 = \mathbf{v}^0|_{\mathcal{V}_i}$. Point 2 expresses that two components must declare the same shared variables for a common tile, and point 3 that on these shared variables, pre- and post-conditions must also be the same. In other words, common tiles only differ by the variables they impact in the components where they appear. So there exists a global $\gamma = \cup_i \gamma_i$ such that $\forall 1 \leq i \leq N, \forall \mathbf{t} \in$

\mathcal{T}_i , $\gamma_i(\mathbf{t}) = \gamma(\mathbf{t}) \cap \mathcal{V}_i$. Similarly, there exist partial functions α and β on $(\cup_i \mathcal{V}_i) \times (\cup_i \mathcal{T}_i)$, the restrictions of which define the α_i and β_i .

With this coherence property, and the associated definitions of α, β, γ , the composition of $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_N$ is simply defined by

$$\mathcal{S}_1 \parallel \mathcal{S}_2 \parallel \dots \parallel \mathcal{S}_N \triangleq (\cup_i \mathcal{V}_i, \cup_i \mathcal{T}_i, \wedge_i \mathbf{v}_i^0, \alpha, \beta, \gamma) \quad (1)$$

In the sequel, by abuse of notations, we may drop the index i in α_i, β_i and γ_i when defining and using coherent components¹.

Given a tile system $\mathcal{S} = (\mathcal{V}, \mathcal{T}, \mathbf{v}^0, \alpha, \beta, \gamma)$ and a subset $\mathcal{V}' \subseteq \mathcal{V}$ of variables, we define the *restriction of \mathcal{S} to variables \mathcal{V}'* as

$$\mathcal{S}_{|\mathcal{V}'} \triangleq (\mathcal{V}', \mathcal{T}', \mathbf{v}_{|\mathcal{V}'}^0, \alpha_{|\mathcal{T}' \times \mathcal{V}'}, \beta_{|\mathcal{T}' \times \mathcal{V}'}, \gamma')$$

where $\mathcal{T}' = \{\mathbf{t} \in \mathcal{T} : \mathcal{V}_t \cap \mathcal{V}' \neq \emptyset\}$ and $\forall t \in \mathcal{T}', \gamma'(t) = \gamma(t) \cap \mathcal{V}'$. In other words, \mathcal{S}' keeps the tiles of \mathcal{S} *truncated* to variables from \mathcal{V}' . For simplicity, we write $\mathcal{S}_{|\mathcal{V}'} = (\mathcal{V}', \mathcal{T}', \mathbf{v}_{|\mathcal{V}'}^0, \alpha, \beta, \gamma')$. For a given \mathcal{S} with variable set \mathcal{V} , let $\mathcal{V}_i \subseteq \mathcal{V}$ for $1 \leq i \leq N$ and define $\mathcal{S}_i^e = \mathcal{S}_{|\mathcal{V}_i}, 1 \leq i \leq N$. Then the \mathcal{S}_i^e are coherent and $\mathcal{S}_1^e \parallel \dots \parallel \mathcal{S}_N^e = \mathcal{S}_{|\cup_i \mathcal{V}_i}$.

Nets. We define a *net* as a triple $\mathcal{N} = (P, T, \rightarrow)$, where P and T are distinct denumerable *place* and *transition* sets, and $\rightarrow \subseteq (P \times T) \cup (T \times P)$ is the *flow relation*. An initial marking $P^0 \subseteq P$ may be added to this definition. A *labeled net* is a net augmented with a labeling map $\lambda : P \cup T \rightarrow \Lambda$, where Λ is a finite set of labels. Elements of $P \cup T$ are also called *nodes*. The transitive and reflexive closure of the flow relation \rightarrow on nodes is denoted by \preceq (\prec for the irreflexive part). For a node x , its *preset* and *post-set* in the net are respectively defined as $\bullet x = \{y : y \rightarrow x\}$ and $x^\bullet = \{y : x \rightarrow y\}$.

A net *homomorphism* from \mathcal{N} to \mathcal{N}' is a map $\phi : P \cup T \rightarrow P' \cup T'$ such that 1/ $\phi(P) \subseteq P'$ (with $\phi(P^0) \subseteq P'^0$), $\phi(T) \subseteq T'$, and 2/ for every transition \mathbf{t} of \mathcal{N} , ϕ restricted to $\bullet \mathbf{t}$ defines a bijection onto $\bullet \phi(\mathbf{t})$, and ϕ restricted to \mathbf{t}^\bullet defines a bijection onto $\phi(\mathbf{t})^\bullet$. If \mathcal{N} and \mathcal{N}' are labeled nets, ϕ is also required to preserve the labeling.

In a net, two nodes x, x' are said to be *in conflict*, denoted by $x \# x'$, iff there exist two transitions $\mathbf{t}, \mathbf{t}' \in T$ such that $\mathbf{t} \preceq x, \mathbf{t}' \preceq x'$ and $\bullet \mathbf{t} \cap \bullet \mathbf{t}' \neq \emptyset$.

A tile system $\mathcal{S} = (\mathcal{V}, \mathcal{T}, \mathbf{v}^0, \alpha, \beta, \gamma)$ can be associated to an equivalent net by taking $P = \mathcal{V} \times \mathcal{D}$, $P^0 = \{(V, \mathbf{v}^0(V)), V \in \mathcal{V}\}$, $T = \mathcal{T}$, and defining \rightarrow by $\forall \mathbf{t} \in \mathcal{T}, \bullet \mathbf{t} = \{(V, \alpha(\mathbf{t}, V)), V \in \mathcal{V}_t\}$ and $\mathbf{t}^\bullet = \{(V, \beta(\mathbf{t}, V)), V \in \mathcal{V}_t\}$. This allows the use of notations $\bullet \mathbf{t}$ and \mathbf{t}^\bullet for tiles. Equivalence trivially extends to runs of the two models. The equivalent net of \mathcal{S} is safe by construction. Conversely, one easily checks that a safe net can be turned into an equivalent tile system.

Occurrence nets. A net $\mathcal{O} = (C, E, \rightarrow)$ is said to be an *occurrence net* iff it satisfies the following properties

1. $\forall x \in C \cup E, \neg(x \prec x) : \prec$ is a partial order (or \rightarrow defines a directed acyclic graph (DAG) on nodes),
2. $\forall x \in C \cup E, |\{y : y \prec x\}| < \infty : \prec$ is well founded,
3. $\forall c \in C, |\bullet c| \leq 1 : \text{each place has at most one input transition,}$
4. $\forall e \in E, |\bullet e| \geq 1 : \text{each transition has at least one input place,}$
5. $\forall x \in C \cup E, \neg(x \# x) : \text{no node is in self-conflict.}$

¹It will be only maintained in γ_i in cases where confusions must be avoided, i.e. when components share tiles.

In an occurrence net, places (C) are called *conditions*, and transitions (E) are called *events*. Observe that minimal nodes of \mathcal{O} are conditions: $\min(\mathcal{O}) \subseteq C$. The partial order \prec defines the *causality relation*. Nodes x, x' are said to be *concurrent*, denoted by $x \perp x'$, iff $x \neq x'$ and neither $x \# x'$ nor $x \prec x'$ nor $x' \prec x$ holds. A *co-set* $X \subseteq C$ is a set of pairwise concurrent conditions. Maximal co-sets (w.r.t. inclusion) are called *cuts*.

$\mathcal{O}' = (C', E', \rightarrow')$ is a *prefix* of \mathcal{O} , denoted by $\mathcal{O}' \sqsubseteq \mathcal{O}$, iff \mathcal{O}' is a subnet of \mathcal{O} , that is the restriction of \mathcal{O} to $C' \cup E'$, where C' and E' are left causally closed subsets of C and E in \mathcal{O} , and satisfy $\min \mathcal{O} \subseteq C'$ and $\forall e \in E, [e \in E' \Rightarrow e^\bullet \in C']$. A *configuration* κ is a conflict-free prefix² of \mathcal{O} . An *extremal* node of κ is a maximal element for \prec (necessarily a condition). Extremal nodes of a configuration form a cut, and conversely the causal closure of a cut is a configuration.

Branching processes. Occurrence nets are useful to represent runs of a tile system under so-called *true concurrency semantics*. The labeled occurrence net $\mathcal{O} = (C, E, \rightarrow, \lambda)$ is a *branching process* of the tile system $\mathcal{S} = (\mathcal{V}, \mathcal{T}, \mathbf{v}^0, \alpha, \beta, \gamma)$ iff

1. λ is a net homomorphism between (C, E, \rightarrow) and the equivalent net of \mathcal{S} ; in particular, conditions are labeled by pairs (variable, value of that variable), and events are labeled by tiles,
2. λ defines a bijection between $\min(\mathcal{O})$ and $\{(V, \mathbf{v}^0(V)), V \in \mathcal{V}\}$,
3. $\forall e, e' \in E, \bullet e = \bullet e'$ and $\lambda(e) = \lambda(e')$ together imply $e = e'$.

With this definition, a branching process of tile system \mathcal{S} is nothing more than a branching process of its safe net version.

A configuration κ of a branching process of \mathcal{S} encodes a run of \mathcal{S} in the following way. Let e_1, e_2, \dots be *any* linear extension of partial order \prec reduced to events of κ . Then the sequence of tiles $\lambda(e_1), \lambda(e_2), \dots$ is firable at the initial state \mathbf{v}^0 of \mathcal{S} . Moreover, whatever the sequence chosen, the final state of \mathcal{S} is the same. Therefore, configurations of a branching process encode *equivalence classes* of sequences of events. It is precisely this use of concurrency which reduces the set of possible runs of a system, and makes partial order techniques attractive for systems with limited interactions.

Two branching processes \mathcal{O}_1 and \mathcal{O}_2 have a maximal common prefix $\mathcal{O}_1 \cap \mathcal{O}_2$, defined up to a unique isomorphism. This common prefix, and the corresponding labeled net isomorphisms between $\mathcal{O}_1 \cap \mathcal{O}_2$ and the \mathcal{O}_i , can be defined recursively, starting at $\min(\mathcal{O}_1)$ and $\min(\mathcal{O}_2)$, and checking the presence of equally labeled successor events for every pair of isomorphic co-sets of conditions. A similar procedure allows to prove that two branching processes $\mathcal{O}_1, \mathcal{O}_2$ of \mathcal{S} are isomorphic iff each configuration of \mathcal{O}_1 is isomorphic to a configuration of \mathcal{O}_2 , and conversely (lemma 1 in [1]). Finally, the union of branching processes \mathcal{O}_1 and \mathcal{O}_2 is defined as the minimal occurrence net having \mathcal{O}_1 and \mathcal{O}_2 as prefixes; it can again be defined recursively. There exists a unique branching process having all branching processes as prefixes; it is called the *unfolding* of the tile system \mathcal{S} , denoted by $\mathcal{U}_{\mathcal{S}}$ (see theorem 23 in [5]).

2.2 Augmentation of occurrence nets and of branching processes

We now propose the notions of augmented occurrence nets and augmented branching processes, in order to describe runs of a system carrying extra causality relations, and to keep track of extra conflict relations as well. These structures are obtained by merging the standard definition of occurrence nets (resp. branching processes) with the definition of prime event structures [2].

²We choose to define configurations as sub-nets, for a matter of homogeneity. This contrasts with the standard definition which rather considers subsets of events.

Augmented occurrence nets. An *augmented occurrence net* (AON) $\dot{\mathcal{O}} = (C, E, \rightarrow, \dot{\prec}, \dot{\#})$ is obtained by adjoining a causality and a conflict relation to the occurrence net $\mathcal{O} = (C, E, \rightarrow)$, in such a way that

1. $\dot{\prec}$ is a partial order relation on E extending the partial order \prec of (C, E, \rightarrow) :
 $\forall e, e' \in E, e \prec e' \Rightarrow e \dot{\prec} e'$
2. $\dot{\prec}$ is well-founded: $\forall e \in E, |\{e' : e' \dot{\prec} e\}| < \infty$
3. $\dot{\#}$ is a symmetric and anti-reflexive relation on E , extending $\#$ of (C, E, \rightarrow) : $\forall e, e' \in E, e \# e' \Rightarrow e \dot{\#} e'$
4. $\dot{\#}$ is inherited *via* causality $\dot{\prec}$: $\forall e, e', e'' \in E, e \dot{\#} e'$ and $e' \dot{\prec} e'' \Rightarrow e \dot{\#} e''$

The augmentation of (C, E, \rightarrow) into $(C, E, \rightarrow, \dot{\prec}, \dot{\#})$ amounts to replacing some concurrency relations between events either by a causality or by a conflict. If $\dot{\prec}$ and $\dot{\#}$ coincide with \prec and $\#$, we will not distinguish $\dot{\mathcal{O}}$ from \mathcal{O} . By abuse of terminology, pairs of events related by $\dot{\#} \setminus \#$ (resp. $\dot{\prec} \setminus \prec$) are called extra conflict relations (resp. extra causality relations).

Relations $\dot{\prec}$ and $\dot{\#}$ are defined on E but naturally propagate to all nodes of $\dot{\mathcal{O}}$, i.e. to $C \cup E$. This propagation is done by transitivity for $\dot{\prec}$, and by inheritance for $\dot{\#}$. As for ordinary occurrence nets, minimal nodes of $\dot{\mathcal{O}}$ are necessarily conditions. Moreover, one has $\min \mathcal{O} = \min \dot{\mathcal{O}}$. As for standard occurrence nets, events e_1 and e_2 are said to be concurrent, denoted by $e_1 \perp e_2$, iff $\neg(e_1 \dot{\prec} e_2)$, $\neg(e_2 \dot{\prec} e_1)$ and $\neg(e_1 \dot{\#} e_2)$. Co-sets and cuts are defined from \perp in the usual way.

The notion of *prefix* differs slightly from the one defined for standard occurrence nets. The AON $\dot{\mathcal{O}}' = (C', E', \rightarrow', \dot{\prec}', \dot{\#}')$ is a prefix of $\dot{\mathcal{O}} = (C, E, \rightarrow, \dot{\prec}, \dot{\#})$, still denoted by $\dot{\mathcal{O}}' \sqsubseteq \dot{\mathcal{O}}$, iff

- $\min \dot{\mathcal{O}} \subseteq C' \subseteq C, E' \subseteq E$,
- C' and E' are left causally closed in $\dot{\mathcal{O}}$ for $\dot{\prec}$,
- $\forall e \in E, [e \in E' \Rightarrow e^\bullet \in C']$,
- \rightarrow' and $\dot{\prec}'$ are the restrictions of \rightarrow and $\dot{\prec}$ to nodes of $\dot{\mathcal{O}}'$, and finally
- $\dot{\#}'$ *contains* the restriction of $\dot{\#}$ to E' .

In other words, $\dot{\mathcal{O}}'$ is a causally closed sub-net of $\dot{\mathcal{O}}$, for $\dot{\prec}$, where the conflict relation is possibly reinforced. It is therefore sufficient to reinforce the conflict relation $\dot{\#}$ of $\dot{\mathcal{O}}$ to obtain a strict prefix of $\dot{\mathcal{O}}$ (notice that this reinforcement must however yield a valid augmented occurrence net). This somehow unusual requirement on conflict relations appears to ensure that $\dot{\mathcal{O}}' \sqsubseteq \dot{\mathcal{O}}$ iff every configuration of $\dot{\mathcal{O}}'$ is (isomorphic to) a configuration of $\dot{\mathcal{O}}$, as we will see below.

An *augmented configuration* κ of $\dot{\mathcal{O}}$ (or simply *configuration*, for short) is a conflict-free (for $\dot{\#}$) prefix of $\dot{\mathcal{O}}$. The augmented configuration κ is thus also a configuration of \mathcal{O} enriched with extra causality relations on its events. But the converse doesn't hold: a configuration of \mathcal{O} , may not be transformable into a configuration of $\dot{\mathcal{O}}$, even if causality relations are added, because $\dot{\#}$ is stronger than $\#$. We denote by κ_e the smallest configuration containing event e ; κ_e is formed by all events in the (left) causal closure of e , plus pre- and post-conditions of these events, and $\min \dot{\mathcal{O}}$.

A homomorphism of occurrence nets from $\dot{\mathcal{O}}$ to $\dot{\mathcal{O}}'$ is a net homomorphism ϕ from (C, E, \rightarrow) to (C', E', \rightarrow') that partially maps relations $\dot{\prec}$ and $\dot{\#}$ to $\dot{\prec}'$ and $\dot{\#}'$ (some of them may be erased, but none can be created), i.e.

$$\begin{aligned} \forall e_1, e_2 \in E, \quad \phi(e_1) \dot{\prec}' \phi(e_2) &\Rightarrow e_1 \dot{\prec} e_2 \\ \text{and } \phi(e_1) \dot{\#}' \phi(e_2) &\Rightarrow e_1 \dot{\#} e_2 \end{aligned}$$

Notice that causality and conflict relations due to \rightarrow are preserved, only extra relations of $\dot{\prec} \setminus \prec$ and $\dot{\#} \setminus \#$ can be erased and transformed into concurrency. If $\dot{\mathcal{O}}$ and $\dot{\mathcal{O}}'$ are isomorphic, then there exists a one to one mapping ϕ between them which exactly maps $\dot{\prec}$ to $\dot{\prec}'$ and $\dot{\#}$ to $\dot{\#}'$.

Canonical representation. To reduce the inhomogeneity of representations for relations $\dot{\prec}$ and $\dot{\#}$ with respect to \prec and $\#$, it is convenient to represent extra relations on E by means of dummy conditions incorporated to \mathcal{O} , as in figure 1. This allows to represent $\dot{\mathcal{O}}$ as a standard occurrence net $(C \cup \dot{C}, E, \rightarrow)$. A canonical representation of $\dot{\mathcal{O}}$ can be obtained by the following procedure :

1. for every pair of events (e, e') such that $e \dot{\prec} e'$ and $\neg(e \prec e')$, create a new condition c in \dot{C} such that $e \rightarrow c \rightarrow e'$,
2. for every pair of events (e, e') such that $e \dot{\#} e'$ and $\neg(e \# e')$, create a new condition c in \dot{C} such that $c \rightarrow e$ and $c \rightarrow e'$,
3. remove from \dot{C} every condition c representing a causality link that can be obtained by transitivity (of \rightarrow on the net),
4. remove from \dot{C} every condition c representing an inherited conflict.

One easily checks that, whatever the ordering in which useless conditions of \dot{C} are removed, the same final net is obtained, which proves the canonicity of this representation.




Figure 1: *Representation of extra causality and conflict relations by means of extra conditions, relating pairs of events.*

Augmented branching processes. Let $\dot{\mathcal{O}} = (C, E, \rightarrow, \lambda, \dot{\prec}, \dot{\#})$ be a labeled augmented occurrence net. $\dot{\mathcal{O}}$ is an *augmented branching process* (ABP) of the tile system $\mathcal{S} = (\mathcal{V}, \mathcal{T}, \mathbf{v}^0, \alpha, \beta, \gamma)$ iff

1. λ is a net homomorphism between (C, E, \rightarrow) and the equivalent net of \mathcal{S} ; in particular, conditions are labeled by pairs (variable,value of that variable), and events are labeled by tiles,
2. λ defines a bijection between $\min(\dot{\mathcal{O}})$ and $\{(V, \mathbf{v}^0(V)), V \in \mathcal{V}\}$,
3. let κ, κ' be two isomorphic finite configurations of $\dot{\mathcal{O}}$, then $\kappa = \kappa'$.

Point 3 assumes the isomorphism between κ and κ' preserves the labeling, so this condition generalizes point 3 in the definition of ordinary branching processes of \mathcal{S} , and is equivalent to it when $\dot{\prec} = \prec$ and $\dot{\#} = \#$. Nevertheless, observe that condition 3 in the definition of standard branching processes is no longer valid for an *augmented* branching process. There may exist two events e and e' with identical labels and connected to the same co-set X if, for example, e and e' are not related to other events of $\dot{\mathcal{O}}$ in the same manner for $\dot{\prec}$. Fig. 2 gives an example of such a situation, where e, e' are connected to the same configuration κ , but assume different relations to other events of κ . As a consequence of this remark, keeping the $(C, E, \rightarrow, \lambda)$ part in an augmented branching process $\dot{\mathcal{O}}$ of \mathcal{S} doesn't yield a standard branching process of \mathcal{S} . However, every configuration κ of $(C, E, \rightarrow, \lambda)$ is isomorphic to

a configuration κ' of $\mathcal{U}_{\mathcal{S}}$, which means that $(C, E, \rightarrow, \lambda)$ can be trimmed into a branching process \mathcal{O} of \mathcal{S} . Specifically, there exists a unique branching process \mathcal{O} of \mathcal{S} such that there exists a surjective homomorphism $\phi : (C, E, \rightarrow, \lambda) \rightarrow \mathcal{O}$.

As for ordinary branching processes, in the sequel we do not distinguish isomorphic ABPs.

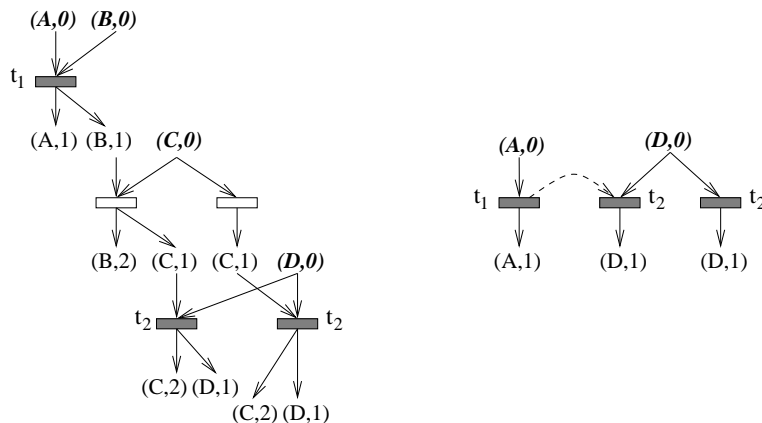


Figure 2: The ordinary branching process (left) is restricted to events impacting variables A and D . This yields the augmented branching process on the right, where two events labeled t_2 appear after condition labeled $(D, 0)$, one concurrent with t_1 , the other causally related to t_1 (dashed arrow).

Remarks.

1. Given an augmented branching process $\dot{\mathcal{O}} = (C, E, \rightarrow, \lambda, \dot{\prec}, \dot{\#})$, its restriction $(E, \dot{\prec}, \dot{\#}, \lambda)$ is an ordinary (labeled) prime event structure. Moreover, this correspondence is bijective, i.e. given the event structure for a system \mathcal{S} and a labeling of events, one can reconstruct conditions C and the flow relation \rightarrow . Nevertheless, for practical purposes rather than theoretical ones, we prefer to preserve a reference to conditions and to the flow relation. In the sequel, conditions are central in the definition of interactions between components. And the flow relation allows a direct distinction between structural relations (causality and conflict) in a component, and relations due to neighboring components. Finally, this way of presenting results allows a direct link with part 1 of this work.
2. In the sequel, we will describe operations on ABPs using the form $\dot{\mathcal{O}} = (C, E, \rightarrow, \lambda, \dot{\prec}, \dot{\#})$. In practice, most of these operations should be translated in terms of canonical representations $(C, \dot{C}, E, \rightarrow, \lambda)$. However, this paper focuses on formal aspects of computations and technicalities due to canonical forms will be left aside.

3 Properties and basic operations on ABPs

3.1 Basic properties

We first prove that if there are several identically labeled events connected to some co-set X , this can only be due to the fact that these events have different sets of predecessors for $\dot{\prec}$, that we call their “pasts.” In other words, it cannot be due to the fact that these events participate in different extra conflict relations.

Let $\dot{\mathcal{O}}$ be an augmented branching process of \mathcal{S} and e_1, e_2 be two events of $\dot{\mathcal{O}}$ such that $\bullet_{e_1} = X = \bullet_{e_2}$ and $\lambda(e_1) = \lambda(e_2)$. Let κ_{e_i} be the smallest configuration containing event e_i , and define κ'_{e_i} , the strict past of e_i , as the smallest configuration containing $\{e \in E : e \dot{\prec} e_i\}$. For obvious reasons, we can write $\kappa'_{e_i} = \kappa_{e_i} \setminus e_i$. If e_1 and e_2 have identical pasts, i.e. κ'_{e_1} and κ'_{e_2} coincide everywhere, it is

straightforward to check that κ_{e_1} and κ_{e_2} are isomorphic. By point 3 in the definition of an ABP, they are thus identical which proves $e_1 = e_2$.

To help intuition, one could then summarize things in the following manner: extra causality relations allow to augment the structure of an ordinary branching process by connecting several occurrences of the same tile after some co-set of conditions, provided these events assume different pasts. And, by contrast, extra conflict relations prevent reading out unwanted configurations on this augmented structure.

By contrast with ordinary branching processes, however, there may be an infinite number of identically labeled events e_1, e_2, \dots connected to some co-set X , each associated to a different past. Let us define the height of an augmented configuration as the maximal size over subsets $\{e_1, e_2, \dots, e_n\}$ of its events satisfying $e_1 \dot{\prec} e_2 \dot{\prec} \dots \dot{\prec} e_n$. Similarly, define the height of an event as the height of its past (including itself). The height of an augmented configuration is thus the maximal height of its events.

Lemma 1 *Any augmented branching process $\dot{\mathcal{O}}$ of a tile system \mathcal{S} has finite width: for all $n \geq 0$, the restriction of $\dot{\mathcal{O}}$ to nodes belonging to a configuration of height lower than n yields a finite augmented branching process.*

Proof. The fact that $\dot{\mathcal{O}}$ restricted to events of height lower than n is an augmented branching process is obvious, so we only consider the finiteness.

The lemma obviously holds for $n = 0$, so assume the result holds for height n . Consider an event e of height $n + 1$, $\kappa_e \setminus e$ is necessarily a configuration of height n . There is a finite number of such configurations, a finite number of tiles one can connect to each configuration, and a finite number of possible extra causality relations between the new event and this configuration, since the latter is finite. (Recall that extra conflict relations do not allow to connect more events.) So there exists a finite number of events of height $n + 1$. \square

This result allows to consider an augmented branching process as the supremum of an increasing sequence of finite prefixes, which legitimates some recursive proofs for ABPs, as in lemma 2 below.

One can add a conflict relation between e_1 and e_2 , in some augmented branching process $\dot{\mathcal{O}}$, and still get another ABP provided: 1/ $e_1 \not\# e_2$ in $\dot{\mathcal{O}}$, 2/ e_1 and e_2 have no common successors, i.e. $\{e \in E : e_1 \dot{\prec} e, e_2 \dot{\prec} e\} = \emptyset$, and 3/ the new conflict relation is propagated to successors of e_1 and e_2 . The converse operation is more interesting.

Lemma 2 *Let $\dot{\mathcal{O}} = (C, E, \rightarrow, \lambda, \dot{\prec}, \#)$ be an augmented branching process of \mathcal{S} , then $\dot{\mathcal{O}}' = (C, E, \rightarrow, \lambda, \dot{\prec}, \#)$, obtained by removing all extra conflict relations of $\# \setminus \#$ in $\dot{\mathcal{O}}$, is still an augmented branching process of \mathcal{S} . We call it the structure of $\dot{\mathcal{O}}$.*

Proof. We first show that the result holds if a single extra $\#$ conflict relation is removed. Assume $e_1 \# e_2$ is a *primary* extra conflict relation, i.e. $e_1 \# e_2$ doesn't hold, and $e_1 \# e_2$ is not inherited. So removing only relation $e_1 \# e_2$ yields $\dot{\mathcal{O}}'$ which is still an augmented occurrence net. We want to prove that $\dot{\mathcal{O}}'$ remains an augmented branching process, i.e. that point 3 in the definition still holds. Let κ, κ' be two isomorphic finite configurations on $\dot{\mathcal{O}}'$, and assume κ was not a configuration of $\dot{\mathcal{O}}$. So κ contains events e_1 and e_2 . We have to prove $\kappa = \kappa'$. Since e_1 and e_2 don't have common successors in $\dot{\mathcal{O}}$, κ decomposes as $\kappa_1 \cup \kappa_2$, where κ_i contains e_i and is a configuration of $\dot{\mathcal{O}}$, $i = 1, 2$. Let κ'_i be the image of κ_i under the isomorphism relating κ and κ' . Then $\kappa' = \kappa'_1 \cup \kappa'_2$, and the κ'_i are configurations of $\dot{\mathcal{O}}$. Using the fact that $\dot{\mathcal{O}}$ is an augmented branching process, one has $\kappa_i = \kappa'_i$, $i = 1, 2$, so $\kappa = \kappa'$.

By recursion, one can remove all extra conflict relations in any finite prefix of $\dot{\mathcal{O}}$, and still get an ABP. This allows to build an increasing sequence of finite ABPs (for the prefix relation), the supremum of which is $\dot{\mathcal{O}}' = (C, E, \rightarrow, \lambda, \dot{\prec}, \#)$. So $\dot{\mathcal{O}}'$ is an ABP. \square

Remark. This type of argument, to go from finite to infinite ABPs, will be used implicitly in the sequel: most proofs will assume finite ABPs, even if results are stated for possibly infinite ABPs.

Lemma 3 *Let $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ be two augmented branching processes of \mathcal{S} . If there exists a covering set of configurations of $\dot{\mathcal{O}}_1$ such that each configuration in this set is isomorphic to a configuration of $\dot{\mathcal{O}}_2$, then there exists a prefix of $\dot{\mathcal{O}}_2$ which is structurally isomorphic to $\dot{\mathcal{O}}_1$. If the converse property also holds, $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ are structurally isomorphic.*

Proof. We prove the lemma for finite ABPs (then concluding as in lemma 2).

As for standard branching processes, the proof consists in gluing the different isomorphisms ϕ relating pairs (κ_1, κ_2) of configurations of $\dot{\mathcal{O}}_1 \times \dot{\mathcal{O}}_2$, where κ_1 is taken in a covering set of $\dot{\mathcal{O}}_1$. Let $\phi : \kappa_1 \rightarrow \kappa_2$ and $\phi' : \kappa'_1 \rightarrow \kappa'_2$ be two such isomorphisms. ϕ and ϕ' necessarily coincide on $\min \dot{\mathcal{O}}_1$ (point 2 in the definition of an ABP), and, by recursion on events, they also coincide on $\kappa_1 \cap \kappa'_1$ (otherwise $\dot{\mathcal{O}}_2$ violates point 3 in the definition of an ABP), so they can be merged into a mapping between $\kappa_1 \cup \kappa'_1$ and $\kappa_2 \cup \kappa'_2$. Let Φ be the global mapping obtained by gluing all such isomorphisms. Φ is injective, and its image in $\dot{\mathcal{O}}_2$ is causally closed. So Φ defines a bijection between $\dot{\mathcal{O}}_1$ and some prefix of $\dot{\mathcal{O}}_2$, preserves the labeling, transforms \rightarrow_1 into \rightarrow_2 and thus $\#_1$ into $\#_2$, and transforms also \prec_1 into \prec_2 since it maps configurations.

If the criterion of the lemma holds in both directions, then there exists a Φ' from $\dot{\mathcal{O}}_2$ to some prefix of $\dot{\mathcal{O}}_1$, with the same properties. This is more than enough to prove that $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ have isomorphic structures: Φ and Φ' are inverses of one another. \square

Lemma 4 *Let $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ be two augmented branching processes of \mathcal{S} . If each configuration of $\dot{\mathcal{O}}_1$ is isomorphic to a configuration of $\dot{\mathcal{O}}_2$, then $\dot{\mathcal{O}}_1$ is isomorphic to a prefix of $\dot{\mathcal{O}}_2$. If the converse property also holds, then $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ are isomorphic.*

Proof. By lemma 3, we already know that the structure of $\dot{\mathcal{O}}_1$ is mapped to the structure of some prefix of $\dot{\mathcal{O}}_2$. To prove that the Φ defined in the previous proof is a homomorphism, we only have to check that every conflict relation $\#_2$ in $\Phi(\dot{\mathcal{O}}_1)$ corresponds to a conflict in $\dot{\mathcal{O}}_1$. Assume this is not the case: $\Phi(e_1) \#_2 \Phi(e'_1)$ in $\dot{\mathcal{O}}_2$, but $\neg(e_1 \#_1 e'_1)$ in $\dot{\mathcal{O}}_1$. Since $e_1 \prec_1 e'_1$ and $e'_1 \prec_1 e_1$ are both impossible (recall that Φ preserves causality), one necessarily has $e_1 \perp_1 e'_1$. So there exists a configuration of $\dot{\mathcal{O}}_1$ containing both e_1 and e'_1 . Its image thus contains $\Phi(e_1)$ and $\Phi(e'_1)$, and is a configuration, whence a contradiction. In summary, $\Phi : \dot{\mathcal{O}}_1 \rightarrow \dot{\mathcal{O}}_2$ preserves \prec_1 but may lose relations $\#_1$. By reinforcing $\#_2$ in $\Phi(\dot{\mathcal{O}}_1)$, we can thus make Φ an isomorphism between $\dot{\mathcal{O}}_1$ and its image.

When the criterion of the lemma holds in both directions, Φ is a bijection and preserves the structure \mathcal{S} (lemma 3), but also $\#_1$ is exactly mapped to $\#_2$, so Φ is an isomorphism between $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$. \square

3.2 Intersection

Intersection and union of ordinary BPs are the key to develop the notion of unfolding. For augmented BPs, these operations are less easy to define in general, essentially because of the extra conflict relations. However, the structures of ABPs behave like standard BPs.

Lemma 5 *Let $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ be two augmented branching processes of \mathcal{S} , there exists a unique augmented branching process $\dot{\mathcal{O}}$ of \mathcal{S} satisfying:*

1. *if κ_1, κ_2 are isomorphic configurations of $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$, there exists a configuration κ of $\dot{\mathcal{O}}$ isomorphic to both of them,*
2. *and, conversely, each configuration κ of $\dot{\mathcal{O}}$ is isomorphic to a configuration of each $\dot{\mathcal{O}}_i$.*

$\dot{\mathcal{O}}$ is called the intersection of $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$, and denoted by $\dot{\mathcal{O}}_1 \cap \dot{\mathcal{O}}_2$.

Proof. Assume $\dot{\mathcal{O}} = \dot{\mathcal{O}}_1 \cap \dot{\mathcal{O}}_2$, and consider the associated structures $\dot{\mathcal{O}}'$, $\dot{\mathcal{O}}'_1$ and $\dot{\mathcal{O}}'_2$. We first prove that $\dot{\mathcal{O}}' = \dot{\mathcal{O}}'_1 \cap \dot{\mathcal{O}}'_2$.

Observe that $\dot{\mathcal{O}}'$ satisfies the two conditions of the lemma with respect to $\dot{\mathcal{O}}'_1$ and $\dot{\mathcal{O}}'_2$ for configurations formed by the causal closure of some event (such configurations are common to an ABP and to its structure). Since this family of configurations forms a covering set for $\dot{\mathcal{O}}'$, lemma 3 implies that $\dot{\mathcal{O}}'$ is isomorphic to a prefix of $\dot{\mathcal{O}}'_i$, $i = 1, 2$, which proves point 2. For point 1, let κ_1, κ_2 be two isomorphic configurations of $\dot{\mathcal{O}}'_1, \dot{\mathcal{O}}'_2$. Since each κ_i is covered by the causal closure of its events, with the isomorphism extension argument used to prove lemma 3, one easily checks that both κ_i are isomorphic to the same configuration κ of $\dot{\mathcal{O}}'$. This proves point 1, so one has $\dot{\mathcal{O}}' = \dot{\mathcal{O}}'_1 \cap \dot{\mathcal{O}}'_2$.

To define $\dot{\mathcal{O}}$, we thus start by building $\dot{\mathcal{O}}'$ and the isomorphisms ϕ_i between a prefix of $\dot{\mathcal{O}}'_i$ and $\dot{\mathcal{O}}'$. We proceed by recursion, starting with $\min \dot{\mathcal{O}}'$ isomorphic to $\min \dot{\mathcal{O}}'_i$, and connect one event at a time (and its successive conditions) until all pairs of isomorphic configurations of $\dot{\mathcal{O}}'_1 \times \dot{\mathcal{O}}'_2$ have been considered. Specifically, let (κ_1, κ_2) and (κ'_1, κ'_2) be pairs of isomorphic configurations, satisfying $\kappa'_i = \kappa_i \setminus e_i$ for some event e_i . Assume there exists already a configuration $\kappa' = \phi_1(\kappa'_1) = \phi_2(\kappa'_2)$ in $\dot{\mathcal{O}}'$, but ϕ_i is not defined on e_i , $i = 1, 2$. Then connect an extra event e (and successive conditions) to co-set $\phi_1(\bullet e_1) = \phi_2(\bullet e_2)$ of κ' , set $e' \prec e$ for every event e' in κ' such that $e' = \phi_i(e'_i)$ and $e'_i \prec_i e_i$, and finally extend ϕ_i to e_i and $e \bullet_i$, in such a way that κ_i is isomorphic to $\kappa = \kappa' \cup \{e\} \cup e \bullet$ through ϕ_i . When the process stops, $\dot{\mathcal{O}}'$ satisfies the conditions of the lemma, by construction. This proves existence and uniqueness of $\dot{\mathcal{O}}'$.

The occurrence net $\dot{\mathcal{O}}'$ obtained in that way is obviously an ABP of \mathcal{S} . Any pair (κ_1, κ_2) of isomorphic configurations of $\dot{\mathcal{O}}_1 \times \dot{\mathcal{O}}_2$ has a representative κ on $\dot{\mathcal{O}}'$. To ensure the converse property, one must prevent reading out from $\dot{\mathcal{O}}'$ configurations which have no counterpart on one of the $\dot{\mathcal{O}}_i$. We thus define $\dot{\mathcal{O}}$ from $\dot{\mathcal{O}}'$ by setting $e \# e'$ iff $\phi_1^{-1}(e) \#_1 \phi_1^{-1}(e')$ or $\phi_2(e) \#_2 \phi_2^{-1}(e')$. \square




Figure 3: *Two ABPs (left and right) and their intersection (center). Only primary extra conflict relations are represented for clarity; so the inherited extra conflict relation $\mathbf{t}_1 \# \mathbf{t}_4$ is not represented.*

The intersection defined above is associative. Notice also that it selects configurations with *identical* structures. If $\dot{\mathcal{O}}_1$ contains $\mathbf{t}_1 \perp \mathbf{t}_2$ and $\dot{\mathcal{O}}_2$ contains $\mathbf{t}_1 \rightarrow \mathbf{t}_2$, then the intersection reduces to \mathbf{t}_1 and not to $\mathbf{t}_1 \rightarrow \mathbf{t}_2$, even if $\dot{\mathcal{O}}_1$ allows to fire \mathbf{t}_1 then \mathbf{t}_2 . This will contrast with the product of augmented branching processes, that we define in section 4 (see fig. 6 for a comparison). Finally, observe that $\dot{\mathcal{O}}_1 \cap \dot{\mathcal{O}}_2 \sqsubseteq \dot{\mathcal{O}}_i$, as usually; $\dot{\mathcal{O}}_1 \cap \dot{\mathcal{O}}_2$ can actually be defined as the maximal common prefix of the $\dot{\mathcal{O}}_i$ (see below).

3.3 Union

Lemma 6 *Let $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ be two augmented branching processes of \mathcal{S} , there exists a unique ABP $\dot{\mathcal{O}} \triangleq \dot{\mathcal{O}}_1 \cup \dot{\mathcal{O}}_2$ of \mathcal{S} satisfying*

1. every configuration κ_i of $\dot{\mathcal{O}}_i$ is isomorphic to a configuration κ of $\dot{\mathcal{O}}$, $i = 1, 2$,
2. for every configuration κ of $\dot{\mathcal{O}}$ there exists $i \in \{1, 2\}$ such that κ is isomorphic to a configuration κ_i of $\dot{\mathcal{O}}_i$.

Proof. The proof is constructive, as for the definition of intersection. Assuming such a $\dot{\mathcal{O}}$ exists, let us consider its structure $\dot{\mathcal{O}}'$. Similarly, let $\dot{\mathcal{O}}'_i$ denote the structure of $\dot{\mathcal{O}}_i$, $i = 1, 2$. $\dot{\mathcal{O}}'$ is characterized by its set of configurations κ_e for e ranging over all its events (lemma 3). A configuration κ_e of $\dot{\mathcal{O}}'$ is also a configuration of $\dot{\mathcal{O}}$, and thus is isomorphic to a configuration κ_{e_i} of $\dot{\mathcal{O}}_i$, for $i = 1$ or $i = 2$ (point 2 of the lemma), which is also a configuration of the structure $\dot{\mathcal{O}}'_i$. Conversely, every κ_{e_i} of $\dot{\mathcal{O}}'_i$ must have an isomorphic counterpart in $\dot{\mathcal{O}}'$, by the first condition of the lemma. We thus have a complete description of all configurations κ_e that must appear in $\dot{\mathcal{O}}'$. We now show that there exists a $\dot{\mathcal{O}}'$ satisfying these conditions, by constructing one. (Again, by lemma 3, if it exists, it is unique.)

The construction follows the same lines as the recursive construction of $\dot{\mathcal{O}}'_1 \cap \dot{\mathcal{O}}'_2$. After all pairs (κ_1, κ_2) of isomorphic configurations have been considered, “private events” e_i of $\dot{\mathcal{O}}'_i$ are explored, i.e. events such that κ_{e_i} has no isomorphic configuration in the other structure $\dot{\mathcal{O}}'_j$. A new event e is then connected to $\dot{\mathcal{O}}'$ (with extension of ϕ_i) to ensure that κ_{e_i} is isomorphic to $\kappa_e = \phi_i(\kappa_{e_i})$, until all private events have been considered. Observe that with this construction, the ϕ_i are defined on the whole $\dot{\mathcal{O}}'_i$ instead of a prefix, and $\phi_1(\dot{\mathcal{O}}'_1) \cap \phi_2(\dot{\mathcal{O}}'_2)$ yields $\dot{\mathcal{O}}'_1 \cap \dot{\mathcal{O}}'_2$.

To complete the proof, we must show that there exists a unique way to define extra conflict relations over events of $\dot{\mathcal{O}}'$ such that the two conditions of the lemma are satisfied. We first have to prevent reading out from $\dot{\mathcal{O}}$ configurations using private events of $\dot{\mathcal{O}}_1$ and of $\dot{\mathcal{O}}_2$. So necessarily $e \# e'$ for pairs of events (e, e') in $[\phi_1(\dot{\mathcal{O}}'_1) \setminus \phi_2(\dot{\mathcal{O}}'_2)] \times [\phi_2(\dot{\mathcal{O}}'_2) \setminus \phi_1(\dot{\mathcal{O}}'_1)]$. Observe that such pairs (e, e') necessarily correspond to concurrent events in $\dot{\mathcal{O}}'$. Secondly, let us set $e \# e'$ as soon as $\phi_i^{-1}(e) \#_i \phi_i^{-1}(e')$ for all i (at least one of these mappings is defined), since e and e' can never appear in the same configuration. With this definition of $\#$, every configuration κ of $\dot{\mathcal{O}}$ lies in (at least) one of the images $\phi_i(\dot{\mathcal{O}}_i)$, and is isomorphic to (at least) one of the $\phi_i^{-1}(\kappa)$, so point 2 is satisfied. Moreover, the ϕ_i are structure preserving morphisms (only extra conflict relations can vanish), so point 1 is satisfied. This proves the existence of $\dot{\mathcal{O}}$. It is easily checked that any stronger conflict relations on $\dot{\mathcal{O}}'$ would violate the conditions of the lemma, since it would yield a strict prefix of $\dot{\mathcal{O}}$. This proves the uniqueness of $\dot{\mathcal{O}}$. \square

This proof deserves several comments. First, by construction of morphisms ϕ_i , we have structure preserving mappings, so the $\dot{\mathcal{O}}_i$ are isomorphic to a prefix of $\dot{\mathcal{O}}$. $\dot{\mathcal{O}}$ is actually the smallest ABP of \mathcal{S} (w.r.t. \sqsubseteq) satisfying this property³. Secondly, we have obtained $\dot{\mathcal{O}} = \dot{\mathcal{O}}_1 \cup \dot{\mathcal{O}}_2$ by defining first a union of structures $\dot{\mathcal{O}}'_1$ and $\dot{\mathcal{O}}'_2$, and then taking a prefix by introducing extra conflict relations. This reveals that a notion of unfolding could be defined for ABPs of \mathcal{S} , that we could denote as $\dot{\mathcal{U}}_{\mathcal{S}}$. $\dot{\mathcal{U}}_{\mathcal{S}}$ can be defined as the union of all structures of ABPs of \mathcal{S} . Since by definition every ABP is a prefix of its structure, one thus has that $\dot{\mathcal{U}}_{\mathcal{S}}$ is also the supremum of all ABPs of \mathcal{S} . $\dot{\mathcal{U}}_{\mathcal{S}}$ is a huge object, containing for example all linear extensions of a given configuration. Therefore its existence only has a theoretical interest.

Proposition 1 *Every augmented branching process is isomorphic to the union of its configurations.*

Proof. This is an obvious application of lemma 4 to the definition of union given in lemma 6. \square

Important remark. As a consequence of this result, an ABP can be equivalently described as a set of configurations, closed for the prefix relation: $[\kappa \in \dot{\mathcal{O}}, \kappa' \sqsubseteq \kappa] \Rightarrow \kappa' \in \dot{\mathcal{O}}$. So the intersection of two ABPs is the standard intersection of sets of configurations. The union of two ABPs is the standard union of sets of configurations. And the prefix of an ABP corresponds to a subset of configurations, closed for the prefix relation on configurations.

3.4 Trimming

Generalized ABP. Let $\dot{\mathcal{O}}$ be a labeled augmented occurrence net (of finite width) satisfying points 1 and 2 in the definition of an augmented branching process of \mathcal{S} , but failing at point 3. That is, $\dot{\mathcal{O}}$

³The verification is left to the reader. Hint: check first that $\dot{\mathcal{O}}$ has the minimal structure, then the maximal conflict.

describes augmented configurations corresponding to runs of tile system \mathcal{S} , but may contain several isomorphic copies of some configurations. Slightly abusing terminology, we will say that $\dot{\mathcal{O}}$ is a *generalized* ABP (GABP) of \mathcal{S} . The *trimming* consists in eliminating the redundant copies of isomorphic configurations, in order to get a true ABP: $Trim(\dot{\mathcal{O}})$ is defined as the union of all configurations of $\dot{\mathcal{O}}$. This definition entails the existence of a canonical homomorphism from $\dot{\mathcal{O}}$ to $Trim(\dot{\mathcal{O}})$.

We briefly describe below a trimming procedure for a generalized ABP $\dot{\mathcal{O}} = (C, E, \rightarrow, \dot{\prec}, \dot{\#}, \lambda)$ of \mathcal{S} . Its first task consists in building a correct ABP structure, by recursively merging identically labeled events connected on top of the same configuration (this is very similar to a classical procedure on standard labelled occurrence nets, to obtain a branching process). The second task updates the set of extra conflict relations, when two events are merged.

Procedure 1

- recursion (while the following condition is satisfied):
 - if there exists a pair (e_1, e_2) of events such that $\lambda(e_1) = \lambda(e_2)$, $\bullet e_1 = \bullet e_2 = X$ and $\kappa_{e_1} \setminus e_1 = \kappa_{e_2} \setminus e_2$
 - merge e_1 and e_2 into a single event e , with $\lambda(e) = \lambda(e_i)$, $\bullet e = X$ (merge also conditions of e_1^\bullet and e_2^\bullet having the same label),
 - for all events $e' \in \kappa_{e_i} \setminus e_i$, set $e' \dot{\prec} e$,
 - for all events $e' \in E$, if $e_1 \dot{\prec} e'$ or $e_2 \dot{\prec} e'$ then set $e \dot{\prec} e'$,
 - for all events $e' \in E$, if $e' \dot{\#} e_1$ and $e' \dot{\#} e_2$ then set $e' \dot{\#} e$.

$\kappa_{e_i} \setminus e_i$ denotes the configuration containing all events of κ_{e_i} but e_i . Procedure 1 converges for finite AONs, doesn't depend on the ordering of operations, and yields the desired ABP (proof left to the reader). Moreover, $Trim(\dot{\mathcal{O}})$ is isomorphic to a prefix of $\dot{\mathcal{O}}$.

4 Product of ABPs

Let $\mathcal{S} = (\mathcal{V}, \mathcal{T}, \mathbf{v}^0, \alpha, \beta, \gamma)$ be a tile system, and consider restrictions $\mathcal{S}_{|\mathcal{V}'}$ with $\mathcal{V}' \subseteq \mathcal{V}$. We define the product in the category of *generalized* ABPs of restrictions of \mathcal{S} (as we will see, the product is not stable in the sub-category of ABPs). The morphisms of this category are given as follows. Let $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$ be GABPs of $\mathcal{S}_{|\mathcal{V}_1}$ and $\mathcal{S}_{|\mathcal{V}_2}$ respectively, with $\mathcal{V}_2 \subseteq \mathcal{V}_1$. A morphism $\phi : \dot{\mathcal{O}}_1 \rightarrow \dot{\mathcal{O}}_2$ is defined as a *partial* map; it operates on conditions c satisfying $\lambda(c) \in \mathcal{V}_2 \times \mathcal{D}$ and events e satisfying $\mathcal{V}_t \cap \mathcal{V}_2 \neq \emptyset$ for $t = \lambda(e)$. ϕ is required to be a standard homomorphism of labeled AONs over its domain of definition. As a consequence, the restriction⁴ $\phi : \min(\dot{\mathcal{O}}_1) \rightarrow \min(\dot{\mathcal{O}}_2)$ is bijective.

Let us make two remarks about this category. First, given $\dot{\mathcal{O}}$, one can determine the variable set \mathcal{V}' such that $\dot{\mathcal{O}}$ is a GABP of $\mathcal{S}_{|\mathcal{V}'}$ by observing labels of $\min \dot{\mathcal{O}}$. Secondly, observe that there may exist non trivial automorphisms over a GABP $\dot{\mathcal{O}}$, but if $\dot{\mathcal{O}}$ is an ABP, its only automorphism is the identity.

In the first part of this work [1], a product \wedge was defined in the sub-category of branching processes of restrictions of \mathcal{S} . Let $\mathcal{V}_i \subseteq \mathcal{V}$ for $1 \leq i \leq N$, the following factorization property was proved for unfoldings

$$\mathcal{U}_{\mathcal{S}_{|\mathcal{V}_1 \cup \dots \cup \mathcal{V}_N}} = \mathcal{U}_{\mathcal{S}_{|\mathcal{V}_1}} \wedge \dots \wedge \mathcal{U}_{\mathcal{S}_{|\mathcal{V}_N}} \quad (2)$$

We now extend this definition of \wedge , and prove that it is a categorical product.

⁴When mentioning the restriction of a partial function ψ_i to a set S , we consider only its domain of definition in that set, i.e. $S \cap \lambda^{-1}(\mathcal{V}_2 \times \mathcal{D})$ in this example.

4.1 Definition

Let $\dot{\mathcal{O}}_i = (C_i, E_i, \rightarrow_i, \lambda_i, \dot{\prec}_i, \dot{\#}_i)$ be a generalized augmented branching process of $\mathcal{S}_i^e = \mathcal{S}_{|\mathcal{V}_i} = (\mathcal{V}_i, \mathcal{T}_i^e, \mathbf{v}_i^0, \alpha, \beta, \gamma_i)$. The product $\dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$, is defined as the labeled augmented occurrence net $\dot{\mathcal{O}} = (C, E, \rightarrow, \lambda, \dot{\prec}, \dot{\#})$ satisfying

1. events are labeled by tiles of $\cup_i \mathcal{T}_i^e$, and conditions by pairs of $(\cup_i \mathcal{V}_i) \times \mathcal{D}$,
2. $\forall i \in \{1, \dots, N\}$, there exists a morphism $\psi_i : \dot{\mathcal{O}} \rightarrow \dot{\mathcal{O}}_i$
3. $\forall \kappa$ a configuration of $\dot{\mathcal{O}}$ and $\forall X$ an extremal co-set⁵ of κ ,
 $\forall \mathbf{t} \in \cup_i \mathcal{T}_i^e$ and $I = \{i : 1 \leq i \leq N, \mathbf{t} \in \mathcal{T}_i^e\}$,
 if $\forall i \in I, \exists e_i \in E_i$ such that $\lambda_i(e_i) = \mathbf{t}$, $\psi_i(X) = \bullet e_i$,
 $\psi_i(\kappa)$ is a configuration and $\kappa_{e_i} \setminus e_i \sqsubseteq \psi_i(\kappa)$,
 then there exists a unique event e in E satisfying
 - (a) $\lambda(e) = \mathbf{t}$ and $\bullet e = X$,
 - (b) $\forall i \in I, \psi_i(e) = e_i$ and the restriction $\psi_i : e^\bullet \rightarrow e_i^\bullet$ is bijective,
 - (c) $\forall e' \in \kappa \cap E, [\exists i \in I : \psi_i(e') \dot{\prec}_i \psi_i(e)] \Leftrightarrow e' \dot{\prec} e$,
 - (d) $\forall e' \in E, [\exists i \in I : \psi_i(e') \dot{\#}_i \psi_i(e)] \Leftrightarrow e' \dot{\#} e$.

This definition contrasts with the product of standard branching processes of [1] on two main aspects. First, not only the local flow relations \rightarrow_i but also the extra causality and conflict relations of factors $\dot{\mathcal{O}}_i$ are mapped to the product structure $\dot{\mathcal{O}}$. Secondly, an event e is connected both to a co-set X and to a configuration κ (instead of a co-set only). The connection to a co-set captures the causality due to the product flow structure, and the configuration κ ensures that extra causality relations in each factor are also satisfied.

As for standard branching processes, this definition leads to a constructive procedure of any finite prefix of the product. The latter is then obtained as the supremum of all finite prefixes given by procedure 2. Naturally, this assumes that the factors $\dot{\mathcal{O}}_i$ are finite width GABPs, property which is preserved by \wedge .

Procedure 2

- initialization :
 - $E = \emptyset$ ($\rightarrow, \dot{\prec}$ and $\dot{\#}$ are empty as well)
 - create $|\cup_i \mathcal{V}_i|$ initial conditions in C , with labels satisfying $\lambda(C) = \cup_i \lambda_i(\min(\dot{\mathcal{O}}_i))$
 - for $1 \leq i \leq N$, define the partial maps $\psi_i : C \rightarrow \min(\dot{\mathcal{O}}_i)$ so that they are surjective and preserve λ
- recursion :
 - for κ a configuration of $\dot{\mathcal{O}}$ and X an extremal co-set of κ ,
 - for $\mathbf{t} \in (\cup_i \mathcal{T}_i^e)$, $I = \{i : 1 \leq i \leq N, \mathbf{t} \in \mathcal{T}_i^e\}$,
 - for $\{e_i : i \in I, e_i \in E_i\}$ a set of events labeled by tile \mathbf{t} , satisfying
 $\forall i \in I, \psi_i(X) = \bullet e_i$ and $\psi_i(\kappa)$ is a configuration containing $\kappa_{e_i} \setminus e_i$,
 - create e in E such that $\bullet e = X$ and $\lambda(e) = \mathbf{t}$ (if there isn't such an event in E already), and define $\psi_i(e) = e_i$ for $i \in I$,
 - then create $|\cup_i \gamma_i(\mathbf{t})|$ new conditions c in C , with $\bullet c = e$, and assign labels to have $\lambda(e^\bullet) = \cup_{i \in I} \lambda_i(e_i^\bullet)$; extend partial maps $\psi_i, i \in I$, so that $\psi_i : e^\bullet \rightarrow e_i^\bullet$ is λ -preserving and surjective,

⁵By extremal co-set we refer to a co-set of extremal conditions of κ .

- $\forall e' \in \kappa \cap E$, set $e' \prec e$ if $\exists i \in I : \psi_i(e') \prec_i \psi_i(e)$,
- $\forall e' \in E$, set $e' \# e$ if $\exists i \in I : \psi_i(e') \#_i \psi_i(e)$.

$\dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$ is a generalized ABP of $\mathcal{S}_{|\mathcal{V}_1 \cup \dots \cup \mathcal{V}_N}$; the proof is straightforward by recursion in procedure 2. Similarly, when applied to standard branching processes, this product coincides with the one defined in [1] and yields a standard BP.

Notation. At point 3 in the definition, when I contains a unique index i , we say that event e_i is *local* to $\dot{\mathcal{O}}_i$. Otherwise, we say that events $\{e_i, i \in I\}$ *synchronize* into event e .

Remark. Given configurations $\kappa_1, \dots, \kappa_N$ in $\dot{\mathcal{O}}_1, \dots, \dot{\mathcal{O}}_N$ resp., there exists a unique maximal configuration κ of $\dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$ such that $\psi_i(\kappa) \sqsubseteq \kappa_i$. We write $\kappa = \kappa_1 \wedge \dots \wedge \kappa_N$.

4.2 Categorical product

Proposition 2 *The operator \wedge defined in the category of GABPs of restrictions of \mathcal{S} is a product in this category.*

Proof. Without loss of generality, we consider the product of two GABPs $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$ of $\mathcal{S}_{|\mathcal{V}_1}$ and $\mathcal{S}_{|\mathcal{V}_2}$ respectively. Let $\dot{\mathcal{O}} = \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2$, with associated morphisms $\psi_i : \dot{\mathcal{O}} \rightarrow \dot{\mathcal{O}}_i$, and denote $\mathcal{V}_{\dot{\mathcal{O}}} = \mathcal{V}_1 \cup \mathcal{V}_2$, $\mathcal{T}_{\dot{\mathcal{O}}} = \mathcal{T}_1^e \cup \mathcal{T}_2^e$. Assume there exists another $\dot{\mathcal{O}}'$ in the category, with morphisms $\psi'_i : \dot{\mathcal{O}}' \rightarrow \dot{\mathcal{O}}_i$. We have to prove the existence of a unique morphism $\phi : \dot{\mathcal{O}}' \rightarrow \dot{\mathcal{O}}$ such that $\psi'_i = \psi_i \circ \phi$ (universal property of a categorical product).

As a GABP of a restriction of \mathcal{S} , conditions of $\dot{\mathcal{O}}'$ are labeled by $\mathcal{V}_{\dot{\mathcal{O}}'}$, and events are labeled by $\mathcal{T}_{\dot{\mathcal{O}}'}$, the tile set of $\mathcal{S}_{|\mathcal{V}_{\dot{\mathcal{O}}'}}$. The existence of ψ'_i implies that $\mathcal{V}_i \subseteq \mathcal{V}_{\dot{\mathcal{O}}'}$, so $\mathcal{V}_{\dot{\mathcal{O}}} \subseteq \mathcal{V}_{\dot{\mathcal{O}}'}$. Similarly, $\mathcal{T}_{\dot{\mathcal{O}}} \subseteq \mathcal{T}_{\dot{\mathcal{O}}'}$. Since ϕ has to be defined on conditions and events of $\dot{\mathcal{O}}'$ labeled by $\mathcal{V}_{\dot{\mathcal{O}}}$ and $\mathcal{T}_{\dot{\mathcal{O}}}$ respectively, let us consider the restriction of $\dot{\mathcal{O}}'$ to these nodes. One easily checks that this restriction is a GABP of $\mathcal{S}_{|\mathcal{V}_{\dot{\mathcal{O}}}}$. So, without loss of generality, we can assume that $\mathcal{V}_{\dot{\mathcal{O}}} = \mathcal{V}_{\dot{\mathcal{O}}'}$ and $\mathcal{T}_{\dot{\mathcal{O}}} = \mathcal{T}_{\dot{\mathcal{O}}'}$.

We build ϕ recursively, following the scheme of procedure 2. Consider conditions of $\min \dot{\mathcal{O}}'$: they are in bijective correspondence with $\min \dot{\mathcal{O}}$, which defines ϕ on $\min \dot{\mathcal{O}}'$. The universal property (UP) is obviously satisfied by ϕ on $\min \dot{\mathcal{O}}'$: if c in $\min \dot{\mathcal{O}}'$ is labeled by variable V of \mathcal{V}_i , then there exists a unique condition in $\min \dot{\mathcal{O}}_i$ labeled by V , and similarly for $\min \dot{\mathcal{O}}$. We thus have $\psi'_i(c) = \psi_i(\phi(c))$.

Assume ϕ has been defined on some prefix $\dot{\mathcal{O}}'_0$ of $\dot{\mathcal{O}}'$, and satisfies the UP on this prefix. Let e' be an event of $\dot{\mathcal{O}}'$ such that $X' = \bullet e'$ is an extremal co-set of $\dot{\mathcal{O}}'_0$, and $\kappa_{e'} \setminus e'$ is a configuration of $\dot{\mathcal{O}}'_0$. We want to extend ϕ to e' . Assume e' is labeled by a tile of $\mathcal{T}_1^e \cap \mathcal{T}_2^e$ for example (the two other cases are examined in a similar manner). Let $e_i = \psi'_i(e')$, $i = 1, 2$, and consider $\kappa = \phi(\kappa_{e'} \setminus e')$, $X = \phi(X')$. Since ϕ satisfies the UP on $\dot{\mathcal{O}}'_0$, e_1, e_2, X and κ satisfy conditions of point 3 in the definition of $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2$, so there exists a unique event e of $\dot{\mathcal{O}}$ with $\psi_i(e) = e_i$, $\bullet e = X$, $(\kappa \setminus e) \sqsubseteq \kappa$. Let us set $\phi(e') = e$, and extend also ϕ from e'^{\bullet} to e^{\bullet} . This extends the UP of ϕ to e' and e'^{\bullet} .

By construction, the flow relation is preserved by ϕ around event e' . The last thing to prove is that ϕ remains a homomorphism, i.e. doesn't create extra conflict relations, nor extra causality relations. But such relations in $\phi(\dot{\mathcal{O}})$ are directly derived from the $\psi'_i(\dot{\mathcal{O}}')$, by definition of the product. And the ψ'_i can only erase, not create, extra causality/conflict relations of $\dot{\mathcal{O}}'$. \square

Corollary 1 *As a categorical product, \wedge is naturally associative.*

Since \wedge is stable in the sub-category of standard branching processes, and coincides with the product defined in [1], this corollary gives another proof of the associativity of \wedge on BPs.

At this point, it is worth relating the product in the category of GABPs to existing product definitions on event structures. Although GABPs carry events, conditions, labels, a flow and extra relations, they are definitely labeled prime event structures. Actually, given a labeled prime event structure $(C, E, \prec, \#, \lambda)$ and knowing that it corresponds to $\mathcal{S}_{|\mathcal{V}'}$, it is possible to recursively recover

all the underlying GABP structure. Morphisms between labeled prime event structures are defined as (label preserving) partial functions that map configurations to configurations (and thus that can only “forget” conflict or causality), which corresponds to our definition. So GABPs can be considered as a sub-category of labeled prime event structures.

As underlined in [14], labels have a minor role in the definition of the product of event structures : a label algebra [3] can be incorporated afterwards. If labels are put aside, there exist two ways of defining the product of prime event structures. One is recursive (or inductive), as in [13], while the other one directly defines legal configurations [14, 4, 15]. We are definitely on the side of the recursive definition, which has the advantage to suggest its implementation, although we have abandoned the heavy naming of events by backward pointers (used in [13] and also in [7]).

4.3 Relations to trimming

One could suspect that applying the product to ABPs (instead of GABPs) yields a trimmed structure. This is false in general : the result of the product is generally not trimmed⁶, as illustrated by the example of fig. 4. The phenomenon appears on events labeled by t_1 and t_2 on this figure, but we have added extra events in their futures to illustrate the non trivial trimming operation on the product.




Figure 4: Two ABPs \dot{O}_1 (top left) and \dot{O}_2 (top right) of components S_1^e and S_2^e sharing only variable A . Observe that t_1 operates on A, B and t_2 on B, C, D . \dot{O}_1 allows $t_1 \perp t_2$ and $t_1 \rightarrow t_2$ (these two possibilities have different futures), whereas \dot{O}_2 imposes $t_1 \rightarrow t_2$. The product $\dot{O}_1 \wedge \dot{O}_2$ is given by the bottom left graph, with its trimmed version on the right: the two occurrences of t_2 must be merged, and their conflict relation is pushed forward to t_3 and t_4 .

Despite this drawback, the product of GABPs of \mathcal{S} is insensitive to trimming in the following sense :

⁶This is the main motivation for introducing the category of generalized ABPs.

Proposition 3 Let \dot{O}_i be a GABP of $\mathcal{S}_i^e = \mathcal{S}_{|\mathcal{V}_i}$, $\mathcal{V}_i \subseteq \mathcal{V}$, $1 \leq i \leq N$, then

$$\text{Trim}(\dot{O}_1 \wedge \dots \wedge \dot{O}_N) = \text{Trim}[\text{Trim}(\dot{O}_1) \wedge \dots \wedge \text{Trim}(\dot{O}_N)] \quad (3)$$

Proof. Let κ be a configuration of $\dot{O}_1 \wedge \dots \wedge \dot{O}_N$, then κ is isomorphic to $\kappa_1 \wedge \dots \wedge \kappa_N$ with $\kappa_i = \psi_i(\kappa)$. And conversely, taking any $\kappa_i \in \dot{O}_i$, $1 \leq i \leq N$, the product $\kappa_1 \wedge \dots \wedge \kappa_N$ yields a configuration of $\dot{O}_1 \wedge \dots \wedge \dot{O}_N$. Relying on proposition 1, one can thus write

$$\text{Trim}(\dot{O}_1 \wedge \dots \wedge \dot{O}_N) = \bigcup_{\kappa_1 \in \dot{O}_1, \dots, \kappa_N \in \dot{O}_N} \kappa_1 \wedge \dots \wedge \kappa_N \quad (4)$$

where the union merges isomorphic configurations. In $\kappa = \kappa_1 \wedge \dots \wedge \kappa_N$, let us replace each κ_i by an isomorphic configuration κ'_i of \dot{O}_i , $1 \leq i \leq N$. One obtains κ' which is obviously isomorphic to κ , whence the result. \square




Figure 5: The category of GABPs is represented as a plane, and the sub-category of ABPs as a line. The figure illustrates that the trimmed product can be defined on classes of trim-equivalent GABPs.

Proposition 3 suggests to define the operator $\dot{\wedge}$ in the sub-category of ABPs by including a trimming in the definition of \wedge :

$$\dot{O}_1 \dot{\wedge} \dots \dot{\wedge} \dot{O}_N \triangleq \text{Trim}(\dot{O}_1 \wedge \dots \wedge \dot{O}_N) \quad (5)$$

However, by projecting $\dot{O}_1 \wedge \dots \wedge \dot{O}_N$ to the sub-category of ABPs, one loses morphisms ψ_i relating the product to its factors. In other words, $\dot{\wedge}$ is not a categorical product in the sub-category of ABPs. Therefore, in the sequel, we use the untrimmed version \wedge for proofs, and extend results to $\dot{\wedge}$. In practice (section 6), all computations must be performed with $\dot{\wedge}$, in order to reduce the size of the nets. However, let us keep in mind the following useful result:

Corollary 2 Define $\dot{O} \doteq \dot{O}'$ iff $\text{Trim}(\dot{O}) = \text{Trim}(\dot{O}')$. Let \dot{O}_i, \dot{O}'_i be GABPs of $\mathcal{S}_i^e = \mathcal{S}_{|\mathcal{V}_i}$, $1 \leq i \leq N$, then

$$[\forall 1 \leq i \leq N, \dot{O}_i \doteq \dot{O}'_i] \Rightarrow \dot{O}_1 \wedge \dots \wedge \dot{O}_N \doteq \dot{O}'_1 \wedge \dots \wedge \dot{O}'_N \quad (6)$$

4.4 Other properties

We have already mentioned that the product \wedge defined for GABPs generalizes the product of standard branching processes proposed in [1]. But some properties are lost. For example, if \dot{O}_1, \dot{O}_2 are two ABPs of the *same* tile system \mathcal{S} , one has $\dot{O}_1 \cap \dot{O}_2 \sqsubseteq \text{Trim}(\dot{O}_1 \wedge \dot{O}_2)$, instead of the equality that takes place in the case of standard branching processes. Figure 6 gives a counter-example where intersection and product differ. Other obvious properties of the product defined on BPs remain valid for ABPs or GABPs. We briefly mention them without proof. Let \dot{O}_i and \dot{O}'_i be GABPs of the \mathcal{S}_i^e , $1 \leq i \leq N$, then




Figure 6: *Top*: two ABPs. *Bottom*, from left to right: their product, their intersection and their union.

1. $[\forall 1 \leq i \leq N, \dot{O}'_i \sqsubseteq \dot{O}_i] \Rightarrow \wedge_i \dot{O}'_i \sqsubseteq \wedge_i \dot{O}_i$
2. $\dot{O} = \wedge_i \dot{O}_i$ and $\dot{O}'_i = \psi_i(\dot{O}) \Rightarrow \dot{O} = \wedge_i \dot{O}'_i$
3. $\dot{O} = \wedge_i \dot{O}_i$ and $\dot{O}' \sqsubseteq \dot{O} \Rightarrow \psi_i(\dot{O}') \sqsubseteq \psi_i(\dot{O}) \sqsubseteq \dot{O}_i$

Finally, let us mention the following important result :

Lemma 7 *Let $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathcal{V}$ be two subsets of variables, and let \dot{O}_1 be a GABP of $\mathcal{S}_{\mathcal{V}_1}$, then*

$$\dot{O}_1 \wedge \mathcal{U}_{\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}} = \dot{O}_1 \quad (7)$$

Proof. Consider morphisms ψ_1, ψ_2 from the product $\dot{O}_1 \wedge \mathcal{U}_{\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}}$ to factors \dot{O}_1 and $\mathcal{U}_{\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}}$ respectively. ψ_1 is a morphism between two GABPs of $\mathcal{S}_{\mathcal{V}_1}$, and one readily has $\psi_1(\dot{O}_1 \wedge \mathcal{U}_{\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}}) \sqsubseteq \dot{O}_1$ by definition of the product. We show that ψ_1 is actually an isomorphism. Let κ be a configuration of the product, then $\kappa_2 = \psi_2(\kappa)$ is obtained by restricting κ to events and conditions labeled by $\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}$ and removing all extra causality relations (only structural relations due to tiles of $\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}$ are preserved). Since $\mathcal{V}_1 \cap \mathcal{V}_2 \subseteq \mathcal{V}_1$, κ_2 can also be obtained by applying the same procedure to κ_1 . So let κ, κ' be two configurations of the product such that $\psi_1(\kappa) = \psi_1(\kappa') = \kappa_1$, then also $\psi_2(\kappa) = \psi_2(\kappa') = \kappa_2$, which proves $\kappa = \kappa_1 \wedge \kappa_2 = \kappa'$. So ψ_1 is injective. To check that ψ_1 is also surjective, select any κ_1 of \dot{O}_1 , and build configuration κ_2 of $\mathcal{U}_{\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}}$ by the above procedure. Then it is easily checked (by recursion in procedure 2) that $\psi_1(\kappa_1 \wedge \kappa_2) = \kappa_1$. So morphism ψ_1 is a bijective mapping. To prove that it is actually an isomorphism, it remains to check that no extra conflict or causality relation is lost by ψ_1 , which is obvious since in $\dot{O}_1 \wedge \mathcal{U}_{\mathcal{S}_{\mathcal{V}_1 \cap \mathcal{V}_2}}$ all such relations come from \dot{O}_1 . \square

As a consequence of this lemma, we define \mathbb{I} as the unfolding of the trivial system \mathcal{S}_{\emptyset} . \mathbb{I} is a trivial GABP, with no condition and no event, and satisfies

$$\forall \mathcal{V}' \subseteq \mathcal{V}, \quad \forall \dot{O} \text{ a GABP of } \mathcal{S}_{\mathcal{V}'}, \quad \dot{O} \wedge \mathbb{I} = \dot{O} \quad (8)$$

5 Projection of ABPs

5.1 Standard projection of BP : its weaknesses

Let \mathcal{S} be given by $\mathcal{S}_1 \parallel \dots \parallel \mathcal{S}_N$, where component \mathcal{S}_i operates of variables \mathcal{V}_i , and define the extended components $\mathcal{S}_i^e \triangleq \mathcal{S}_{\mathcal{V}_i}$. By equation (2), one has $\mathcal{U}_{\mathcal{S}} = \mathcal{U}_{\mathcal{S}_1^e} \wedge \dots \wedge \mathcal{U}_{\mathcal{S}_N^e}$. When dealing with large distributed systems, building $\mathcal{U}_{\mathcal{S}}$ may be of little practical interest : $\mathcal{U}_{\mathcal{S}}$ can be a huge object (even if restricted to a complete finite prefix, which we do not consider here), and one is not necessarily

interested in global behaviors of \mathcal{S} . On the contrary, a relevant question is “How does the behavior of, say, component \mathcal{S}_i changes once \mathcal{S}_i is inserted into \mathcal{S} ?” We have shown in [1] that the question amounts to determining $\Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$, the projection of $\mathcal{U}_{\mathcal{S}}$ on behaviors of \mathcal{S}_i^e . $\Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$ is obtained by restricting $\mathcal{U}_{\mathcal{S}}$ to conditions and events labeled by elements of \mathcal{S}_i^e , and trimming the result to merge isomorphic configurations. $\Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$ is a prefix of $\mathcal{U}_{\mathcal{S}_i^e}$, i.e. represents the restriction of $\mathcal{U}_{\mathcal{S}_i^e}$ to configurations that synchronize with configurations of the other components. Moreover, the collection $\Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$, $1 \leq i \leq N$, defines a minimal product covering of $\mathcal{U}_{\mathcal{S}}$, i.e. $\mathcal{U}_{\mathcal{S}} = \Pi_{\mathcal{V}_1}(\mathcal{U}_{\mathcal{S}}) \wedge \dots \wedge \Pi_{\mathcal{V}_N}(\mathcal{U}_{\mathcal{S}})$.

The central result of [1] is that minimal factors $\Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$ of $\mathcal{U}_{\mathcal{S}}$ can be obtained without computing $\mathcal{U}_{\mathcal{S}}$ itself, by means of modular computations. In the simple case of two components, these modular computations reduce to:

$$\Pi_{\mathcal{V}_1}(\mathcal{U}_{\mathcal{S}}) = \mathcal{U}_{\mathcal{S}_1^e} \wedge \Pi_{\mathcal{V}_1}(\mathcal{U}_{\mathcal{S}_2^e}) \quad (9)$$

(and symmetrically for $\Pi_{\mathcal{V}_2}(\mathcal{U}_{\mathcal{S}})$), where only “small” BPs appear on the RHS of (9), compared to $\mathcal{U}_{\mathcal{S}}$. This kind of relation holds on branching processes *provided projections are not “misleading”*, i.e. if every configuration of $\Pi_{\mathcal{V}_1}(\mathcal{U}_{\mathcal{S}_2^e})$ is the restriction of a configuration of $\mathcal{U}_{\mathcal{S}_2^e}$, and if causality relations between selected events are not lost (an example of a misleading projection is given in fig. 7). There exists a small family of distributed systems for which such a property is guaranteed, namely those interacting through a single variable only. The projection we define below for GABPs releases this limitation.

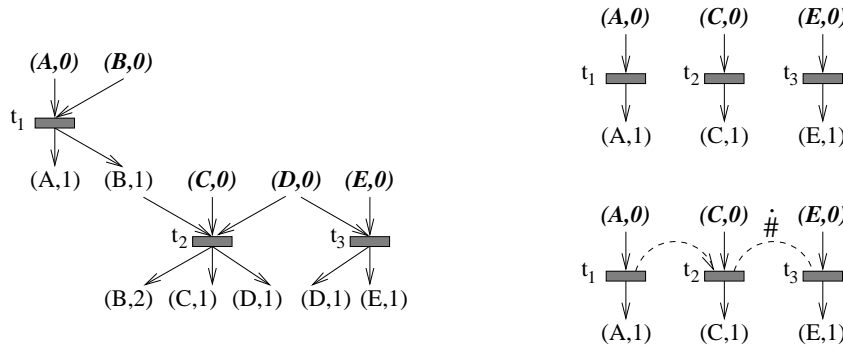


Figure 7: A branching process \mathcal{O} (left) and its standard projection $\Pi_{\{A,C,E\}}(\mathcal{O})$ (top, right). This projection is misleading: \mathbf{t}_2 and \mathbf{t}_3 appear as concurrent, but there is no configuration of \mathcal{O} containing both of them. Similarly, the causality $\mathbf{t}_1 \rightarrow \mathbf{t}_2$ is lost. In the category of ABPs, these features can be preserved (bottom, right).

5.2 Projection of ABPs

Let $\dot{\mathcal{O}}_1$ be a GABP of $\mathcal{S}_{|\mathcal{V}_1}$, and let $\mathcal{V}_2 \subseteq \mathcal{V}$. We define $\dot{\mathcal{O}}_{1|\mathcal{V}_2}$ as the *restriction* of $\dot{\mathcal{O}}_1$ to conditions labeled by $(\mathcal{V}_1 \cap \mathcal{V}_2) \times \mathcal{D}$ and to events labeled by $\mathcal{T}_1^e \cap \mathcal{T}_2^e = \{\mathbf{t} \in \mathcal{T} : \mathcal{V}_{\mathbf{t}} \cap \mathcal{V}_1 \cap \mathcal{V}_2 \neq \emptyset\}$. Observe that, although some conditions vanish, all conflict and causality relations on the remaining events are preserved. So every configuration of $\dot{\mathcal{O}}_{1|\mathcal{V}_2}$ is the restriction of a configuration of $\dot{\mathcal{O}}_1$. It is easily checked that $\dot{\mathcal{O}}_{1|\mathcal{V}_2}$ is a GABP of $\mathcal{S}_{|\mathcal{V}_1 \cap \mathcal{V}_2}$, but $\dot{\mathcal{O}}_{1|\mathcal{V}_2}$ is not trimmed in general, even if $\dot{\mathcal{O}}_1$ is an ABP. We define the *projection* of $\dot{\mathcal{O}}_1$ on variables \mathcal{V}_2 as

$$\dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1) \triangleq \text{Trim}(\dot{\mathcal{O}}_{1|\mathcal{V}_2}) \quad (10)$$

Observe that we make a distinction between $\Pi_{\mathcal{V}'}$, a family of projectors in the category of BP (defined in [1]), and $\dot{\Pi}_{\mathcal{V}'}$, a family of projectors in the category of ABPs.

There exists a canonical morphism between $\dot{\mathcal{O}}_1$ and its restriction $\dot{\mathcal{O}}_1|_{\mathcal{V}_2}$, and between a GABP and its trimmed version. Their composition defines the canonical morphism $\pi_{\mathcal{V}_2} : \dot{\mathcal{O}}_1 \rightarrow \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$ associated to this projection.

One could suspect the existence of a universal property defining $\dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$, looking like: for all ABP $\dot{\mathcal{O}}_2$ of $\mathcal{S}|_{\mathcal{V}_1 \cap \mathcal{V}_2}$ and morphism $\psi : \dot{\mathcal{O}}_1 \rightarrow \dot{\mathcal{O}}_2$, there exists a unique morphism $\phi : \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1) \rightarrow \dot{\mathcal{O}}_2$ such that $\psi = \phi \circ \pi_{\mathcal{V}_2}$. This property doesn't hold however, essentially because of the trimming.




Figure 8: A branching process \mathcal{O} (left) and its projection $\dot{\Pi}_{\{A,D\}}(\mathcal{O})$ (right). The standard projection $\Pi_{\{A,D\}}(\mathcal{O})$ defined in [1] would give a single occurrence of t_1 and of t_2 , these two events appearing as concurrent: $\Pi_{\{A,D\}}(\mathcal{O})$ would be misleading.

Lemma 8 Let $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathcal{V}$, then operators $\dot{\Pi}_{\mathcal{V}_i}$ defined above satisfy

$$\dot{\Pi}_{\mathcal{V}_1} \circ \dot{\Pi}_{\mathcal{V}_2} = \dot{\Pi}_{\mathcal{V}_1 \cap \mathcal{V}_2} \quad (11)$$

Proof. The result is obvious if one only considers restrictions: $(\dot{\mathcal{O}}|_{\mathcal{V}_1})|_{\mathcal{V}_2} = \dot{\mathcal{O}}|_{\mathcal{V}_1 \cap \mathcal{V}_2}$. Taking the trimming into account, one has

$$\text{Trim}(\dot{\mathcal{O}}|_{\mathcal{V}_1 \cap \mathcal{V}_2}) = \text{Trim}[(\dot{\mathcal{O}}|_{\mathcal{V}_1})|_{\mathcal{V}_2}] = \text{Trim}\{\text{Trim}(\dot{\mathcal{O}}|_{\mathcal{V}_1})|_{\mathcal{V}_2}\}$$

The last equality uses the relation $\dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}') = \dot{\Pi}_{\mathcal{V}_2}[\text{Trim}(\dot{\mathcal{O}}')]$ for a GABP $\dot{\mathcal{O}}'$, which can be obtained with proposition 1 for example. \square

6 ABP calculus on trees

This section studies the algebraic relations between the two operations defined on ABPs: product and projection. For a compound system \mathcal{S} , we show that these relations can be combined to compute projections of $\mathcal{U}_{\mathcal{S}}$ on components of \mathcal{S} by means of small size operations.

6.1 A key property

All results presented so far didn't make any assumption on the components of \mathcal{S} . From now, we consider systems built from components satisfying the

Structural assumption (SA): Let $\mathcal{S}_1, \dots, \mathcal{S}_N$ be coherent components, with $\mathcal{S}_i = (\mathcal{V}_i, \mathcal{T}_i, \mathbf{v}_i^0, \alpha_i, \beta_i, \gamma_i)$. For all i, j in $\{1, \dots, N\}$, $i \neq j$, if $\mathbf{t} \in \mathcal{T}_i \cap \mathcal{T}_j$ then $\mathcal{V}_{\mathbf{t}} \cap \mathcal{V}_i \cap \mathcal{V}_j \neq \emptyset$ where $\mathcal{V}_{\mathbf{t}}$ denotes variables impacted by \mathbf{t} in $\mathcal{S}_1 \parallel \dots \parallel \mathcal{S}_N$.

In words, if two components share some tile \mathbf{t} , then this tile has some impact on their shared variables, or, conversely, the variables shared by two components make all their shared tiles “visible.”

Observe that the SA is obviously satisfied when components \mathcal{S}_i have no common tile, which means that, in \mathcal{S} , all tiles have a limited span (each \mathcal{V}_t is necessarily contained in some \mathcal{V}_i), and interactions between components are only due to shared variables. This corresponds to the assumption made in [1]. But the SA is weaker than this particular case, and allows tiles covering several components. Notice also that the SA is preserved by extended components $\mathcal{S}_i^e = \mathcal{S}|_{\mathcal{V}_i}$, with $\mathcal{S} = \mathcal{S}_1 \parallel \dots \parallel \mathcal{S}_N$.

Proposition 4 *Let $\mathcal{S}_1^e, \mathcal{S}_2^e$ be restrictions of \mathcal{S} to $\mathcal{V}_1, \mathcal{V}_2$ resp., and assume these components satisfy the structural assumption. Let $\mathcal{O}_1, \mathcal{O}_2$ be GABPs of $\mathcal{S}_1^e, \mathcal{S}_2^e$ respectively, then*

$$\forall \mathcal{V}_3 \supseteq \mathcal{V}_1 \cap \mathcal{V}_2, \quad (\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)|_{\mathcal{V}_3} = \dot{\mathcal{O}}_1|_{\mathcal{V}_3} \wedge \dot{\mathcal{O}}_2|_{\mathcal{V}_3} \quad (12)$$

Proof. We first show that there exists an injective morphism from $\phi : (\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)|_{\mathcal{V}_3} \rightarrow \dot{\mathcal{O}}_1|_{\mathcal{V}_3} \wedge \dot{\mathcal{O}}_2|_{\mathcal{V}_3}$. Consider the diagram in fig. 9. Morphisms $\psi_i : \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2 \rightarrow \dot{\mathcal{O}}_i$ come from the definition of the product \wedge , and similarly for the $\phi_i : \dot{\mathcal{O}}_1|_{\mathcal{V}_3} \wedge \dot{\mathcal{O}}_2|_{\mathcal{V}_3} \rightarrow \dot{\mathcal{O}}_i|_{\mathcal{V}_3}$. The pair $(\psi_1, \psi_2) : \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2 \rightarrow \dot{\mathcal{O}}_1 \times \dot{\mathcal{O}}_2$ is injective, by definition of the product. This property remains true between restrictions $(\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)|_{\mathcal{V}_3}$ and $\dot{\mathcal{O}}_1|_{\mathcal{V}_3} \times \dot{\mathcal{O}}_2|_{\mathcal{V}_3}$. Since there exist morphisms ψ_i from $(\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)|_{\mathcal{V}_3}$ to the $\dot{\mathcal{O}}_i|_{\mathcal{V}_3}$, there exists also a unique morphism $\phi : (\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)|_{\mathcal{V}_3} \rightarrow \dot{\mathcal{O}}_1|_{\mathcal{V}_3} \wedge \dot{\mathcal{O}}_2|_{\mathcal{V}_3}$ which makes the diagram commutative (universal property of the categorical product). ϕ is necessarily injective, otherwise the pair (ψ_1, ψ_2) couldn't be injective on $(\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)|_{\mathcal{V}_3}$.




Figure 9: A commutative diagram of morphisms. Dashed arrows indicate restrictions of GABPs to events and conditions involving variables of \mathcal{V}_3 .

We now show that ϕ is also surjective. To do so, we prove that for any configuration $\kappa = \kappa_1 \wedge \kappa_2$ of $\dot{\mathcal{O}}_1|_{\mathcal{V}_3} \wedge \dot{\mathcal{O}}_2|_{\mathcal{V}_3}$ (with $\kappa_i = \phi_i(\kappa)$), one can find a configuration $\kappa' = \kappa'_1 \wedge \kappa'_2$ of $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2$, such that $\phi(\kappa'|_{\mathcal{V}_3}) = \kappa$. Let us select any configuration κ'_i of $\dot{\mathcal{O}}_i$ such that $\kappa'_i|_{\mathcal{V}_3} = \kappa_i$ (there always exist such configurations, because the restriction doesn't lose conflict relations), and consider $\kappa' = \kappa'_1 \wedge \kappa'_2$ in $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2$. One has $\phi(\kappa'|_{\mathcal{V}_3}) \sqsubseteq \kappa$. The point is thus to prove that equality holds.

In a construction of $\kappa = \kappa_1 \wedge \kappa_2$ based on procedure 2, let $<$ be the total order in which events are created. Let $<_i$ be the image of $<$ under ϕ_i ; $<_i$ defines a total order on events of κ_i which is a linear extension of \prec_i (the causality relation on κ_i). Notice that the pair $(<_1, <_2)$ can be used as a guide to build κ from κ_1 and κ_2 in procedure 2:

- a) Start with $\kappa = \mathbf{v}_1^0 \wedge \mathbf{v}_2^0$ and go over events of κ_1, κ_2 with indexes i_1, i_2 obeying $<_1, <_2$ (resp.).
- b) As long as e_{i_1} is a local event of κ_1 (i.e. doesn't correspond to a shared tile), build its image in κ and increase i_1 by 1.
- c) Then do the same for κ_2 .

- d) Events e_{i_1}, e_{i_2} of κ_1 and κ_2 are now shared events, associated to the same tile; so build the corresponding product event e of κ , and increase both indexes by 1. Then go back to b).

Let us now consider configurations κ'_i . There exists a total ordering $<'_i$ of events of κ'_i which coincides with $<_i$ on κ_i (because the restriction doesn't lose transitive causality relations). Let us build $\kappa' = \kappa'_1 \wedge \kappa'_2$, with $(<'_1, <'_2)$ as a guide. The structural assumption indicates that every event of $\kappa'_i \setminus \kappa_i$ is necessarily a local event. So all event synchronizations in the construction of κ' are due to events shared by κ_1 and κ_2 . Therefore the construction cannot block before all events of κ'_1 and κ'_2 have found a corresponding event in κ' . Since $\psi_i(\kappa') = \kappa'_i$, one has $\psi_i(\kappa'_{|\mathcal{V}_3}) = \kappa_i$, which proves $\phi(\kappa'_{|\mathcal{V}_3}) = \kappa$.

We have proved that morphism ϕ is a bijective mapping. To prove that ϕ is actually an isomorphism, it remains to prove that it preserves all extra conflict and extra causality relations, which is straightforward by definition of the product (left to the reader). \square

Isolating the first part of the proof, which doesn't make use of the structural assumption, one gets

Corollary 3 *Let $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$ be GABPs of $\mathcal{S}_{|\mathcal{V}_1}, \mathcal{S}_{|\mathcal{V}_2}$ resp., and let $\mathcal{V}_3 \subseteq \mathcal{V}$ be any variable set, then*

$$(\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)_{|\mathcal{V}_3} \sqsubseteq \dot{\mathcal{O}}_{1|\mathcal{V}_3} \wedge \dot{\mathcal{O}}_{2|\mathcal{V}_3} \quad (13)$$

Observe also that, given a pair $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$ of GABPs, conditions to get equality can be weakened: equality holds as far as \mathcal{V}_3 captures all events of $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_2$ that must synchronize, i.e. for a \mathcal{V}_3 satisfying

$$\forall \mathbf{t} \in \mathcal{T}_1^e \cap \mathcal{T}_2^e, [\exists e_1 \in \dot{\mathcal{O}}_1, \lambda_1(e_1) = \mathbf{t} \text{ or } \exists e_2 \in \dot{\mathcal{O}}_2, \lambda_2(e_2) = \mathbf{t}] \Rightarrow \mathcal{V}_t \cap \mathcal{V}_3 \neq \emptyset$$

Theorem 1 *Let $\mathcal{S}_1^e, \mathcal{S}_2^e$ be restrictions of \mathcal{S} to $\mathcal{V}_1, \mathcal{V}_2$ resp., and assume these components satisfy the structural assumption. Let $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$ be GABPs of $\mathcal{S}_1^e, \mathcal{S}_2^e$ resp., then*

$$\forall \mathcal{V}_3 \supseteq \mathcal{V}_1 \cap \mathcal{V}_2, \quad \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_1 \dot{\wedge} \dot{\mathcal{O}}_2) = \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_1) \dot{\wedge} \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_2) \quad (14)$$

Proof. Notice first that (14) uses $\dot{\wedge}$. The same equality with \wedge instead of $\dot{\wedge}$ doesn't hold in general, since the LHS term is necessarily trimmed whereas the RHS term may not be trimmed.

The proof relies on proposition 4. By trimming both sides of (12), one gets

$$\begin{aligned} \text{Trim}[(\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)_{|\mathcal{V}_3}] &= \text{Trim}(\dot{\mathcal{O}}_{1|\mathcal{V}_3} \wedge \dot{\mathcal{O}}_{2|\mathcal{V}_3}) \\ &= \text{Trim}[\text{Trim}(\dot{\mathcal{O}}_{1|\mathcal{V}_3}) \wedge \text{Trim}(\dot{\mathcal{O}}_{2|\mathcal{V}_3})] \\ &= \text{Trim}[\dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_1) \dot{\wedge} \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_2)] \\ &= \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_1) \dot{\wedge} \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_2) \end{aligned} \quad (15)$$

where the second equality uses proposition 3. In the proof of lemma 8, we have already shown that $\text{Trim}(\dot{\mathcal{O}}_{|\mathcal{V}_3}) = \text{Trim}[\text{Trim}(\dot{\mathcal{O}})_{|\mathcal{V}_3}]$ for any GABP $\dot{\mathcal{O}}$. Applying this property to the LHS term above yields

$$\begin{aligned} \text{Trim}[(\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)_{|\mathcal{V}_3}] &= \text{Trim}[\text{Trim}(\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2)_{|\mathcal{V}_3}] \\ &= \text{Trim}[(\dot{\mathcal{O}}_1 \dot{\wedge} \dot{\mathcal{O}}_2)_{|\mathcal{V}_3}] \\ &= \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_1 \dot{\wedge} \dot{\mathcal{O}}_2) \end{aligned}$$

which proves the theorem. \square

Corollary 4 *Under the conditions of corollary 3, i.e. without any assumption on \mathcal{S} nor on \mathcal{V}_3 , one has*

$$\dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_1 \dot{\wedge} \dot{\mathcal{O}}_2) \sqsubseteq \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_1) \dot{\wedge} \dot{\Pi}_{\mathcal{V}_3}(\dot{\mathcal{O}}_2) \quad (16)$$

6.2 Modular computations on tree-shaped systems

Let us come back to the problem outlined in the introduction of section 5. Given $\mathcal{S} = \mathcal{S}_1 \parallel \dots \parallel \mathcal{S}_N$, one has (theorem 1 in [1]) $\mathcal{U}_{\mathcal{S}} = \mathcal{U}_{\mathcal{S}_1^e} \wedge \dots \wedge \mathcal{U}_{\mathcal{S}_N^e} = \Pi_{\mathcal{V}_1}(\mathcal{U}_{\mathcal{S}}) \wedge \dots \wedge \Pi_{\mathcal{V}_N}(\mathcal{U}_{\mathcal{S}})$, where $\Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}}) \sqsubseteq \mathcal{U}_{\mathcal{S}_i^e}$ represents runs of \mathcal{S}_i^e that remain possible once this component is connected to all others. More generally, for \mathcal{O} a BP of \mathcal{S} defined as $\mathcal{O} = \mathcal{O}_1 \wedge \dots \wedge \mathcal{O}_N$, where \mathcal{O}_i is a BP of \mathcal{S}_i^e , one has $\mathcal{O} = \Pi_{\mathcal{V}_1}(\mathcal{O}) \wedge \dots \wedge \Pi_{\mathcal{V}_N}(\mathcal{O})$. And $\Pi_{\mathcal{V}_i}(\mathcal{O}) \sqsubseteq \mathcal{O}_i$ represents runs of the specification \mathcal{O}_i that are compatible with runs specified by the other components. Therefore, it is of great practical interest to determine these minimal factors $\Pi_{\mathcal{V}_i}(\mathcal{O})$ directly from the \mathcal{O}_i , without computing \mathcal{O} itself, which can be a huge object. We have shown in [1] that this was possible under two conditions: the first one concerns the structure of \mathcal{S} , and the second one requires that projectors Π be not misleading. In this section, we drop the quite restrictive condition on projectors, by moving to the larger category of ABPs. But before, let us establish connections between $\Pi_{\mathcal{V}_i}$ and $\dot{\Pi}_{\mathcal{V}_i}$.

Lemma 9 *Let \mathcal{O} be a standard branching process of \mathcal{S} , and define $\dot{\mathcal{O}}_i = \dot{\Pi}_{\mathcal{V}_i}(\mathcal{O}) = (C, E, \rightarrow, \lambda, \prec, \#)$, then $\mathcal{O}_i = \text{Trim}(C, E, \rightarrow, \lambda, \prec, \#) = \Pi_{\mathcal{V}_i}(\mathcal{O})$.*

Proof. By definition of $\dot{\Pi}_{\mathcal{V}_i}$, $\dot{\mathcal{O}}_i$ is obtained by restricting \mathcal{O} to nodes labeled by \mathcal{S}_i^e , and trimming the result. In the restriction, the flow relation \rightarrow of \mathcal{O} is reduced to relations coming from tiles of \mathcal{S}_i^e , but all causality and conflict relations that were present in \mathcal{O} are preserved, under the form of “extra” relations. By contrast, $\Pi_{\mathcal{V}_i}$ does the same restriction, but only preserves causality and conflict relations due to tiles of \mathcal{S}_i^e . So $\Pi_{\mathcal{V}_i}(\mathcal{O})$ is obtained as the union of configurations of $(C, E, \rightarrow, \lambda, \prec, \#)$. As a consequence, $\Pi_{\mathcal{V}_i}(\mathcal{O})$ and $\text{Trim}(C, E, \rightarrow, \lambda, \prec, \#)$ have the same configurations. Since the latter has no extra conflict relation, nor extra causality relation, $\text{Trim}(C, E, \rightarrow, \lambda, \prec, \#)$ is an ordinary BP, and is thus isomorphic to $\Pi_{\mathcal{V}_i}(\mathcal{O})$. \square

As a consequence of this lemma, the minimal factors $\Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$ (resp. $\Pi_{\mathcal{V}_i}(\mathcal{O})$) proposed in [1] can be derived from the $\dot{\Pi}_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$ (resp. $\dot{\Pi}_{\mathcal{V}_i}(\mathcal{O})$). The latter contain a richer information, since they preserve causality and conflict relations due to the other components of \mathcal{S} , so we concentrate on them. Notice however that the factorization result on $\mathcal{U}_{\mathcal{S}}$ is lost when replacing Π by $\dot{\Pi}$:

$$\mathcal{U}_{\mathcal{S}} \sqsubseteq \dot{\Pi}_{\mathcal{V}_1}(\mathcal{U}_{\mathcal{S}}) \dot{\wedge} \dots \dot{\wedge} \dot{\Pi}_{\mathcal{V}_N}(\mathcal{U}_{\mathcal{S}}) \quad (17)$$

This can be proved by checking that every configuration of $\mathcal{U}_{\mathcal{S}}$ can be reconstructed in the product, and concluding with lemma 4. But equality doesn't hold in general: there may remain extra causality relations in the RHS term, for example (see section 7.1).

Let us now gather properties of $\dot{\Pi}$ and $\dot{\wedge}$.

Proposition 5 *Let \mathcal{S} be a tile system, with \mathcal{V} as variable set, and denote by $\dot{\mathcal{O}}$ an augmented branching process of some restriction of \mathcal{S} .*

$$\forall \dot{\mathcal{O}}, \exists \mathcal{V}' \subseteq \mathcal{V}, \quad \dot{\mathcal{O}} = \dot{\Pi}_{\mathcal{V}'}(\dot{\mathcal{O}}) \quad (18)$$

$$\forall \mathcal{V}', \mathcal{V}'' \subseteq \mathcal{V}, \quad \dot{\Pi}_{\mathcal{V}'} \circ \dot{\Pi}_{\mathcal{V}''} = \dot{\Pi}_{\mathcal{V}' \cap \mathcal{V}''} \quad (19)$$

$$\forall \dot{\mathcal{O}}, \quad \dot{\mathcal{O}} \dot{\wedge} \mathbb{I} = \dot{\mathcal{O}} \quad \text{and} \quad \mathbb{I} = \dot{\Pi}_{\emptyset}(\mathbb{I}) \quad (20)$$

More importantly, let $\dot{\mathcal{O}}', \dot{\mathcal{O}}''$ be ABPs of $\mathcal{S}_{|\mathcal{V}'}$ and $\mathcal{S}_{|\mathcal{V}''}$ respectively, and assume these two systems satisfy the structural assumption, then

$$\forall \mathcal{W} \supseteq \mathcal{V}' \cap \mathcal{V}'', \quad \dot{\Pi}_{\mathcal{W}}(\dot{\mathcal{O}}' \dot{\wedge} \dot{\mathcal{O}}'') = \dot{\Pi}_{\mathcal{W}}(\dot{\mathcal{O}}') \dot{\wedge} \dot{\Pi}_{\mathcal{W}}(\dot{\mathcal{O}}'') \quad (21)$$

Proof. (18) expresses that $\dot{\mathcal{O}}$ is an ABP of $\mathcal{S}_{|\mathcal{V}'}$ for some \mathcal{V}' : the minimal \mathcal{V}' is determined by labels of $\min(\dot{\mathcal{O}})$. (19) comes from lemma 8. (20) expresses the definition of \mathbb{I} and property (8). (21) rephrases theorem 1. \square

These four properties form the core of an axiomatic framework that allows the development of modular algorithms to determine the $\dot{\Pi}_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$. This framework was studied in [20] and applied to our problem in [1], section 5.3, under the assumption of non-misleading projections. We briefly recall the main features of this procedure.

As a first step, we associate a *connection graph* to a system $\mathcal{S} = \mathcal{S}_1^e \parallel \dots \parallel \mathcal{S}_N^e$, where components $\mathcal{S}_i^e = \mathcal{S}_{|\mathcal{V}_i}$ satisfy the structural assumption. This graph \mathcal{G}^{cnx} has $\{1, 2, \dots, N\}$ as vertices, and (i, j) is an edge iff $\mathcal{V}_i \cap \mathcal{V}_j \neq \emptyset$, that is if components $\mathcal{S}_i^e, \mathcal{S}_j^e$ share variables, or interact (recall that there is no interaction outside shared variables, by the SA). A *communication graph* \mathcal{G}^c is derived from \mathcal{G}^{cnx} by recursively removing redundant edges, where edge (i, j) is redundant iff there exists a path i, k_1, \dots, k_L, j such that $\mathcal{V}_i \cap \mathcal{V}_j \subseteq \mathcal{V}_{k_l}$ for $1 \leq l \leq L$, *i.e.* if the interaction between \mathcal{S}_i^e and \mathcal{S}_j^e is also captured by another path. There generally exist several communication graphs for a given system, but if one of them is a tree, then all of them are trees. We say that \mathcal{S} is a *tree-shaped system*, or that \mathcal{S} *lives on a tree*.

For a tree-shaped system, projections $\dot{\Pi}_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$ can be determined by a message passing algorithm (MPA) based on a communication graph \mathcal{G}^c of \mathcal{S} . Each edge (i, j) of \mathcal{G}^c carries a message $\dot{\mathcal{M}}_{i,j}$ from component \mathcal{S}_i^e to component \mathcal{S}_j^e , and a message $\dot{\mathcal{M}}_{j,i}$ in the reverse direction. Messages are recursively updated until convergence:

Procedure 3

1. Initialization

$$\dot{\mathcal{M}}_{i,j} = \mathbb{I}, \quad \forall (i, j) \in \mathcal{G}^c \quad (22)$$

2. Until stability of messages, select an edge (i, j) and apply the update rule

$$\dot{\mathcal{M}}_{i,j} := \dot{\Pi}_{\mathcal{V}_i \cap \mathcal{V}_j}[\mathcal{U}_{\mathcal{S}_i^e} \hat{\wedge} (\hat{\wedge}_{k \in \mathcal{N}(i) \setminus j} \dot{\mathcal{M}}_{k,i})] \quad (23)$$

3. Termination

$$\dot{\mathcal{O}}'_i = \mathcal{U}_{\mathcal{S}_i^e} \hat{\wedge} (\hat{\wedge}_{k \in \mathcal{N}(i)} \dot{\mathcal{M}}_{k,i}), \quad 1 \leq i \leq N \quad (24)$$

In (23, 24), $\mathcal{N}(i)$ denotes neighbors of i in \mathcal{G}^c . The message $\dot{\mathcal{M}}_{i,j}$ is progressively refined until it has collected all information from systems lying in the subtree located beyond i from the standpoint of j . Procedure 3 is “chaotic,” in the sense that the exact ordering of updates is left unspecified. Nevertheless, it has been proved in [20] that this procedure converges in a finite number of useful updates (*i.e.* updates that effectively change the message), and that the $\dot{\mathcal{O}}'_i$ obtained at convergence correspond to the desired $\dot{\Pi}_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$. The same algorithms can be used to compute the $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$, where $\dot{\mathcal{O}}$ is a product GABP of \mathcal{S} : $\dot{\mathcal{O}} = \dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$. In that case, terms $\mathcal{U}_{\mathcal{S}_i^e}$ must be replaced by $\dot{\mathcal{O}}_i$.

Remark. Procedure 3 operates on ABPs: a trimming is performed at each step, in projections $\dot{\Pi}$ and in products $\hat{\wedge}$. This offers the advantage to minimize the size of objects, at the expense of a heavier computation load due to numerous trimmings. Relying on corollary 2, the trade-off between memory size and computation time can be positioned differently: trimmings need not be performed at each operation, or need not be complete. They can actually be delayed and placed at any step of the procedure.

7 Equivalence of ABPs and involutivity

Let \mathcal{O}_1 be a standard branching process of some restriction of \mathcal{S} , say $\mathcal{S}_{|\mathcal{V}_1}$, and let $\mathcal{V}_2 \subseteq \mathcal{V}$, the following *involutivity property* holds with the projection defined in [1]:

$$\mathcal{O}_1 \wedge \Pi_{\mathcal{V}_2}(\mathcal{O}_1) = \mathcal{O}_1 \quad (25)$$

which means that composing an BP with “part of itself” doesn’t change this BP. This strong property has an important role in the convergence of modular algorithms, for systems not living on a tree [20]. Unfortunately, it is only partly preserved in the case of ABPs.

7.1 Sub-involutivity of GABPs

Proposition 6 *Let $\dot{\mathcal{O}}_1$ be a GABP of $\mathcal{S}_{|\mathcal{V}_1}$, and let $\mathcal{V}_2 \subseteq \mathcal{V}$, then*

$$\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1) \sqsupseteq \dot{\mathcal{O}}_1 \quad (26)$$

The term “sub-involutivity” refers to the prefix relation replacing equality in (26).

Proof. We first show that $\dot{\mathcal{O}}_1$ is isomorphic to a prefix of $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_{1|\mathcal{V}_2}$. Let ψ_1, ψ_2 be the morphisms relating $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_{1|\mathcal{V}_2}$ to factors $\dot{\mathcal{O}}_1$ and $\dot{\mathcal{O}}_{1|\mathcal{V}_2}$ resp. Consider events e of $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_{1|\mathcal{V}_2}$ such that $\psi_1(e) = e_1$ and $\psi_2(e) = e_{1|\mathcal{V}_2}$ where $e_{1|\mathcal{V}_2}$ is the image of event e_1 in the restriction $\dot{\mathcal{O}}_{1|\mathcal{V}_2}$, if this image exists. The restriction of $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_{1|\mathcal{V}_2}$ to such events e (and the related pre- and post-conditions) forms a prefix $\dot{\mathcal{O}}$ of $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_{1|\mathcal{V}_2}$, and $\psi_1 : \dot{\mathcal{O}} \rightarrow \dot{\mathcal{O}}_1$ is an isomorphism.

Now, let us consider $\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$, with the associated morphisms ϕ_1, ϕ_2 to factors $\dot{\mathcal{O}}_1, \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$ resp. There exists a canonical morphism $\psi'_2 : \dot{\mathcal{O}}_{1|\mathcal{V}_2} \rightarrow \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$, so the GABP $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_{1|\mathcal{V}_2}$ is related to $\dot{\mathcal{O}}_1$ and $\dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$ by the pair of morphisms $(\psi_1, \psi'_2 \circ \psi_2)$. Hence, by the universal property of the product, there exists $\phi : \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_{1|\mathcal{V}_2} \rightarrow \dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$ that makes the diagram commutative, *i.e.* $\psi_1 = \phi_1 \circ \phi$ and $\psi'_2 \circ \psi_2 = \phi_2 \circ \phi$. As a consequence, $\phi(\dot{\mathcal{O}})$ is isomorphic to $\dot{\mathcal{O}}_1$ through ϕ_1 . \square

Notice that $\dot{\mathcal{O}}_1$ is generally a *strict* prefix of $\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$, or even of $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_1$. And this is not a matter of trimming, as illustrated in fig. 10: an ABP $\dot{\mathcal{O}}_1$ (top, left) and its projection $\dot{\Pi}_{\{A,B\}}(\dot{\mathcal{O}}_1)$ (top right) are represented. In order to avoid confusions, the figure gives names to events, assuming that names with the same index are labeled by the same tile (*e.g.* events e_2, e'_2, \bar{e}_2 both represent the firing of \mathbf{t}_2). The product $\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\{A,B\}}(\dot{\mathcal{O}}_1)$ is represented below, where events are identified by their image in each factor. The product is not trimmed, but even after trimming, it contains a configuration which has no isomorphic counterpart in $\dot{\mathcal{O}}_1$, namely the firing of \mathbf{t}_1 followed by \mathbf{t}_2 followed by \mathbf{t}_3 (thick arrows).

7.2 Pre-order on GABPs

Let $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$ be GABPs of $\mathcal{S}_{|\mathcal{V}_1}, \mathcal{S}_{|\mathcal{V}_2}$ resp., we define relation \sqsubseteq by

$$\dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_2 \Leftrightarrow \dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2 \quad (27)$$

which implies $\mathcal{V}_1 = \mathcal{V}_1 \cup \mathcal{V}_2$, *i.e.* $\mathcal{V}_2 \subseteq \mathcal{V}_1$. This relation expresses that $\dot{\mathcal{O}}_2$ doesn’t constrain behaviors of $\dot{\mathcal{O}}_1$, in the sense that all configurations of the latter are preserved. However, there may exist in $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2$ configurations that were not present in $\dot{\mathcal{O}}_1$: we show below that they are “reinforced” copies of configurations of $\dot{\mathcal{O}}_1$, where some concurrency has been turned into an extra causality relation.

Proposition 7 \sqsubseteq *is a pre-order on GABPs of \mathcal{S} . Moreover,*

$$\forall \dot{\mathcal{O}}, \quad \dot{\mathcal{O}} \sqsubseteq \mathbb{I} \quad (28)$$

$$\forall \dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2, \dot{\mathcal{O}}_3, \quad \dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_2 \Rightarrow \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_3 \sqsubseteq \dot{\mathcal{O}}_2 \wedge \dot{\mathcal{O}}_3 \quad (29)$$

$$\forall \dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2, \forall \mathcal{V}' \subseteq \mathcal{V} \quad \dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_2 \Rightarrow \dot{\Pi}_{\mathcal{V}'}(\dot{\mathcal{O}}_1) \sqsubseteq \dot{\Pi}_{\mathcal{V}'}(\dot{\mathcal{O}}_2) \quad (30)$$




Figure 10: An ABP \dot{O}_1 (top, left) and its projection $\dot{\Pi}_{\{A,B\}}(\dot{O}_1)$ (top right). \dot{O}_1 is a strict prefix of the product $\dot{O}_1 \wedge \dot{\Pi}_{\{A,B\}}(\dot{O}_1)$ (bottom), even after trimming.

Proof. $\dot{O}_1 \in \dot{O}_2 \in \dot{O}_1$ doesn't imply $\dot{O}_1 = \dot{O}_2$, as indicated by the counter-example of fig. 11, so \in can't be a partial order. But transitivity holds. Assume $\dot{O}_1 \in \dot{O}_2 \in \dot{O}_3$, we want to show $\dot{O}_1 \in \dot{O}_3$. From $\dot{O}_2 \sqsubseteq \dot{O}_2 \wedge \dot{O}_3$, one has $\dot{O}_1 \wedge \dot{O}_2 \sqsubseteq \dot{O}_1 \wedge \dot{O}_2 \wedge \dot{O}_3$ (property 1 section 4.4). Taking $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_2$ into account, one gets $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_2 \wedge \dot{O}_3$. Let \dot{O} be the prefix of $\dot{O}_1 \wedge \dot{O}_2 \wedge \dot{O}_3$ which is isomorphic to \dot{O}_1 , and denote by $\psi_1 : \dot{O}_1 \wedge \dot{O}_2 \wedge \dot{O}_3 \rightarrow \dot{O}_1$ the morphism relating the product to factor \dot{O}_1 . Then the restriction $\psi_1 : \dot{O} \rightarrow \dot{O}_1$ is an isomorphism. By the associativity of \wedge , there exists $\psi_{1,3} : \dot{O}_1 \wedge \dot{O}_2 \wedge \dot{O}_3 \rightarrow \dot{O}_1 \wedge \dot{O}_3$. Denote by ψ'_1 the morphism relating product $\dot{O}_1 \wedge \dot{O}_3$ to its factor \dot{O}_1 . Using the universal property of the product, one has $\psi_1(\dot{O}) = \psi'_1(\psi_{1,3}(\dot{O})) = \dot{O}_1$. Since $\psi_1 : \dot{O} \rightarrow \dot{O}_1$ is an isomorphism, one gets that $\psi_{1,3}(\dot{O}) \sqsubseteq \dot{O}_1 \wedge \dot{O}_3$ is isomorphic to \dot{O}_1 through ψ'_1 , so $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_3$, i.e. $\dot{O}_1 \in \dot{O}_3$.

(28) is obvious, by definition of \mathbb{I} .

For (29), observe that $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_2$ and $\dot{O}_3 \sqsubseteq \dot{O}_3 \wedge \dot{O}_3$ (proposition 4) imply $\dot{O}_1 \wedge \dot{O}_3 \sqsubseteq (\dot{O}_1 \wedge \dot{O}_3) \wedge (\dot{O}_2 \wedge \dot{O}_3)$ (property 1, section 4.4).

(30) requires more efforts. From $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_2$, one derives $\dot{O}_{1|\mathcal{V}'} \sqsubseteq (\dot{O}_1 \wedge \dot{O}_2)|_{\mathcal{V}'} \sqsubseteq \dot{O}_{1|\mathcal{V}'} \wedge \dot{O}_{2|\mathcal{V}'}$ where the last relation comes from corollary 3. Trimming extremal terms of this relation yields $\dot{\Pi}_{\mathcal{V}'}(\dot{O}_1) \sqsubseteq \text{Trim}(\dot{O}_{1|\mathcal{V}'} \wedge \dot{O}_{2|\mathcal{V}'})$ and, by proposition 3, $\dot{\Pi}_{\mathcal{V}'}(\dot{O}_1) \sqsubseteq \text{Trim}[\dot{\Pi}_{\mathcal{V}'}(\dot{O}_1) \wedge \dot{\Pi}_{\mathcal{V}'}(\dot{O}_2)]$. We conclude by $\text{Trim}[\dot{\Pi}_{\mathcal{V}'}(\dot{O}_1) \wedge \dot{\Pi}_{\mathcal{V}'}(\dot{O}_2)] \sqsubseteq \dot{\Pi}_{\mathcal{V}'}(\dot{O}_1) \wedge \dot{\Pi}_{\mathcal{V}'}(\dot{O}_2)$. \square




Figure 11: Two ABPs and their product: one has $\dot{O}_1 \in \dot{O}_2 \in \dot{O}_1$, but $\dot{O}_1 \neq \dot{O}_2$.

7.3 Equivalence of GABPs

Relying on proposition 7, one can define an equivalence relation on GABPs

$$\dot{O}_1 \equiv \dot{O}_2 \Leftrightarrow \dot{O}_1 \in \dot{O}_2 \in \dot{O}_1 \quad (31)$$

Proposition 8 *Classes of equivalent GABPs are stable under product \wedge . Equivalence is preserved by trimming. Classes of equivalent ABPs are stable under product $\dot{\wedge}$, under intersection and under union, and thus have a minimal and a maximal element for the prefix relation.*

Proof. Assume GABPs \dot{O}_1 and \dot{O}_2 are equivalent (so they are necessarily GABPs of the same system, i.e. operate on the same variables). From $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_2$, one gets $\dot{O}_1 \wedge \dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_1 \wedge \dot{O}_2$, and since $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_1$, one derives $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge (\dot{O}_1 \wedge \dot{O}_2)$, i.e. $\dot{O}_1 \in \dot{O}_1 \wedge \dot{O}_2$. From $\dot{O}_1 \sqsubseteq \dot{O}_1 \wedge \dot{O}_2$, one also gets $\dot{O}_1 \wedge \dot{O}_2 \sqsubseteq (\dot{O}_1 \wedge \dot{O}_2) \wedge \dot{O}_2$, i.e. $\dot{O}_1 \wedge \dot{O}_2 \in \dot{O}_2$. Gathering relations, $\dot{O}_1 \in \dot{O}_1 \wedge \dot{O}_2 \in \dot{O}_2$, and since $\dot{O}_2 \in \dot{O}_1$ we conclude $\dot{O}_1 \equiv \dot{O}_1 \wedge \dot{O}_2$.

Remark: We can say more on these equivalence classes. When $\dot{O}_1 \equiv \dot{O}_2$, as GABPs of the same system, \dot{O}_1 and \dot{O}_2 don't have local events. So, in the product $\dot{O}_1 \wedge \dot{O}_2$, every event of \dot{O}_1 has to synchronize with an event of \dot{O}_2 . As a consequence, let κ_1 be a configuration of \dot{O}_1 . Since κ_1 is isomorphic to some configuration κ of $\dot{O}_1 \wedge \dot{O}_2$, there exists a configuration $\kappa_2 = \psi_2(\kappa)$ in \dot{O}_2 which is identical to κ_1 , up to some extra causality relations that are turned into concurrency. We say that κ_2 is a *weakened* version of κ_1 , and conversely that κ_1 is a *reinforced* version of κ_2 .

Let \dot{O} be a GABP, and define the ABP $\dot{O}' = \text{Trim}(\dot{O})$. Let κ be a configuration of \dot{O} , there exists a configuration κ' in \dot{O}' which is isomorphic to κ . So $\dot{O} \sqsubseteq \dot{O} \wedge \dot{O}'$, i.e. $\dot{O} \in \dot{O}'$. The same arguments holds in the opposite direction, so $\dot{O} \equiv \dot{O}'$: equivalence is preserved by trimming.

Consider now the ABPs of an equivalence class. Combining the stability by product and by trimming, one gets that this sub-class is stable under product $\hat{\wedge}$.

From here, it is easier to consider ABPs as configuration sets, as evidenced by proposition 1 and the subsequent remark. We first prove stability by intersection. Let \dot{O}_1, \dot{O}_2 be finite equivalent ABPs. Since $\dot{O}_1 \cap \dot{O}_2 \sqsubseteq \dot{O}_1$ for example, one readily has $\dot{O}_1 \cap \dot{O}_2 \sqsubseteq (\dot{O}_1 \cap \dot{O}_2) \hat{\wedge} \dot{O}_1$. Conversely, let κ_1 be a configuration of \dot{O}_1 . From the above remark, there exists κ_2 in \dot{O}_2 which is a weakened version of κ_1 , so configuration $\kappa_1 \wedge \kappa_2$ in $\dot{O}_1 \hat{\wedge} \dot{O}_2$ is isomorphic to κ_1 (we can write $\kappa_1 \wedge \kappa_2 = \kappa_1$ since isomorphic ABPs are not distinguished). Similarly, there exists $\kappa_3 \in \dot{O}_1$ which is a weakened version of κ_2 , so $\kappa_1 \wedge \kappa_3 = \kappa_1$. The recursion finally builds a chain $\kappa_1, \kappa_2, \dots, \kappa_n, \dots$ of weaker and weaker configurations. Since \dot{O}_1 and \dot{O}_2 are finite, this chain stabilizes to a configuration $\kappa \in \dot{O}_1 \cap \dot{O}_2$ that still satisfies $\kappa_1 \wedge \kappa = \kappa_1$. This proves $\dot{O}_1 \sqsubseteq (\dot{O}_1 \cap \dot{O}_2) \hat{\wedge} \dot{O}_1$, and so $\dot{O}_1 \equiv \dot{O}_1 \cap \dot{O}_2$.

For the stability of a class by union, we use an auxiliary result. Let \dot{O} be an ABP, we build the configuration set K by selecting some configurations of \dot{O} , and reinforcing them. Let K' be obtained as the closure of K for the prefix relation, and let $\dot{O}' = \dot{O} \cup K'$. As a configuration set closed under prefix relation, \dot{O}' defines a valid ABP, and $\dot{O} \sqsubseteq \dot{O}'$. So one has $\dot{O} \sqsubseteq \dot{O} \hat{\wedge} \dot{O} \sqsubseteq \dot{O} \hat{\wedge} \dot{O}'$. In addition, every configuration κ' of K' is isomorphic to a configuration of $\dot{O} \hat{\wedge} \dot{O}'$: there exists $\kappa \in \dot{O}$ which is weaker than κ' , so $\kappa' = \kappa \wedge \kappa'$. This proves $\dot{O}' \sqsubseteq \dot{O} \hat{\wedge} \dot{O}'$, whence $\dot{O} \equiv \dot{O}'$: by incorporating reinforced configurations into an ABP, one stays in the same class. Now, let \dot{O}_1, \dot{O}_2 be finite equivalent ABPs, and take $\dot{O} = \dot{O}_1 \cap \dot{O}_2$. Since $\dot{O} \equiv \dot{O}_i$, every configuration κ_i of $\dot{O}_i \setminus \dot{O}$ is a reinforced version of some configuration κ of \dot{O} . So $\dot{O}_1 \cup \dot{O}_2$ can be obtained by incorporating reinforced configurations to \dot{O} , which proves $\dot{O}_1 \cup \dot{O}_2 \equiv \dot{O} \equiv \dot{O}_1 \equiv \dot{O}_2$.

To extend these stability results to infinite ABPs, observe first that, inside a class, all ABPs have the same height. And by truncating ABPs of a given class to their prefix of height n , one gets another equivalence class, which is finite. We conclude by letting n go to infinity. \square

Notice that the maximal element of a finite equivalence class can either be obtained as the union or as the product of all elements of this class. Also, the proof of the proposition reveals that the ABPs of a class only differ by the way they reinforce some configurations of the minimal element of that class. Therefore, if $\dot{O}_1 \equiv \dot{O}_2$, with $\dot{O}_i = (C_i, E_i, \rightarrow_i, \lambda_i, \prec_i, \#_i)$, $i = 1, 2$, define $\mathcal{O}_i = \text{Trim}(C_i, E_i, \rightarrow_i, \lambda_i, \prec_i, \#_i)$, where all extra causality relations are removed. Then the \mathcal{O}_i form ordinary branching processes, and $\mathcal{O}_1 = \mathcal{O}_2$.

7.4 Product, projection and equivalence

Proposition 9 *Let $\dot{O}_1 \equiv \dot{O}'_1$ and $\dot{O}_2 \equiv \dot{O}'_2$ be four ABPs, then $\dot{O}_1 \hat{\wedge} \dot{O}_2 \equiv \dot{O}'_1 \hat{\wedge} \dot{O}'_2$ and $\dot{\Pi}_{\nu_2}(\dot{O}_1) \equiv \dot{\Pi}_{\nu_2}(\dot{O}'_1)$, i.e. equivalence is preserved by product $\hat{\wedge}$ and by projection.*

Proof. It is enough to prove $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2 \equiv \dot{\mathcal{O}}'_1 \wedge \dot{\mathcal{O}}'_2$, since equivalence is preserved under trimming. Let $\kappa' = \kappa'_1 \wedge \kappa'_2$ be a configuration of $\dot{\mathcal{O}}'_1 \wedge \dot{\mathcal{O}}'_2$. There exists $\kappa_i \in \dot{\mathcal{O}}_i$ which is a weakened version of κ'_i , $i = 1, 2$. So $\kappa = \kappa_1 \wedge \kappa_2 \in \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2$ is a weakened version of κ' , whence $\kappa' \wedge \kappa = \kappa'$. This proves $(\dot{\mathcal{O}}'_1 \wedge \dot{\mathcal{O}}'_2) \sqsubseteq (\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2) \wedge (\dot{\mathcal{O}}'_1 \wedge \dot{\mathcal{O}}'_2)$, i.e. $\dot{\mathcal{O}}'_1 \wedge \dot{\mathcal{O}}'_2 \in \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2$. By symmetry, one has also $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2 \in \dot{\mathcal{O}}'_1 \wedge \dot{\mathcal{O}}'_2$, whence the equivalence. The proof of the second relation of the lemma is similar. \square

Given an ABP $\dot{\mathcal{O}}$, let us denote by $\overline{\dot{\mathcal{O}}}$ the maximal element of its class, and by $\underline{\dot{\mathcal{O}}}$ the minimal element. $\overline{\dot{\mathcal{O}}}$ is a huge object formed by all possible reinforcements of configurations of $\dot{\mathcal{O}}$. Given two ABPs $\dot{\mathcal{O}}_1, \dot{\mathcal{O}}_2$, one has $\overline{\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2} = \overline{\dot{\mathcal{O}}_1} \wedge \overline{\dot{\mathcal{O}}_2}$ and $\dot{\Pi}_{\mathcal{V}_2}(\overline{\dot{\mathcal{O}}_1}) = \overline{\dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)}$, but these properties have little interest. Properties of minimal elements are potentially much more useful.




Figure 12: $\underline{\dot{\mathcal{O}}_1}, \underline{\dot{\mathcal{O}}_2}$ are minimal ABPs, but their product is not minimal.

By proposition 9, one has $\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2 \equiv \underline{\dot{\mathcal{O}}_1} \wedge \underline{\dot{\mathcal{O}}_2}$, but unfortunately the minimality is not stable under product: $\underline{\dot{\mathcal{O}}_1} \wedge \underline{\dot{\mathcal{O}}_2} = \underline{\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_2}$ doesn't hold in general, as illustrated by figure 12. In the same way, one has $\dot{\Pi}_{\mathcal{V}_2}(\underline{\dot{\mathcal{O}}_1}) \equiv \underline{\dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)}$ by proposition 9, but unfortunately again the minimality is not stable under projection: Figure 13 gives a counter-example where $\dot{\Pi}_{\mathcal{V}_2}(\underline{\dot{\mathcal{O}}_1}) \neq \underline{\dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)}$.




Figure 13: A minimal ABP $\underline{\dot{\mathcal{O}}}$ on the left-hand side, and its projection $\underline{\dot{\Pi}_{\{A,B\}}(\dot{\mathcal{O}})}$ on the right-hand side. The latter is not minimal.

Despite these drawbacks, the parsimony of minimal ABPs is a desirable feature. Indeed, the latter offer the advantage to discard extra causality relations as soon as they do not lead to different futures. Therefore, it is tempting to consider that a minimization operation should also be incorporated to \wedge and $\dot{\Pi}$, just like the trimming. This presents no theoretical difficulty; however, we draw the reader's attention to the fact that, by contrast with the trimming operation, computing a minimal representative $\underline{\dot{\mathcal{O}}}$ from $\dot{\mathcal{O}}$ cannot be done recursively. To decide whether some given extra causality link $e \dot{\prec} e'$ can be removed, one must ensure that every configuration κ containing this link is isomorphic to another κ' which doesn't contain this link (isomorphism is meant "up to the relation $e \dot{\prec} e'$ " and its proper consequences, which vanish). And this test requires information arbitrarily far away in the future of events e and e' . This is suggested by figure 14.




Figure 14: The first two ABPs are equivalent, but are not equivalent to the last one. This shows that the position of t_3 is relevant to decide whether the extra relation $t_1 \dot{\prec} t_2$ can be removed or not in the computation of the minimal ABP of a class. This computation may thus need events arbitrarily far away in the future.

7.5 Involutivity of minimal ABPs

Proposition 10 Let $\dot{\mathcal{O}}_1$ be a GABP of $\mathcal{S}_{|\mathcal{V}_1}$, and let $\mathcal{V}_2 \subseteq \mathcal{V}$, then

$$\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1) \equiv \dot{\mathcal{O}}_1 \quad (32)$$

Proof. We already know $\dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$ from (26), whence $\dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_1 \wedge [\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)]$, so $\dot{\mathcal{O}}_1 \in \dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)$. For the converse part, one has $\dot{\mathcal{O}}_1 \sqsubseteq \dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}_1$, so $\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1) \sqsubseteq \dot{\mathcal{O}}_1 \wedge [\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)]$, i.e. $\dot{\mathcal{O}}_1 \wedge \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1) \in \dot{\mathcal{O}}_1$. \square

If $\dot{\mathcal{O}}_1$ is an ABP, this relation becomes $\dot{\mathcal{O}}_1 \dot{\wedge} \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1) \equiv \dot{\mathcal{O}}_1$. This proposition expresses that involutivity of ABPs is recovered if equality is replaced by equivalence, to which we refer as *weak involutivity* in the sequel. Further, since

$$\underline{\dot{\mathcal{O}}_1} = \underline{\underline{\dot{\mathcal{O}}_1 \dot{\wedge} \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)}} \triangleq \underline{\underline{\dot{\mathcal{O}}_1 \dot{\Delta} \dot{\Pi}_{\mathcal{V}_2}(\dot{\mathcal{O}}_1)}} \quad (33)$$

the involutivity in the “strict sense” is recovered, provided computations are projected to the subcategory of *minimal* ABPs.

7.6 Minimal product covering of an ABP

For a branching process \mathcal{O} of \mathcal{S} , and without any assumption but $\mathcal{V} = \mathcal{V}_1 \cup \dots \cup \mathcal{V}_N$, one has

$$\mathcal{O} \sqsubseteq \Pi_{\mathcal{V}_1}(\mathcal{O}) \wedge \dots \wedge \Pi_{\mathcal{V}_N}(\mathcal{O}) \quad (34)$$

The RHS term corresponds to the *minimal product covering* of \mathcal{O} by branching processes of the $\mathcal{S}_{|\mathcal{V}_i}$: every other product covering of that kind for \mathcal{O} , say $\mathcal{O} \sqsubseteq \mathcal{O}_1 \wedge \dots \wedge \mathcal{O}_N$, satisfies $\Pi_{\mathcal{V}_i}(\mathcal{O}) \sqsubseteq \mathcal{O}_i$. Moreover, if \mathcal{O} is already a product BP of \mathcal{S} , $\mathcal{O} = \mathcal{O}_1 \wedge \dots \wedge \mathcal{O}_N$ with \mathcal{O}_i a BP of $\mathcal{S}_{|\mathcal{V}_i}$, then equality holds: $\mathcal{O} = \Pi_{\mathcal{V}_1}(\mathcal{O}) \wedge \dots \wedge \Pi_{\mathcal{V}_N}(\mathcal{O})$.

Unfortunately, this minimality property is lost for augmented branching processes. Let $\dot{\mathcal{O}}$ be an ABP of \mathcal{S} , then $\dot{\mathcal{O}} \sqsubseteq \dot{\Pi}_{\mathcal{V}_1}(\dot{\mathcal{O}}) \dot{\wedge} \dots \dot{\wedge} \dot{\Pi}_{\mathcal{V}_N}(\dot{\mathcal{O}})$ still holds (by proposition 1). However, by taking a strict prefix of some of these factors, one may preserve this relation⁷. Moreover, even if no prefix can be taken in any factor without breaking the covering property, unicity of the factors is still not granted. Nevertheless, a weaker notion of minimality can be defined on product coverings of an ABP.

Lemma 10 Let $\dot{\mathcal{O}}$ be an ABP of \mathcal{S} , and assume $\mathcal{V} = \mathcal{V}_1 \cup \dots \cup \mathcal{V}_N$. Projections $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$ define a product covering of $\dot{\mathcal{O}}$ which is minimal for \in : if $\dot{\mathcal{O}} \sqsubseteq \dot{\wedge}_i \dot{\mathcal{O}}_i$ is another product covering of $\dot{\mathcal{O}}$, then $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \in \dot{\mathcal{O}}_i$.

⁷In this section, we let the reader build his own examples and counter-examples, taking inspiration from figures 10 or 11 for ex.

Proof. Let $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$ and $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}}'_1 \wedge \dots \wedge \dot{\mathcal{O}}'_N$ be two product coverings of $\dot{\mathcal{O}}$. As $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}} \wedge \dot{\mathcal{O}}$, one has $\dot{\mathcal{O}} \sqsubseteq (\dot{\mathcal{O}}_1 \wedge \dot{\mathcal{O}}'_1) \wedge \dots \wedge (\dot{\mathcal{O}}_N \wedge \dot{\mathcal{O}}'_N)$, which gives another product covering of $\dot{\mathcal{O}}$. Notice that each new factor lives in an inferior equivalence class of \equiv since $(\dot{\mathcal{O}}_i \wedge \dot{\mathcal{O}}'_i) \in \dot{\mathcal{O}}_i, \dot{\mathcal{O}}'_i$. Therefore, by taking a product of appropriately chosen product coverings⁸ of $\dot{\mathcal{O}}$, one gets another product covering of $\dot{\mathcal{O}}$ where factors $\dot{\mathcal{O}}_i$ lie in *minimal* equivalence classes, for the partial order \in . If one considers the product of *all* product coverings of $\dot{\mathcal{O}}$, these factors $\dot{\mathcal{O}}_i$ result in the maximal elements of these minimal classes ($\dot{\mathcal{O}}_i = \overline{\dot{\mathcal{O}}_i}$).

Let us denote by $\dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$ a minimal product covering of $\dot{\mathcal{O}}$ in the sense of equivalence classes: if $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}}'_1 \wedge \dots \wedge \dot{\mathcal{O}}'_N$, then $\dot{\mathcal{O}}_i \in \dot{\mathcal{O}}'_i$. How does it relate to $\dot{\mathcal{O}} \sqsubseteq \dot{\Pi}_{\mathcal{V}_1}(\dot{\mathcal{O}}) \wedge \dots \wedge \dot{\Pi}_{\mathcal{V}_N}(\dot{\mathcal{O}})$? By definition, $\dot{\mathcal{O}}_i \in \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$. Conversely, from $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$ one has $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \sqsubseteq \dot{\mathcal{O}}_i \wedge \dot{\mathcal{O}}''_i$ with $\dot{\mathcal{O}}''_i = \dot{\Pi}_{\mathcal{V}_i}(\wedge_{j \neq i} \dot{\mathcal{O}}_j)$ (see cor. 4). So

$$\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \sqsubseteq \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \wedge \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \sqsubseteq \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \wedge \dot{\mathcal{O}}_i \wedge \dot{\mathcal{O}}''_i$$

from which one gets $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \in \dot{\mathcal{O}}_i \wedge \dot{\mathcal{O}}''_i$. And since $\dot{\mathcal{O}}_i \wedge \dot{\mathcal{O}}''_i \in \dot{\mathcal{O}}_i \wedge \mathbb{I} = \dot{\mathcal{O}}_i$ (prop. 7), we derive $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \in \dot{\mathcal{O}}_i$. So finally $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \equiv \dot{\mathcal{O}}_i$ and factors $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$ define another product covering of $\dot{\mathcal{O}}$ which is minimal in the sense of equivalence classes. (Notice however that the covering property is generally lost if some factor is replaced by another member of its class.) \square

Product covering of ABPs apparently have weaker properties than product coverings of standard branching processes. But, as a consequence of the previous lemma, they recover a classical behavior when expressed on equivalence classes of ABPs, *i.e.* when the prefix relation \sqsubseteq is weakened into \in .

Lemma 11 *Let $\dot{\mathcal{O}}$ be an ABP of \mathcal{S} , and assume $\mathcal{V} = \mathcal{V}_1 \cup \dots \cup \mathcal{V}_N$. Projections $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$ define the unique minimal product covering of (the class of) $\dot{\mathcal{O}}$ in the sense of equivalence classes: $\dot{\mathcal{O}} \in \wedge_i \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$, and if $\dot{\mathcal{O}} \in \wedge_i \dot{\mathcal{O}}_i$ is another product covering, then $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \in \dot{\mathcal{O}}_i$. Moreover, if $\dot{\mathcal{O}}$ belongs to a product class, *i.e.* $\dot{\mathcal{O}} \equiv \wedge_i \dot{\mathcal{O}}_i$, then $\dot{\mathcal{O}} \equiv \wedge_i \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$.*

Proof. Since $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}}'$ implies $\dot{\mathcal{O}} \in \dot{\mathcal{O}}'$, the proof of lemma 10 applies, with \in replacing \sqsubseteq . For the last part of the lemma, assume $\dot{\mathcal{O}}$ is a product class $\dot{\mathcal{O}} \equiv \wedge_i \dot{\mathcal{O}}_i$. From the minimal product covering property of projections, one has $\dot{\mathcal{O}} \in \wedge_i \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \in \wedge_i \dot{\mathcal{O}}_i \equiv \dot{\mathcal{O}}$, so $\dot{\mathcal{O}} \equiv \wedge_i \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$. \square

This result applies of course to a product ABP $\dot{\mathcal{O}} = \wedge_i \dot{\mathcal{O}}_i$. Specializing this lemma to minimal ABPs, one has $\underline{\dot{\mathcal{O}}} \in \underline{\dot{\Pi}_{\mathcal{V}_1}(\dot{\mathcal{O}})} \underline{\wedge} \dots \underline{\wedge} \underline{\dot{\Pi}_{\mathcal{V}_N}(\dot{\mathcal{O}})}$ in general, and

$$\underline{\dot{\mathcal{O}}} = \underline{\dot{\Pi}_{\mathcal{V}_1}(\dot{\mathcal{O}})} \underline{\wedge} \dots \underline{\wedge} \underline{\dot{\Pi}_{\mathcal{V}_N}(\dot{\mathcal{O}})} \quad (35)$$

when $\dot{\mathcal{O}}$ belongs to a product class.

Notice however that the stronger relation $\underline{\dot{\mathcal{O}}} \sqsubseteq \underline{\dot{\Pi}_{\mathcal{V}_1}(\dot{\mathcal{O}})} \wedge \dots \wedge \underline{\dot{\Pi}_{\mathcal{V}_N}(\dot{\mathcal{O}})}$ is tempting, but doesn't hold in general, essentially because $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}}' \not\Rightarrow \underline{\dot{\mathcal{O}}} \sqsubseteq \underline{\dot{\mathcal{O}}}'$ (or in other words because \in is a strictly weaker relation than \sqsubseteq).

8 Local computations on systems with cycles

8.1 Summary of results obtained on trees

At this point, it is worth gathering results that are accessible by ABP calculus. Let $\mathcal{S} = \mathcal{S}_1 \parallel \dots \parallel \mathcal{S}_N$ be a compound system satisfying the structural assumption and living on a tree.

The unfolding $\mathcal{U}_{\mathcal{S}}$ satisfies $\mathcal{U}_{\mathcal{S}} = \mathcal{U}_{\mathcal{S}_1^e} \wedge \dots \wedge \mathcal{U}_{\mathcal{S}_N^e}$. Procedure 3 initialized with ordinary BPs $\dot{\mathcal{O}}_i = \mathcal{U}_{\mathcal{S}_i^e}$ yields the ABPs $\dot{\mathcal{O}}_i = \dot{\Pi}_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}})$, from which standard projections $\mathcal{O}'_i = \Pi_{\mathcal{V}_i}(\mathcal{U}_{\mathcal{S}}) \sqsubseteq \mathcal{U}_{\mathcal{S}_i^e}$ can

⁸By the usual argument, this is easy to do on finite height ABPs, and we extend it to infinite ABPs by letting the height go to infinity.

be derived (lemma 9). The \mathcal{O}'_i correspond to the minimal product covering of \mathcal{U}_S by BPs of the \mathcal{S}_i^e : $\mathcal{U}_S \sqsubseteq \mathcal{O}'_1 \wedge \dots \wedge \mathcal{O}'_N$. And since $\mathcal{O}'_1 \wedge \dots \wedge \mathcal{O}'_N \sqsubseteq \mathcal{U}_S$, one has $\mathcal{U}_S = \mathcal{O}'_1 \wedge \dots \wedge \mathcal{O}'_N$.

More generally, consider a product branching process of \mathcal{S} : $\mathcal{O} = \mathcal{O}_1 \wedge \dots \wedge \mathcal{O}_N$ where \mathcal{O}_i is a BP of \mathcal{S}_i^e for all i . Initialized with the \mathcal{O}_i , procedure 3 yields the ABPs $\dot{\mathcal{O}}'_i = \dot{\Pi}_{\mathcal{V}_i}(\mathcal{O})$, whence the $\mathcal{O}'_i = \Pi_{\mathcal{V}_i}(\mathcal{O})$ by lemma 9. And $\mathcal{O} \sqsubseteq \mathcal{O}'_1 \wedge \dots \wedge \mathcal{O}'_N$ is the minimal product covering of \mathcal{O} by BPs of the \mathcal{S}_i^e . But since $\mathcal{O}'_i \sqsubseteq \mathcal{O}_i$ and \sqsubseteq is preserved by \wedge , one has $\mathcal{O} \sqsubseteq \mathcal{O}'_1 \wedge \dots \wedge \mathcal{O}'_N \sqsubseteq \mathcal{O}_1 \wedge \dots \wedge \mathcal{O}_N = \mathcal{O}$, whence again $\mathcal{O} = \mathcal{O}'_1 \wedge \dots \wedge \mathcal{O}'_N$. Being able to compute efficiently this minimal factorization of an ordinary BP of \mathcal{S} was one of the objectives of the present paper.

Finally, in the most general case, let $\dot{\mathcal{O}} = \dot{\mathcal{O}}_1 \dot{\wedge} \dots \dot{\wedge} \dot{\mathcal{O}}_N$ be a product ABP of \mathcal{S} , where $\dot{\mathcal{O}}_i$ is an ABP of \mathcal{S}_i^e for all i . Procedure 3 initialized with the $\dot{\mathcal{O}}_i$ yields the $\dot{\mathcal{O}}'_i = \dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$, and one readily has $\dot{\mathcal{O}} \sqsubseteq \dot{\mathcal{O}}'_1 \dot{\wedge} \dots \dot{\wedge} \dot{\mathcal{O}}'_N$. But equality is not granted: $\dot{\mathcal{O}}'_i \sqsubseteq \dot{\mathcal{O}}_i$ may not hold. However, we can replace relations \sqsubseteq and $=$ by their weaker versions \in and \equiv . One has $\dot{\mathcal{O}}'_i \in \dot{\mathcal{O}}_i$, so $\dot{\mathcal{O}} \sqsubseteq \dot{\wedge}_i \dot{\mathcal{O}}'_i \in \dot{\wedge}_i \dot{\mathcal{O}}_i = \dot{\mathcal{O}}$ from which one gets $\dot{\mathcal{O}} \equiv \dot{\wedge}_i \dot{\mathcal{O}}'_i$ and $\underline{\dot{\mathcal{O}}} = \underline{\dot{\wedge}_i \dot{\mathcal{O}}'_i}$. This last result of course applies to the previous cases: if $\dot{\mathcal{O}}$ is actually a BP \mathcal{O} , one has $\mathcal{O} = \underline{\dot{\wedge}_i \dot{\mathcal{O}}'_i}$.

8.2 Weak convergence for cyclic systems

In general, a given compound system $\mathcal{S} = \mathcal{S}_1 \parallel \dots \parallel \mathcal{S}_N$ does not live on a tree. But this does not mean that modular computations should rashly be removed from the tool box. To obtain minimal factors of \mathcal{U}_S (or more generally projections of $\dot{\mathcal{O}} = \dot{\mathcal{O}}_1 \wedge \dots \wedge \dot{\mathcal{O}}_N$) at a reasonable cost, a first strategy consists in aggregating some components \mathcal{S}_i in order to obtain a tree-shaped system with larger components (the structural assumption is preserved by aggregation). However, this doesn't save all situations.

Another track consists in applying procedure 3 to a communication graph \mathcal{G}^c of \mathcal{S} , *despite the presence of cycles*. This idea was explored with great success for the approximate decoding of some error correcting codes, under the name of “turbo⁹ decoding.” Natural questions are: 1/ Does it converge? 2/ What are the properties of the limit? 3/ How does it relate to the desired minimal factors? There exist useful answers to these three questions.

Proposition 11 *For a system \mathcal{S} satisfying the structural assumption but not living on a tree, procedure 3 run on a communication graph \mathcal{G}^c of \mathcal{S} converges in the sense of equivalence classes. Limit classes do not depend on the ordering of updates in procedure 3. They do not depend on the choice of the communication graph either.*

Moreover, the ABPs $\dot{\mathcal{O}}'_i$ computed at each step of procedure 3 form a decreasing sequence for \in which is lower bounded by projections $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}})$, i.e. $\dot{\Pi}_{\mathcal{V}_i}(\dot{\mathcal{O}}) \in \dot{\mathcal{O}}'_i$. And relation $\dot{\mathcal{O}} \equiv \dot{\wedge}_i \dot{\mathcal{O}}'_i$ is preserved at each step.

Proof. First of all, the properties of relation \in mentioned in proposition 7 ensure that messages $\dot{\mathcal{M}}_{i,j}$ are decreasing for \in in a run of procedure 3 (lemma 7 in [20]): $\dot{\mathcal{M}}_{i,j}^{k+1} \in \dot{\mathcal{M}}_{i,j}^k \in \dots \in \dot{\mathcal{M}}_{i,j}^0 = \mathbb{I}$, where $\dot{\mathcal{M}}_{i,j}^k$ denotes message $\dot{\mathcal{M}}_{i,j}$ at its k -th update. If procedure 3 is initialized with ABPs $\dot{\mathcal{O}}_i$ of height lower than h , the height of all messages remains lower than h . There is a finite number of such ABPs, for every restriction of \mathcal{S} , so the sequence of messages $\dot{\mathcal{M}}_{i,j}^k$ cannot be infinitely strictly decreasing, and thus stabilizes in a equivalence class of \equiv after a finite number of updates¹⁰. Moreover, this equivalence class is the same, for all runs of procedure 3 (theorem 3 in [20]). It is not known however whether the procedure stabilizes inside these equivalence classes, or if there exist limit cycles. We conjecture that convergence takes place, although different elements of the same class may be

⁹The adjective “turbo” stresses the fact that when a MPA is run on a graph with loops, the information sent by some system will eventually come back to it.

¹⁰Again, the result extends to infinite ABPs by proving convergence for higher and higher prefixes. This reveals also that procedure 3 converges progressively: stability of classes for the prefix of height h is reached before stability for the prefix of height $h + 1$.

reached for different runs of procedure 3. The invariance of the limit by graph changes relies on the weak involutivity property, and is proved in lemma 11 of [20].

Since messages are decreasing for \sqsubseteq , the ABPs \dot{O}'_i computed by (24) at each step (instead of at the end only) also form a decreasing sequence, for all i . To prove that $\dot{O} \equiv \dot{\wedge}_i \dot{O}'_i$ is preserved at each step, notice that this relation is true at initialization since $\dot{O}'_i = \dot{O}_i$. And it is easily checked that the update equation preserves this relation (see also theorem 8 in [20]). Finally, $\dot{\Pi}_{\mathcal{V}_i}(\dot{O}) \equiv \dot{\Pi}_{\mathcal{V}_i}(\dot{\wedge}_j \dot{O}'_j) = \dot{O}'_i \dot{\wedge} \dot{\Pi}_{\mathcal{V}_i}(\dot{\wedge}_{j \neq i} \dot{O}'_j) \sqsubseteq \dot{O}'_i$. \square

Observe that the equivalence classes of $\dot{\Pi}_{\mathcal{V}_i}(\dot{O})$ may not be reached by procedure 3, which generally stops before this lower bound, *i.e.* into strictly less constrained classes, or larger classes w.r.t. \sqsubseteq . Nevertheless, proposition 11 reveals that procedure 3 removes from the \dot{O}_i events that certainly do not participate in \dot{O} , and at the same time does not remove too many events, since $\dot{O} \equiv \dot{\wedge}_i \dot{O}'_i$. In the particular case where \dot{O} is a standard branching process \mathcal{O} , the limit classes have a certain practical interest: since $\mathcal{O} = \dot{\wedge}_i \dot{O}'_i$, the behaviors they forbid are also forbidden by \mathcal{O} . This is made clearer in the following lemma.

Lemma 12 *Let the $\dot{M}_{i,j}$ be messages obtained at convergence of procedure 3 (in the sense of equivalence classes): one has $\dot{O}'_i = \text{Trim}[\dot{O}_i \dot{\wedge} (\dot{\wedge}_{k \in \mathcal{N}(i)} \dot{M}_{k,i})]$. Define the ABPs $\dot{O}_{i,0} = \psi_i[\dot{O}_i \dot{\wedge} (\dot{\wedge}_{k \in \mathcal{N}(i)} \dot{M}_{k,i})]$. These prefixes $\dot{O}_{i,0}$ of the \dot{O}_i do not depend on the convergence point of procedure 3, and satisfy $\dot{O} = \dot{\wedge}_i \dot{O}_{i,0}$.*

Proof. Let the $\dot{M}_{i,j}$ and $\dot{M}'_{i,j}$ be messages obtained at two convergence points of procedure 3: one has $\dot{M}_{i,j} \equiv \dot{M}'_{i,j}$, so $\dot{M} \triangleq \dot{\wedge}_{k \in \mathcal{N}(i)} \dot{M}_{k,i} \equiv \dot{\wedge}_{k \in \mathcal{N}(i)} \dot{M}'_{k,i} \triangleq \dot{M}'$. We thus have to prove that $\psi_i(\dot{O}_i \dot{\wedge} \dot{M})$ is isomorphic to $\psi'_i(\dot{O}_i \dot{\wedge} \dot{M}')$ when $\dot{M} \equiv \dot{M}'$. Let $\kappa_i \dot{\wedge} \kappa$ be a configuration of $\dot{O}_i \dot{\wedge} \dot{M}$, with $\psi_i(\kappa_i \dot{\wedge} \kappa) = \kappa_i$. There exists κ' in \dot{M}' which is a weaker version of κ , so $\kappa_i \dot{\wedge} \kappa'$ in $\dot{O}_i \dot{\wedge} \dot{M}'$ satisfies $\psi'_i(\kappa_i \dot{\wedge} \kappa') = \kappa_i$. So $\psi_i(\dot{O}_i \dot{\wedge} \dot{M})$ and $\psi'_i(\dot{O}_i \dot{\wedge} \dot{M}')$ have isomorphic configurations, and we conclude by proposition 1.

For the last relation, observe that $\dot{\wedge}_i \dot{O}_{i,0} \sqsubseteq \dot{\wedge}_i \dot{O}_i = \dot{O}$. For the converse prefix relation, let $\kappa = \dot{\wedge}_i \kappa_i$ be a configuration of \dot{O} . One easily checks by recursion in procedure 3 that $\kappa_{j|\mathcal{V}_i \cap \mathcal{V}_j}$ is always present in message $\dot{M}_{j,i}$ after its first update. So κ_i is also present in $\dot{O}_{i,0}$, and we conclude again by proposition 1. \square

In the case of systems not living on a tree, procedure 3 appears as an approximate reduction strategy: it removes from the initial factors \dot{O}_i of \dot{O} events that make no contribution to \dot{O} . There exist several other interesting properties of limit classes, which prove that procedure 3 doesn't perform a trivial task and heavily prunes the initial factors \dot{O}_i . In particular, every configuration of a \dot{O}'_i can be extended into a larger configuration of $\dot{\wedge}_{j \in J} \dot{O}'_j$, where the index set J defines a spanning tree around \mathcal{S}_i on graph \mathcal{G}^e . All these properties rely on the fact that procedure 3 is actually an iterative constraint solving algorithm. The reader is referred to [20] for more details.

9 Discussion and conclusion

We have proposed a framework to perform modular computations on event structures related to a possibly large distributed system \mathcal{S} . This calculus operates on a factorized representation of runs of system \mathcal{S} , and is based on two operations: product and projection. Our target objects are typically branching processes of components of \mathcal{S} , but computations must be performed with richer structures. We have introduced four nested categories of event structures (fig. 15). Augmented branching processes (ABPs) were introduced to keep track of causality and conflict relations in the projection operation. This category is not stable under product, whence the introduction of generalized augmented branching processes, and of the trimming operation, projecting the result back to the category of ABPs. The computations we have presented consider projections and product as elementary operations, which

may look abusive since they potentially operate on infinite objects. However, this presents no practical difficulty since product, projection and trimming can be further decomposed into recursive procedures.




Figure 15: *The four nested categories we have introduced, with the effects of product \wedge , of trimming and of minimization (inside an equivalence class). The validity domain of the involutivity property is also depicted.*

The extra category of minimal ABPs introduced in the last sections appears as the most interesting one in many aspects. First of all, it captures extra causality or conflict relations in the most compact manner: a reinforced configuration is preserved, with respect to a more permissive version, only if it leads to a different “future.” Secondly, minimal ABPs bring back the involutivity property: this property is a desirable feature as it expresses that composing an event structure with part of itself doesn’t increase the amount of information carried by this structure. Involutivity is also crucial to get the convergence of iterative message passing procedures, and provides several other properties on the limit that we have not detailed here. Observe however that the involutivity of minimal ABPs takes the form of an involutivity in the sense of equivalence classes in the category of ABPs, with (\subseteq, \equiv) replacing $(\sqsubseteq, =)$. Finally, minimal product coverings in the category of minimal ABPs have the same and intuitive properties than minimal product coverings of standard branching processes. These features suggest that minimal ABPs are in fact the correct generalization of ordinary branching processes, in order to perform modular computations. Their main drawback comes from the difficulty to compute the minimal element of an equivalence class of ABPs, which cannot be done in a recursive manner, by contrast with trimming. This appears in particular in the fact that a prefix of a minimal ABP is not a minimal ABP: the “correct” prefix relation for minimal ABPs is \subseteq rather than \sqsubseteq . In summary, whereas the relevant event structures are certainly minimal ABPs, computations must be done in the larger category of ABPs, and performed “up to equivalence.”

The event structures we have used all along this paper are closely related to branching processes associated to tile systems. They clearly separate “structural” causality and conflict relations, due to tiles operating on a given component, from “inherited” relations, coming from the other components (specifically after projection). This distinction is not crucial however, and product, trimming and projection could probably be rephrased in terms of labeled prime event structures. By contrast, it is less easy to get rid of the notion of variable. Their first important role appears in the definition of the composition of systems, since components interact both by sharing variables and tile names (corresponding to a synchronization). However, a careful look at the definition of extended components reveals that the essential function of this transform is to gather all interactions under the form of a standard synchronous product (variables can then be considered as “private” to each extended component). This is the key to the factorization result on unfoldings of a compound system, as shown by Winskel in [3]. The second and most important role of variables appears in theorem 1, that forms the basis of modular computations. Specifically, we require that variables shared by two components

make visible *all* interactions between these components. This way of displaying interaction events is probably stronger than necessary, and we suspect the same function can be ensured in a lighter manner. For example, when the product of event structures is defined by a label algebra, by isolating subsets of labels where interaction does occur between components. All in all, the context of our results can probably be reshaped. But we believe that, in one form or another, the factorized representation of unfoldings together with modular computation techniques provide a key to deal with large systems.

Finally, let us come back to the case of systems not living on a tree. As we have seen, modular computations do not give the exact projections of a product branching process of the global system. Nevertheless, they yield goods approximations, where discarded configurations in one component are certainly not possible in the global system, while however some of the remaining local configurations may also be impossible to observe. This suggests that low cost approximate analysis strategies for distributed systems may already be interesting.

References

- [1] E. Fabre, Factorization of Unfoldings for Distributed Tile Systems, Part 1 : Limited Interaction Case, Inria research report no. 4829, April 2003, submitted for publication.
- [2] M. Nielsen, G. Plotkin, G. Winskel, Petri nets, event structures and domains, Theoretical Computer Science 13(1), 1981, pp. 85-108.
- [3] G. Winskel, Categories of models for concurrency, Seminar on Concurrency, Carnegie-Mellon Univ. (July 1984), LNCS 197, pp. 246-267, 1985.
- [4] G. Winskel, Event structure semantics of CCS and related languages , LNCS 140, 1982, also as report PB-159, Aarhus Univ., Denmark, April 1983.
- [5] J. Engelfriet, Branching Processes of Petri Nets, Acta Informatica 28, 1991, pp. 575-591.
- [6] K.L. McMillan, Using unfoldings to avoid the state explosion problem in the verification of asynchronous circuits, in Proc. 4th Workshop of Computer Aided Verification, Montreal, 1992, pp. 164-174.
- [7] J. Esparza, S. Römer, An unfolding algorithm for synchronous products of transition systems, in Proc. of CONCUR'99, LNCS 1664, Springer Verlag, 1999.
- [8] J. Esparza, S. Römer, W. Vogler, An improvement of McMillan's unfolding algorithm, in Proc. of TACAS'96, LNCS 1055, pp. 87-106.
- [9] J. Esparza, S. Römer, W. Vogler, An Improvement of McMillan's Unfolding Algorithm, Formal Methods in System Design 20(3), pp. 285-310, May 2002. Extended version of [8].
- [10] J. Esparza, Model checking using net unfoldings, Science of Computer Programming 23, pp. 151-195, 1994.
- [11] J. Esparza, C. Schröter, Reachability Analysis Using Net Unfoldings, Workshop of Concurrency, Specification and Programming, volume II of Informatik-Bericht 140, pp. 255-270, Humboldt-Universität zu Berlin, 2000.
- [12] S. Melzer, S. Römer, Deadlock checking using net unfoldings, CAV'97, LNCS 1254, pp. 352-363.
- [13] P. Degano, R. De Nicola, U. Montanari, On the Consistency of "Truly Concurrent" Operational and Denotational Semantics, in proc. Symposium on Logic in Computer Science (LICS) 1988, pp. 133-141.
- [14] F. W. Vaandrager, A simple definition for parallel composition of prime events structures, Report CS-R8903, CWI, Amsterdam, March 1989.
- [15] I. Castellani, G.-Q. Zhang, Parallel product of event structures, Theoretical Computer Science, no. 179, pp. 203-215, 1997.
- [16] J.-M. Couvreur, S. Grivet, D. Poirineau, Unfolding of Products of Symmetrical Petri Nets, 22nd International Conference on Applications and Theory of Petri Nets (ICATPN 2001), Newcastle upon Tyne, UK, June 2001, LNCS 2075, pp. 121-143.
- [17] A. Arnold, Finite Transition Systems, Prentice Hall, 1992.
- [18] L. Lamport, N. Lynch, Distributed Computing: Models and Methods, in Handbook of Theoretical Computer Science, vol. B: Formal Models and Semantics, Jan van Leeuwen ed., Elsevier (1990), pp. 1157-1199.
- [19] M. Raynal, Distributed algorithms and protocols, Wiley & Sons, 1988.
- [20] E. Fabre, Convergence of the turbo algorithm for systems defined by local constraints, Iria research report no. PI 1510, May 2003.
- [21] A. Benveniste, E. Fabre, S. Haar, C. Jard, Diagnosis of asynchronous discrete event systems, a net unfolding approach, IEEE Trans. on Automatic Control, vol. 48, no. 5, pp. 714-727, May 2003.
- [22] E. Fabre, Compositional Models of Distributed and Asynchronous Dynamical Systems, 41st Conf. on Decision and Control, Las Vegas, Dec. 2002, pp. 1-6.
- [23] E. Fabre, V. Pigourier, Monitoring distributed systems with distributed algorithms, 41st Conf. on Decision and Control, Las Vegas, Dec. 2002, pp. 411-416.
- [24] A. Benveniste, S. Haar, E. Fabre, C. Jard, Distributed and Asynchronous Discrete Event Systems Diagnosis, in Proc. 42nd Conf. on Decision and Control, Hawaii, Dec. 2003.
- [25] A. Benveniste, S. Haar, E. Fabre, C. Jard, Distributed monitoring of concurrent and asynchronous systems (plenary address), in Proc. of CONCUR'2003, Marseille, LNCS no. 2761, Pages 1-26, 2003.
- [26] E. Fabre, Distributed Diagnosis for Large Discrete Event Dynamic Systems, in preparation.



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

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