

Measuring Concurrency of Regular Distributed Computations

Cyrille Bareau, Benoit Caillaud, Claude Jard, René Thoraval

► **To cite this version:**

Cyrille Bareau, Benoit Caillaud, Claude Jard, René Thoraval. Measuring Concurrency of Regular Distributed Computations. [Research Report] RR-2394, INRIA. 1994. <inria-00074281>

HAL Id: inria-00074281

<https://hal.inria.fr/inria-00074281>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Measuring Concurrency of Regular Distributed Computations

Cyrille Bateau, Benoît Caillaud, Claude Jard, René Thoraval

N° 2394

Octobre

PROGRAMME 1



*Rapport
de recherche*



Measuring Concurrency of Regular Distributed Computations

Cyrille Bareau, Benoît Caillaud*, Claude Jard, René Thoraval**

Programme 1 — Architectures parallèles, bases de données, réseaux et systèmes distribués
Projet Pampa

Rapport de recherche n° 2394 — Octobre — 23 pages

Abstract: In this paper, we present a concurrency measure that is especially adapted to distributed programs that exhibit regular run-time behaviours. Such programs are frequently obtained by automatic parallelization of sequential code. This measure is based on the antichain lattice of the partial order that models the distributed execution under consideration. We show under which condition the measure is computable on an infinite execution which is the repetition of a finite pattern. The measure can then be computed by considering only a bounded number of patterns, this bound being at most the number of processors.

Key-words: Concurrency measure, Automated parallelization, Antichain lattice, Labeled partial orders, Regular executions.

(Résumé : tsvp)

This work was partially supported by the project Trace of the French Ministère de l'Enseignement Supérieur et de la Recherche.

E-mail adress of the authors: {name}@irisa.fr

*LFCS, JCMB, King's Buildings, The University of Edinburgh, Edinburgh, EH9 3JZ, UK

**Université de Nantes, Section Informatique, 2 rue de la Houssinière F-44072 Nantes cedex 03

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

Mesure de concurrence sur des exécutions réparties régulières

Résumé : Dans ce rapport, nous présentons une mesure de concurrence adaptée aux programmes répartis ayant un comportement régulier à l'exécution, tels que les programmes obtenus par parallélisation automatique de programmes séquentiels. Cette mesure se calcule sur le treillis des antichaînes de l'ordre partiel modélisant l'exécution répartie considérée. Nous montrons sous quelles conditions cette mesure est calculable sur une exécution infinie qui est la répétition d'un motif fini. La mesure peut alors être calculée en ne considérant qu'un nombre borné de motifs, cette borne étant au plus le nombre de processeurs de l'exécution.

Mots-clé : Mesure de concurrence, Parallélisation automatique, Treillis des antichaînes, Ordres partiels étiquetés, Exécutions régulières.

1 Introduction

The trend towards the use of distributed memory parallel machines is very evident. Clearly, their programming environments have to be significantly improved, especially in the field we are mainly interested in: semi-automated distribution of sequential code for scientific computing. Indeed, programmers need sophisticated performance evaluation tools. However, there is no well-accepted “complexity” criterion for distributed programs, for the behaviours of asynchronous message-passing programs are not yet sufficiently understood. It is also very difficult to design tools that can give relevant performance information by only statically analyzing a distributed code.

A research axis is then to study runs of a distributed program instead of the program itself. Especially, there is work to define concurrency measures, that is, measures that can help revealing the synchronization structure of a computation, whereas traditional ones (message count, for instance) only give quantitative information about the computation. For this, and as it is now usual [13], a distributed execution is seen as a partially ordered set of events that are causally related by process sequentiality and interprocess communication. As far as we know, the first concurrency measure that takes account of causality was proposed by Charron-Bost [4], followed by [8, 10, 15].

It is shown in [10] that either these measures are too inaccurate or their computational complexity makes them impracticable. It is also important to observe that all these measures only deal with finite runs, while executions of reactive programs are usually infinite. Also, executions of distributed programs for scientific applications, even finite, are usually very long. Nevertheless, [4] and [10] also give some encouraging results. The measure in [4] is shown to behave well with respect to a particular kind of concatenation operator on computations. On the other hand, [10] shows there is some hope to obtain practical concurrency measures for particular classes of executions (for instance, [10] considers executions that can be modeled as so-called interval orders).

We address several closely related problems. First, can a measure possibly be defined that gives significant values even in the case of infinite executions? If such a measure exists, another question immediately arises: does there exist a class of infinite computations for which the measure can be computed? Naturally, if such executions do exist, they must exhibit some kind of regularity which the measure must take into account. Moreover, as an infinite execution can be seen as a limit of a sequence of finite ones, the computation of the measure should not depend on the size of the order that models an execution: it must only depend on the size of a bounded subset of this order.

In this paper, we give a first positive answer to these problems. We define a concurrency measure and a class of executions that exhibit a particular kind of regularity: an execution in this class can be modeled as either an infinite or a finite repetition of a finite elementary order we call a basic pattern. For infinite (or very long) regular executions we establish that our measure is bounded if and only if the basic pattern of the execution is enough

connected. In this case the concurrency measure can be computed by taking account of a bounded number a repetitions of the basic pattern. Such an execution is said to be well-synchronized. Finally, we show this property is of interest for semi-automated distribution of sequential programs.

Paper Organization. We first describe the formal framework used throughout the paper and we present a model of an execution of a distributed program as a labeled poset.

Then we define our concurrency measure μ on a distributed execution. It is expressed in terms of the antichain lattice of the associated labeled poset. Measure μ associates a value with each event in the execution. An event with a small value denotes a strong synchronization, that is, an execution bottleneck.

Then we formally define regular executions and well-synchronization. For a regular well-synchronized execution, we show μ is bounded, it reflects regularity and can be computed on at most $2N - 1$ repetitions of the basic pattern (where N is the number of processes). This enables the definition of a measure μ_∞ on the events of the basic pattern. In the case of very long or infinite executions, the computation of this measure suffices to know μ on almost the whole execution. We also show that the antichain lattice of the infinite repetition of the pattern is regular enough so that μ_∞ can be computed on at most N repetitions of this pattern.

Then we compare μ with other measures from the literature.

Finally, we show our measure is especially relevant for automatically distributed programs.

2 Framework

2.1 Definitions and Notations

For an introduction to poset theory, see for instance [5].

A set E associated with a partial order relation \preceq is called a *partially ordered set* (*poset* for short) and is denoted by $\mathcal{E} = \langle E, \preceq \rangle$. Let $x, y \in E$: we say that x and y are *comparable* in \mathcal{E} when either $x \preceq y$ or $y \preceq x$, otherwise x and y are said to be *incomparable*, denoted by $x \parallel y$ (as usual, $x \preceq y \wedge x \neq y$ is denoted by $x \prec y$). A *chain* (resp. an *antichain*) in \mathcal{E} is a subset A of E such that every pair of distinct elements of A are comparable (resp. incomparable). Let $A \subseteq E$, $\max(A) = \{e \in A \mid \forall f \in A, e \not\prec f\}$ is the set of *maximal* elements in A . The *width* of \mathcal{E} is the maximum number of elements in an antichain in \mathcal{E} . The *covering relation* of \preceq is such that f covers e , denoted by $e \prec\prec f$, iff ($e \prec f$ and $\nexists g, e \prec g \prec f$). For each element e of E , we define $\downarrow e = \{f \in E \mid f \preceq e\}$ (the set of predecessors of e), and for each subset F of E , $\downarrow F = \bigcup_{f \in F} (\downarrow f)$.

The *Hasse diagram* of a poset \mathcal{E} is the directed graph whose vertices are the elements of E and the arcs are the elements of $\prec\prec$ (usually, the direction of the arcs is not represented by arrows but must be read bottom-up).

We define a labeled poset Θ as a tuple $\langle E, \preceq, L, \pi \rangle$ consisting of a non-empty poset $\langle E, \preceq \rangle$ of finite width and with no infinitely decreasing chain, a non-empty set L of labels and a labeling function $\pi : E \rightarrow L$.

The set of antichains of Θ is denoted by $\mathcal{A}(\Theta)$. This set is known to be a distributive lattice when equipped with the partial order \sqsubseteq defined as follows: $\forall A, B \in \mathcal{A}(\Theta)$, $A \sqsubseteq B \iff \downarrow A \subseteq \downarrow B$. Moreover it is easy to show that:

Lemma 1 *Let Θ be a labeled poset. Then $\forall A, B \in \mathcal{A}(\Theta)$,*

$$B \text{ covers } A \iff A \sqsubseteq B \wedge |\downarrow B \setminus \downarrow A| = 1$$

Then we can define $\mathcal{G}(\Theta) = \langle \mathcal{A}(\Theta), \Gamma(\Theta) \rangle$ where $\Gamma(\Theta) \subseteq \mathcal{A}(\Theta) \times E \times \mathcal{A}(\Theta)$ is the set of edges (A, e, B) such that B covers A and $(\downarrow B) \setminus (\downarrow A) = \{e\}$. We call this graph the labeled Hasse diagram of $\langle \mathcal{A}(\Theta), \sqsubseteq \rangle$.

2.2 Discrete Model of a Distributed Computation

Let us consider a computation of a distributed program, that is, a parallel run of a family $(P_i)_{i \in \{1, \dots, N\}}$ of N sequential processes that communicate by asynchronously exchanging messages. Let P denote the set $\{1, \dots, N\}$.

We define a discrete model of this computation as a labeled poset $\Theta = \langle E, \preceq, P, \pi \rangle$ that we will call a *distributed order* in the sequel. The elements of E are significant events that occur during the computation. The partial order \preceq indicates how these events are causally related (causality is based on process sequentiality and interprocess communication). The labeling function $\pi : E \rightarrow P$ associates with each event the identifier of the process it occurs on.

Process Sequentiality. $\pi^{-1}(\{i\})$ denotes the set of events that occur on any given process P_i . Since P_i runs sequentially, any two events of $\pi^{-1}(\{i\})$ are causally related, that is, $\pi^{-1}(\{i\})$ is modeled as a chain. Thus, the family $(\pi^{-1}(\{i\}))_{i \in P}$ is a N -chain decomposition of the labeled poset Θ , the width of which is therefore no more than N .

Concurrency. As we intend measuring concurrency throughout a computation, we are interested to know how far processes can simultaneously proceed at any given point of it. In the context of our discrete model Θ of a computation, that means we are interested in all the sets of events that are causally unconnected, i.e. the antichains of Θ .

Thus, the distributive lattice $\mathcal{A}(\Theta)$ of antichains of Θ well describes the dynamics of concurrency throughout the computation. Moreover, as the behaviours $(\pi^{-1}(\{i\}))_{i \in P}$ of the processes constitute a N -chain decomposition of Θ , this lattice can be given a graphical representation (see Fig. 1), for its Hasse diagram can be embedded into a N -dimensional grid with one dimension per process.




Figure 1: Hasse diagram of a distributed execution Θ on 3 processors and its graph $\mathcal{G}(\Theta)$

The left part of Fig. 1 shows the Hasse diagram of a computation Θ of a distributed program consisting of three processes P_1 , P_2 and P_3 . Its right part illustrates $\mathcal{A}(\Theta)$ by showing the labeled directed graph $\mathcal{G}(\Theta)$. In this graph, a path from the bottom to the top represents a linear extension of Θ .

3 A Concurrency Measure on the Antichain Lattice

3.1 Definition

The intuition underlying our measure is that the degree of concurrency of an event e is related to “what can happen simultaneously with e ”. The first idea is to count the number of processors ready to work, i.e. that are not blocked waiting for e . However this criterion is not accurate enough: it does not enable us to distinguish between, for instance, events e and h although e is clearly more concurrent than h (e is completely independent on the run of processors P_1 and P_2 , whereas h blocks the execution of events i , j , k , therefore the other processors). This idea can be refined by computing the number of antichains containing e , i.e. the number of configurations where the processor that performs e works in parallel with other processors (this is a refinement of a global indicator proposed by Charron-Bost [4]).

Definition 1 Let Θ be a distributed order $\langle E, \preceq, P, \pi \rangle$. For each event e of E :

1. $C_\Theta(e) = \{f \in E \mid f \parallel e\}$ is the set of events that are concurrent with e .
2. $\mathcal{A}_\Theta(e) = \{A \in \mathcal{A}(\Theta) \mid e \in A\}$ is the set of antichains that contain e .

The concurrency measure $\mu_\Theta : E \longrightarrow \mathbb{N} \cup \{\omega\}$ is defined as follows:

$$\forall e \in E, \quad \mu_\Theta(e) = |\mathcal{A}_\Theta(e)|$$

A large value of $\mu_\Theta(e)$ (we write $\mu(e)$ when no confusion is possible) means that many things may happen between the first and the last place where e can occur, that is, e has a great “latency” before actually occurring, therefore is “very concurrent”. In contrast, a little value means that this latency is very short: e is in fact a point of strong synchronization, a “bottleneck”.

In the execution of Fig. 1, events g and h for instance have respectively h, e, j, k and g, e as concurrent events; $\mu(h) = |\mathcal{A}_\Theta(h)| = |\{(h), (h, e), (h, g), (h, e, g)\}| = 4$ and $\mu(g) = |\{(g), (g, e), (g, h), (g, e, h), (g, h, j), (g, e, h, j), (g, e, h, j, k)\}| = 7$. The following table gives the values for all events of the execution:

a	b	c	d	e	f	g	h	i	j	k
3	3	3	2	12	7	4	5	6	7	3

In Figure 1, we can see on the lattice that, for instance, events d and f strongly synchronize the execution: nothing happens on processor P_1 simultaneously, and only e can occur on P_3 . In contrast, e does not depend on the computation on processors P_1 and P_2 and therefore is very concurrent because there are several possible configurations where it can be executed.

3.2 Computation

In the previous section, we have defined the measure for an event e in terms of the set of antichains that contain it. To compute this measure, we need to count the edges labeled by e in the labelled Hasse diagram $\mathcal{G}(\Theta)$ of the lattice of the execution:

Proposition 1 Let Θ be a distributed order and $e \in E$. Then

$$\mu_\Theta(e) = |\{B \in \mathcal{A}(\Theta) \mid \exists A \in \mathcal{A}(\Theta), (A, e, B) \in \Gamma(\Theta)\}|$$

Proof

From Lemma 1, $\forall A, B \in \mathcal{A}(\Theta), \forall e \in E, (A, e, B) \in \Gamma(\Theta)$ iff $A = \max((\downarrow B) \setminus \{e\})$. Then $\mathcal{A}_\Theta(e) = \{B \in \mathcal{A}(\Theta) \mid \exists A \in \mathcal{A}(\Theta), (A, e, B) \in \Gamma(\Theta)\}$. \square

This is computationally equivalent to the problem of counting the antichains of an order, which is known to be $\#P$ -complete.

This leads us to look for executions for which the number of antichains to be counted does not depend on the length of the execution.

4 Regular and Well-synchronized Executions

The aim of this section is to study a particular class of regular executions: executions which are finite or infinite repetitions of an elementary one (for instance a loopbody). We will show that for a subclass of these executions, μ is computable even in the infinite case by only taking into account at most N (the number of processes) repetitions of a basic “pattern”.

4.1 Regular Executions

Definition 2 (Regular distributed order) *Let $p \in \mathbb{N}^+ \cup \{\omega\}$ and $\Theta = \langle E, \leq, P, \pi \rangle$ be a finite distributed order. The distributed order $\Theta_p = \langle E_p, \leq_p, P, \pi_p \rangle$ is defined up to order-isomorphism by:*

- $E_p = \bigcup_{i \in [0, p[} \varphi_i(E)$ where $(\varphi_i(E))_{i \in [0, p[}$ is a sequence of mutually disjoint isomorphic copies of E .
- $\forall i \in [0, p[, \pi_p \circ \varphi_i = \pi$
- \leq_p is the least order relation such that:
 - $\forall i \in [0, p[, \forall e, f \in E, \varphi_i(e) \leq_p \varphi_i(f) \iff e \leq f$
 - $\forall i, j \in [0, p[, \forall e, f \in E, (\pi(e) = \pi(f)) \wedge (i < j) \implies \varphi_i(e) \leq_p \varphi_j(f)$

We say that a distributed order Φ is regular if there exists $p \in \mathbb{N}^+ \cup \{\omega\}$ and a finite distributed order Θ such that Φ is order isomorphic to Θ_p .

When in the sequel we consider a regular distributed order Θ_p , \leq_p and π_p are denoted by \leq and π for the sake of clarity. To speak about events of Θ_p more conveniently, we identify $\varphi_0(E)$ with E and we use a function $\lambda : E_p \longrightarrow E_p$ defined as follows: $\forall e \in E, \forall i \in [0, p - 1[, \lambda(\varphi_i(e)) = \varphi_{i+1}(e)$. This allows us to use the non-negative powers of λ instead of φ_i because $\forall e \in E, \forall i \in [0, p[, \varphi_i(e) = \lambda^i(e)$ (see Fig 2.(1) for a very simple example) and, as λ is clearly injective, its negative powers can also be used. Moreover, λ has the following property:

Lemma 2 *Let Θ_p be a regular distributed order. Then $\forall e, f \in E, \forall i, j \in [1, p[,$*

$$\lambda^{i-1}(e) \preceq \lambda^{j-1}(f) \iff \lambda^i(e) \preceq \lambda^j(f)$$




Figure 2: Notations for regular executions

In other words, λ preserves the order relation \preceq : it is an order isomorphism from $\bigcup_{i \in [0, p-1[} \lambda^i(E)$ onto $\bigcup_{i \in [1, p[} \lambda^i(E)$. On the Hasse diagram of Θ_p , λ is represented as a “one pattern upward shift”. Therefore the Hasse diagram of Θ_p is invariant by pattern-wise translations.

Proof

Let $e, f \in E$ and $i, j \in [1, p[$ such that $\lambda^{i-1}(e) \preceq \lambda^{j-1}(f)$. Then there exists a finite path $\lambda^{i-1}(e) = e_0 \prec e_1 \prec \dots \prec e_h = \lambda^{j-1}(f)$. Let us consider any comparability $e_k \prec e_{k+1}$. If $\exists l \in [0, p[$ such that $e_k, e_{k+1} \in \lambda^l(E)$, then $\lambda(e_k) \prec \lambda(e_{k+1})$. Otherwise, from definition of order concatenation, we have $\pi(e_k) = \pi(e_{k+1})$ and $\exists l \in [0, p-1[$ such that $e_k \in \lambda^l(E)$ and $e_{k+1} \in \lambda^{l+1}(E)$, therefore again $\lambda(e_k) \prec \lambda(e_{k+1})$. Thus $\lambda(e_0) \prec \lambda(e_1) \prec \dots \prec \lambda(e_h)$, i.e. $\lambda^i(e) \preceq \lambda^j(f)$. The converse is similar. \square

4.2 Well-synchronized Executions

For a given event e , the number of events that are incomparable with e in an infinite execution can be *a priori* infinite, and so $\mu(e) = \omega$. It is however interesting to see if there exists infinite regular executions for which all events have finite measures.

We will show (Lemma 3) that this property is related to the communication scheme of the basic pattern. We first introduce the notion of *communication graph* of a distributed order.

Definition 3 We call *communication graph* of a distributed order Θ the quotient of its Hasse diagram by the equivalence relation induced by π^{-1} , i.e. the directed graph $\mathcal{Q}(\Theta) = \langle V, C \rangle$ where:

- $V = \pi(E)$
- $C = \{(\alpha, \beta) \in V \times V \mid \exists e \in \pi^{-1}(\{\alpha\}), f \in \pi^{-1}(\{\beta\}), e \not\prec f\}$

A distributed order Θ whose associated communication graph is strongly connected is said to be *well-synchronized*. The diameter of $\mathcal{Q}(\Theta)$ is then denoted by $k_\Theta - 1$.

Note that $k_\Theta \in [1, N]$, N being the number of processes.

Lemma 3 Let Θ be a finite distributed order. The concurrency measure μ_{Θ_ω} of any event in $\Theta_\omega = \langle E_\omega, \preceq_\omega, P, \pi_\omega \rangle$ is finite if and only if Θ is well-synchronized.

Proof

$$\begin{aligned} \exists e \in E_\omega, \mu_{\Theta_\omega}(e) = \omega &\iff \exists e \in E_\omega, |C_{\Theta_\omega}(e)| = \omega \iff (P \text{ is finite}) \exists e \in E_\omega, \\ \exists \alpha \in \pi(E_\omega), |C_{\Theta_\omega}(e) \cap \pi^{-1}(\{\alpha\})| = \omega &\iff (\Theta_\omega \text{ has no infinitely decreasing chain}) \\ \exists e \in E_\omega, \exists \alpha \in \pi(E_\omega), (\forall f \in \pi^{-1}(\{\alpha\}), e \not\prec f) &\iff (\text{Lemma 2}) \exists \alpha, \beta \in \pi(E_\omega), \\ \forall f \in \pi^{-1}(\{\alpha\}), \forall e \in \pi^{-1}(\{\beta\}), e \not\prec f &\iff (\text{Definition 2}) \mathcal{Q}(\Theta) \text{ is not strongly} \\ &\text{connected.} \quad \square \end{aligned}$$

4.3 Boundedness and Regularity of the Measure

A property stronger than Lemma 3 can immediately be drawn: for a well-synchronized regular execution, the measure of the events is not only finite but also bounded and regular.

Lemma 4 Let Θ_p be a regular well-synchronized distributed order.

1. $\forall i \in [0, p[, \forall e \in \lambda^i(E),$

$$C_{\Theta_p}(e) \subseteq \bigcup_{j=\max(0, i-k_\Theta+1)}^{\min(p-1, i+k_\Theta-1)} \lambda^j(E)$$

2. $\forall i \in [k_\Theta - 1, p - k_\Theta[, \forall e \in \lambda^i(E)$,

$$C_{\Theta_p}(\lambda(e)) = \lambda(C_{\Theta_p}(e))$$

Proof

1. Let $i \in [0, p[$ and $e \in \lambda^i(E)$. Let $j \in [i + k_\Theta, p[$, and $f \in \lambda^j(E)$. From Definition 3, since the diameter of $\mathcal{Q}(\Theta)$ is $k_\Theta - 1$, $\exists e' \in \lambda^{i+1}(E)$ with $\pi(e') = \pi(e)$ and $\exists f' \in \lambda^{i+k_\Theta-1}(E)$, $\pi(f') = \pi(f)$ such that $e' \preceq f'$. From Definition 2, $e \preceq e'$ and $f' \preceq f$, therefore e and f are comparable. In the same way, we can show that $\forall j \in [0, i - k_\Theta]$, $\forall f \in \lambda^j(E)$, $f \preceq e$.
2.
 - Let $f \in \lambda(C_{\Theta_p}(e))$, i.e. $f = \lambda(h)$ with $h \parallel e$, hence $h \in \bigcup_{j=\max(0, i-k_\Theta+1)}^{\min(p-1, i+k_\Theta-1)} \lambda^j(E)$ (from Lemma 4(1)). Now, $i < p - k_\Theta \Rightarrow i + k_\Theta - 1 < p - 1$, therefore Lemma 2 can apply: $\lambda(h) \parallel \lambda(e)$ thus $f = \lambda(h) \in C_{\Theta_p}(\lambda(e))$.
 - Let $f \in C_{\Theta_p}(\lambda(e))$, i.e. $f \parallel \lambda(e)$, hence $f \in \bigcup_{j=\max(0, i-k_\Theta+2)}^{\min(p-1, i+k_\Theta)} \lambda^j(E)$ (from Lemma 4(1)) with $i \geq k_\Theta - 1$, hence $\max(0, i - k_\Theta + 2) \geq 1$, therefore $\exists f'$, $f = \lambda(f')$ and Lemma 2 can apply: $f' \parallel e$, thus $f \in \lambda(C_{\Theta_p}(e))$.

□

The measure clearly is bounded: $\forall e \in E$, $\mu(e) \leq |\mathcal{A}(\Theta_{2k_\Theta-1})|$. By Lemma 2, it is also regular: $\forall i, j \in [k_\Theta - 1, p - k_\Theta + 1[, \forall e \in E: \mu_{\Theta_p}(\lambda^i(e)) = \mu_{\Theta_p}(\lambda^j(e))$.

Moreover, μ can be computed on $\mathcal{A}(\Theta_{2k_\Theta-1})$ (even in the case of infinite executions), thus taking at most $2N - 1$ patterns into account (N being the number of processes). For infinite or very long regular well-synchronized executions whose basic patterns have reasonable sizes, μ can be realistically computed.

We have implemented the computation of μ in our distributed environment [11] based on the Estelle specification language. This environment provides a mechanism of vectorial clocks [14], that are traced “on line”. These traces are used as input for our algorithm of construction of the antichain lattice of an order [6, 12].

When given as input a linear extension of $\Theta_{2k_\Theta-1}$, this algorithm has a time complexity of $\mathcal{O}(|\mathcal{A}(\Theta_{2k_\Theta-1})| + |\Gamma(\Theta_{2k_\Theta-1})| + N \times |\Theta_{2k_\Theta-1}|^2)$.

The fact that μ reflects regularity makes useful the definition of a measure μ_∞ on the basic pattern:

Definition 4 *Let Θ be a finite well-synchronized distributed order, we write:*

$$\forall e \in E, \quad \mu_\infty(e) = \mu_{\Theta_\omega}(\lambda^{k_\Theta-1}(e))$$

As clearly, for all event $e \in E$, $\mu_\infty(e) = \mu_{\Theta_{2k_\Theta-1}}(\lambda^{k_\Theta-1}(e))$, μ_∞ can also be computed on $\mathcal{A}(\Theta_{2k_\Theta-1})$.

We also have a stronger result on μ_∞ : it can be computed on a subgraph of the labeled Hasse diagram $\mathcal{G}(\Theta_{k_\Theta})$ of $\mathcal{A}(\Theta_{k_\Theta})$, that is, by only taking account of at most N repetitions

of the basic pattern ¹. This result does not significantly improve the complexity of the computation of μ_∞ (which becomes in N instead of $2N-1$) but enlightens the regular structure of the antichain lattice.

4.4 Regularity of the Antichain Lattice of an Infinite Regular Well-synchronized Execution

From now on, we will always consider an infinite regular well-synchronized distributed order Θ_ω . We first give a formal expression for the regular structure the labeled Hasse diagram $\mathcal{G}(\Theta_\omega)$ of its antichain lattice $\langle \mathcal{A}(\Theta_\omega), \sqsubseteq \rangle$ exhibits ². This regularity, although less trivial than that of the Hasse diagram of Θ_ω , allows a finite representation of $\mathcal{G}(\Theta_\omega)$ that can be computed on a subgraph of $\mathcal{G}(\Theta_{k_\Theta})$ of $\langle \mathcal{A}(\Theta_{k_\Theta}), \sqsubseteq \rangle$. In the sequel, we shall take advantage of this regularity to compute μ_∞ more efficiently, that is, by only taking account of this subgraph of $\mathcal{G}(\Theta_{k_\Theta})$ instead of the whole labeled Hasse diagram of $\langle \mathcal{A}(\Theta_{2k_\Theta-1}), \sqsubseteq \rangle$.

4.4.1 Notations and preliminary results.

We give here some notations for expressing the regularity of the lattice, and present a few technical lemmas.

We first deal with the vertices of $\mathcal{G}(\Theta_\omega)$:

- $\Lambda : \mathcal{A}(\Theta_\omega) \longrightarrow \mathcal{A}(\Theta_\omega)$ denotes the function such that $\forall A \in \mathcal{A}(\Theta_\omega)$, $\Lambda(A) = \max(\lambda(A) \cup E)$. Clearly, Λ is injective and order preserving (Lemma 2). Moreover, $\forall i \in \mathbb{N}$, $i \geq 1$, $\Lambda^i(\emptyset) = \lambda^{i-1}(\max(E))$ and $\Lambda^{i-1}(\emptyset) \sqsubseteq \Lambda^i(\emptyset)$ (see Fig. 2(2)).
- $\perp : \mathcal{A}(\Theta_\omega) \longrightarrow \mathbb{N}$ is such that $\forall A \in \mathcal{A}(\Theta_\omega)$, $\perp(A) = \max(\{i \in \mathbb{N} \mid \Lambda^i(\emptyset) \sqsubseteq A\})$, which is well-defined because $\{i \in \mathbb{N} \mid \Lambda^i(\emptyset) \sqsubseteq A\}$ is nonempty (it contains 0) and bounded (any antichain in $\mathcal{A}(\Theta_\omega)$ is finite).
- $\forall i \in \mathbb{N}$, \mathcal{A}_i denotes the set $\{A \in \mathcal{A}(\Theta_\omega) \mid \perp(A) = i\}$ we call a “slice” of $\mathcal{A}(\Theta_\omega)$. Clearly, $(\mathcal{A}_i)_{i \in \mathbb{N}}$ is a partition of $\mathcal{A}(\Theta_\omega)$. Moreover $\Lambda(\mathcal{A}(\Theta_\omega)) \subseteq \mathcal{A}(\Theta_\omega) \setminus \mathcal{A}_0$.

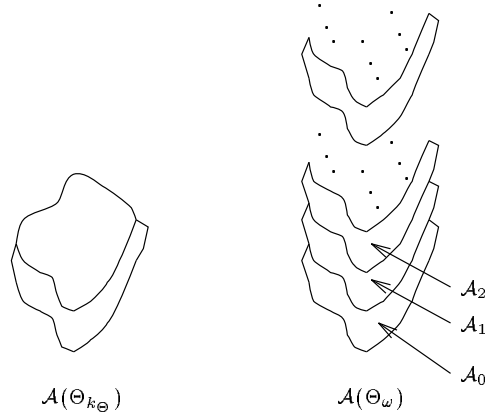
There are relations between the set of patterns of Θ_ω an antichain A intersects and the slice $\mathcal{A}_{\perp(A)}$ of $\mathcal{A}(\Theta_\omega)$ it belongs to. We distinguish between the elements in A according to whether they belong to $\Lambda^{\perp(A)}(\emptyset)$ (the bottom element of $\mathcal{A}_{\perp(A)}$) or not; this distinction corresponds to the one we can make between the elements in $\Lambda(A')$ ($A' \in \mathcal{A}_0$) according to whether they belong to $\Lambda(\emptyset) = \max(E)$ or to $\lambda(A')$.

Lemma 5 For any $A \in \mathcal{A}(\Theta_\omega)$,

$$1. A \cap \Lambda^{\perp(A)}(\emptyset) \subseteq \lambda^{\perp(A)-1}(E)$$

¹Clearly, μ_∞ cannot be computed on $\mathcal{G}(\Theta_p)$ where $0 < p < k_\Theta$.

²A similar but slightly more intricate result can be given for a regular well-synchronized distributed order Θ_p such that $p \in \mathbb{N}$ and $p \geq k_\Theta$.

Figure 3: Hasse diagram of $\mathcal{A}(\Theta_\omega)$

$$2. A \setminus \Lambda^{\perp(A)}(\emptyset) \subseteq \bigcup_{i=\perp(A)}^{\perp(A)+k_\Theta-1} \lambda^i(E)$$

Proof

The first inclusion is obvious and if $A = \Lambda^{\perp(A)}(\emptyset)$ the second one is also obvious. The assumption that $A \setminus \Lambda^{\perp(A)}(\emptyset)$ would intersect at least one on the $\perp(A)$ first patterns of Θ_ω immediately leads to $A \notin \mathcal{A}(\Theta_\omega)$, a contradiction. The assumption that it would intersect at least one pattern of Θ_ω “higher” than $\lambda^{\perp(A)+k_\Theta-1}(E)$ immediately leads (Lemma 4(1)) to $\Lambda^{\perp(A)+1}(\emptyset) \sqsubseteq A$, another contradiction. \square

We now consider the powers of Λ :

Lemma 6 $\forall A \in \mathcal{A}(\Theta_\omega)$,

1. $\forall i \in \mathbb{N}, i \geq 1$, we have $\Lambda^i(A) = \lambda^{i-1}(\Lambda(A))$.
2. $\forall i \in [1, \perp(A)[$, $\Lambda^{-i}(A) = \lambda^{-i}(A)$ and $\Lambda^{-\perp(A)}(A) = \lambda^{-1}(\lambda^{-\perp(A)+1}(A) \setminus \max(E))$.
3. $\forall i \in [-\perp(A), -1] \cup \mathbb{N}$, we have $\perp(\Lambda^i(A)) = \perp(A) + i$.

Proof

1. By induction. The case $i = 1$ is trivial, and for $i = 2$, $\forall A$ in the domain of Λ^2 , $\max(E) \sqsubseteq \Lambda(A) \Rightarrow \lambda(\Lambda(A)) = \Lambda^2(A)$. Let A be in the domain of Λ^{i+1} , $i \in [2, p[$, then $\Lambda^{i+1}(A) = \Lambda^i(\Lambda(A)) = \lambda^{i-1}(\Lambda(\Lambda(A))) = \lambda^{i-1}(\Lambda^2(A)) = \lambda^{i-1}(\lambda(\Lambda(A))) = \lambda^i(\Lambda(A))$.

2. Let $A \in \mathcal{A}(\Theta_p)$ and $i \in [1, \perp(A)]$, then $\lambda^{-1}(\lambda^{-i+1}(A) \setminus \max(E))$ is defined (Lemma 5). Let it be denoted by A' , we have $A' \in \mathcal{A}(\Theta_p)$ and $\Lambda(A')$ is defined (Lemmas 2 and 5). Moreover $\lambda^{-i+1}(A) \in \mathcal{A}(\Theta_p)$ and $\max(E) \sqsubseteq \lambda^{-i+1}(A)$ (Lemma 2), therefore $\Lambda(A') = \lambda^{-i+1}(A)$ hence $\Lambda^i(A') = A$ (Lemma 6(1)). Finally, $\Lambda^{-i}(A)$ is defined and $\Lambda^{-i}(A) = A'$.
3. Immediate from Lemmas 6(1,2) since Λ is order preserving.

□

Now, we have all we need to deal with the vertices of $\mathcal{G}(\Theta_\omega)$, but we also need similar results and notations to deal with its edges :

- $\tilde{\Lambda} : \Gamma(\Theta_\omega) \longrightarrow \Gamma(\Theta_\omega)$ denotes the function such that $\forall (A, e, B) \in \Gamma(\Theta_\omega)$, $\tilde{\Lambda}(A, e, B) = (\Lambda(A), \lambda(e), \Lambda(B))$, which can be easily shown well-defined and injective by using Lemmas 1 and 2 and the fact that Λ is an order isomorphism from $\mathcal{A}(\Theta_\omega)$ to $\mathcal{A}(\Theta_\omega) \setminus \mathcal{A}_0$.
- $\forall i \in \mathbb{N}$, Γ_i denotes the set $\{(A, e, B) \in \Gamma(\Theta_\omega) \mid A \in \mathcal{A}_i\}$. Clearly, $(\Gamma_i)_{i \in \mathbb{N}}$ is a partition of $\Gamma(\Theta_\omega)$, i.e. each Γ_i is a “slice” of $\Gamma(\Theta_\omega)$.

There are relations between the slice $\Gamma_{\perp(A)}$ of $\Gamma(\Theta_\omega)$ an edge (A, e, B) belongs to and the slices $\mathcal{A}_{\perp(A)}$ and $\mathcal{A}_{\perp(B)}$ of $\mathcal{A}(\Theta_\omega)$:

Lemma 7 $\forall (A, e, B) \in \Gamma(\Theta_\omega)$,

1. $\perp(A) \leq \perp(B) \leq \perp(A) + 1$
2. $e \in B \setminus \Lambda^{\perp(A)}(\emptyset)$
3. $\perp(B) = \perp(A) + 1 \iff e \in \Lambda^{\perp(A)+1}(\emptyset)$

Proof

1. The first inequality is obvious and, as $e \in B$ (Lemma 1), the assumption that the second inequality would not hold easily leads to $|(\downarrow B) \setminus (\downarrow A)| > 1$, which contradicts Lemma 1.
2. Assume $e \in \Lambda^{\perp(A)}(\emptyset)$, then this immediately leads to $e \in \downarrow A$, a contradiction.
3. If $\perp(B) = \perp(A) + 1$ then $e \in \Lambda^{\perp(A)+1}(\emptyset)$ from Lemma 1 and the converse follows from Lemma 7(1).

□

Finally, we consider the negative powers of $\tilde{\Lambda}$:

Lemma 8 $\forall (A, e, B) \in \Gamma(\Theta_\omega)$,

$$\tilde{\Lambda}^{-\perp(A)}(A, e, B) = (\Lambda^{-\perp(A)}(A), \lambda^{-\perp(A)}(e), \Lambda^{-\perp(A)}(B))$$

Proof

Let $\gamma = (\Lambda^{-\perp(A)}(A), \lambda^{-\perp(A)}(e), \Lambda^{-\perp(A)}(B))$. Then γ clearly is defined (Lemmas 5, 6 and 7), and we immediately have $\gamma \in \Gamma(\Theta_\omega)$ from Lemmas 2 and 6(1,3), and from the fact that Λ is an order isomorphism from $\mathcal{A}(\Theta_\omega)$ to $\mathcal{A}(\Theta_\omega) \setminus \mathcal{A}_0$. Then $\tilde{\Lambda}^{\perp(A)}(\gamma)$ is defined and clearly is the edge (A, e, B) . \square

4.4.2 Regularity.

Proposition 2 *For any $i \in \mathbb{N}$,*

1. Λ^i is a one-to-one mapping of \mathcal{A}_0 onto \mathcal{A}_i .
2. $\tilde{\Lambda}^i$ is a one-to-one mapping of Γ_0 onto Γ_i .

Hence all the slices of $\mathcal{A}(\Theta_\omega)$ are order isomorphic³.

Proof

Λ and $\tilde{\Lambda}$ are injective, hence Λ^i and $\tilde{\Lambda}^i$ are so. By Lemma 6, we have $\Lambda^i(\mathcal{A}_0) \subseteq \mathcal{A}_i$ and $\tilde{\Lambda}^i(\Gamma_0) \subseteq \Gamma_i$, and we know that Λ^i is a mapping of \mathcal{A}_0 onto \mathcal{A}_i . Finally, using Lemma 8, we also know that $\tilde{\Lambda}^i$ is a mapping of Γ_0 onto Γ_i . \square

If we look at $\mathcal{G}(\Theta_\omega)$, we can see that the subset $\mathcal{A}_0 \cup \{B \in \mathcal{A}_1 \mid B \cap \max(E) \neq \emptyset\}$ of $\mathcal{A}(\Theta_\omega)$ and the subset Γ_0 of $\Gamma(\Theta_\omega)$ define a subgraph of $\mathcal{G}(\Theta_\omega)$ that we denote by \mathcal{G}_0 . Furthermore, by Lemma 5, $\mathcal{A}_0 \cup \{B \in \mathcal{A}_1 \mid B \cap \max(E) \neq \emptyset\}$ is included in $\cup_{i \in [0, k_\Theta]} \lambda^i(E)$. Hence \mathcal{G}_0 is a subgraph of $\mathcal{G}(\Theta_{k_\Theta})$.

Then Proposition 2 shows that all the information about $\mathcal{G}(\Theta_\omega)$ is contained into the subgraph \mathcal{G}_0 of $\mathcal{G}(\Theta_{k_\Theta})$. Consequently, we can compute μ_∞ by only taking \mathcal{G}_0 into account.

4.4.3 Finite representation.

Proposition 2 also shows that $\mathcal{G}(\Theta_\omega)$ can be given a finite representation. Let $\psi(A)$ denote either the antichain A or the antichain $\Lambda^{-1}(A)$ according to whether $A \in \mathcal{A}_0$ or $A \in \{B \in \mathcal{A}_1 \mid B \cap \max(E) \neq \emptyset\}$. Then $\mathcal{G}(\Theta_\omega)$ can clearly be seen as an appropriate unfolding of the labeled directed graph whose set of vertices is \mathcal{A}_0 and whose set of edges is $\{(A, e, \psi(B)) \mid (A, e, B) \in \Gamma_0\}$. That is, $\mathcal{G}(\Theta_\omega)$ can be given a finite representation that can be computed on the subgraph \mathcal{G}_0 of $\mathcal{G}(\Theta_{k_\Theta})$. This representation can even be simplified by labeling the edges by e in place of $\lambda^i(e)$ for any $e \in E$ and any $i \in \mathbb{N}$ (this relabeling causes no problem because it follows from Lemma 1 that it cannot produce edges (A, e, B) and (A, e', B) such that $e \neq e'$).

³Note that for $p \geq k_\Theta$, $p \neq \omega$, all the $p - k_\Theta + 1$ first slices of $\mathcal{A}(\Theta_p)$ are also order isomorphic.

4.5 Computation of μ_∞

We now show how to compute μ_∞ on \mathcal{G}_0 ⁴:

Proposition 3 *Let Θ be a finite well-synchronized distributed order. Then $\forall e \in E$,*

$$\mu_\infty(e) = |\{(A, f, B) \in \Gamma_0 \mid \exists i \in \mathbb{N}, f = \lambda^i(e)\}|$$

Proof

From Proposition 1, it suffices to prove that the sets $\{(A, \lambda^{k_\Theta-1}(e), B) \in \Gamma(\Theta_\omega)\}$ and $\{(A, f, B) \in \Gamma_0 \mid \exists i \in \mathbb{N}, f = \lambda^i(e)\}$ are equipotent. For this, $\forall \gamma = (A, \lambda^{k_\Theta-1}(e), B) \in \Gamma(\Theta_\omega)$, let $\phi(\gamma)$ denote the edge $\Lambda^{-\perp(A)}(\gamma)$ in $\Gamma(\Theta_\omega)$. By Proposition 2(2), ϕ is clearly a well-defined function onto $\{(A, f, B) \in \Gamma_0 \mid \exists i \in \mathbb{N}, f = \lambda^i(e)\}$, and by Lemma 1, it is injective. Therefore ϕ is a one-to-one mapping of $\{(A, \lambda^{k_\Theta-1}(e), B) \in \Gamma(\Theta_\omega)\}$ onto $\{(A, f, B) \in \Gamma_0 \mid \exists i \in \mathbb{N}, f = \lambda^i(e)\}$. \square

This proposition gives us a new algorithm for the computation of μ_∞ . It relies on the computation of the subgraph \mathcal{G}_0 of $\mathcal{G}(\Theta_{k_\Theta})$ instead of the computation of $\mathcal{G}(\Theta_{2k_\Theta-1})$ (see subsection 4.3). Its time complexity is the same as for the computation of $\mathcal{G}(\Theta_{k_\Theta})$, that is, a time complexity of $\mathcal{O}(|\mathcal{A}(\Theta_{k_\Theta})| + |\Gamma(\Theta_{k_\Theta})| + N \times |\Theta_{k_\Theta}|^2)$ [6, 12] where $k_\Theta \leq N$ (recall N is the number of processes).

5 Comparison with Other Measures

In this section, we present some recently proposed concurrency measures. We show that ours can be relevantly compared with each of them, especially if we study the behaviour in case of infinite regular executions.

At first, we present two “global” measures, that is, they give a single value to quantify a whole execution. In our context (see section 6), we are not interested in such an approach, that can be suitable for other distributed programming problems, for it does not enable one to identify the bottlenecks of an execution. However, note that a “local” measure (that associates a value with each event of an execution) can be derived from a global one. The following measures are local ones. They are computed directly on the order of an execution, not on the antichain lattice; the advantage is a better complexity (polynomial in the number of events), but the drawback is that they are less accurate, and not suited to infinite well-synchronized regular executions.

Charron-Bost [4] As far as we know, this is the first attempt to take into account, for a concurrency measure, the causal structure of an execution. In our notations, this measure is:

$$m(\Theta) = \frac{|\mathcal{A}(\Theta)| - |E| - 1}{c(\Theta) - |E| - 1}$$

⁴It could also be computed in a very similar way on the folded representation of $\mathcal{G}(\Theta_\omega)$.

where $c(\Theta)$ is the number of antichains in a totally concurrent execution (i.e. any two events with distinct labels are incomparable) with as many processes and as many events per process as in Θ , that is to say $c(\Theta) = \prod_{i=1, N} (|\pi^{-1}(\{i\})| + 1)$. In fact, μ is quite similar to m : the idea is to count the number of antichains in an execution. The difference is that μ is a global measure and that it is normalized: it ranges from 0 in the worst case ($|E| + 1$ is the number of antichains of a totally sequential execution) to 1 in the best (a totally concurrent execution). But this normalization is unsuited to the infinite case: $c(\Theta_p) = \prod_{i=1, N} (p|\pi^{-1}(\{i\})| + 1) = p^N \prod_{i=1, N} (|\pi^{-1}(\{i\})| + 1/p)$, hence for a well-synchronized execution:

$$m(\Theta_p) = \frac{(p - k_\Theta + 1) |\mathcal{A}_0| + |\mathcal{A}(\Theta_{k_\Theta}) \setminus \mathcal{A}_0| - p|E| - 1}{p^N \prod_{i=1, N} (|\pi^{-1}(\{i\})| + 1/p) - p|E| - 1} \Rightarrow m(\Theta_\omega) = 0$$

Habib et al. [10] They propose a “worst-case-measure” of the concurrency of an execution: the minimal size of the maximal (for inclusion) antichains of its associated poset. It can be seen as the number of processors that can proceed during the worst bottleneck of the execution.

This measure could be made local by computing for each event e the minimal size of the maximal antichains that contain e . As this measure depends on the antichains of a poset, it is obvious from our framework and results that such a derived measure presents the same regularity properties as ours. The difference with our measure is that we consider all the antichains which contain an event, whereas they only consider the maximal antichains.

Fidge [8] Fidge proposes a local measure, that he extends to a global one. The measure β of an event e in an execution Θ is defined as follows:

$$\beta(e) = \frac{|\downarrow e| - 1 - h(e)}{|\downarrow e| - 1 - a}$$

(Fidge proposes two closely related measures, whether $a = 1$ or $a = 1/N$). $h(e)$ is the “height” of e , i.e. the length of the longest chain ending by e .

If we consider a well-synchronized regular execution Θ_ω , we can see that for $p > k_\Theta$, $k_\Theta h'(E) + (p - k_\Theta)h(E) \leq h(\lambda^p(\hat{e})) \leq h(e) + ph(E)$ where $h'(E)$ is the minimal rank of a maximal element in E , and $(p - k_\Theta)|E| + k_\Theta \leq |\downarrow \lambda^p(\hat{e})| \leq (p + 1)|E|$ therefore, clearly, when p tends to infinity, $\beta(\lambda^p(\hat{e}))$ converges:

$$\forall e, \lim_{p \rightarrow \infty} \beta(\lambda^p(e)) = \frac{|E| - h(E)}{|E|}$$

We obtain a finite and computable measure, but it does not preserve regularity and for infinite executions, the measure is identical for almost all events.




Figure 4: Measures for the Jacobi algorithm.

Raynal et al. [15] This measure is a variant of Fidge’s one:

$$\alpha(e) = \frac{|\downarrow e| - 1 - h(e)}{v(\downarrow e) - 1 - h(e)}$$

with $v(\downarrow e)$ defined as the “volume” of the causal past of e , i.e. $v(\downarrow e) = \sum_{i=1}^N (h(e_i) + 1)$ with $e_{\pi(e)} = e$ and for all $i \neq \pi(e)$, e_i is the maximum of the predecessors of e in $\pi^{-1}(\{i\})$.

For well-synchronized executions, if $p > k_{\Theta}$, we have $(p - k_{\Theta})v(E) + k_{\Theta} \leq |v(\lambda^p(\hat{e}))| \leq (p + 1)v(E)$ therefore α converges as well:

$$\lim_{p \rightarrow \infty} \alpha(\lambda^p(e)) = \frac{|E| - h(E)}{v(E) - h(E)}$$

We have computed these measures on a kernel of the Jacobi algorithm, automatically distributed on five processors. The results are presented in Fig. 4.

The values for μ must be read on the right vertical axis, and the values for the three other measures (α , β , m) on the left axis: from 0 to 1 because these measures are normalized. Therefore, for comparison sake, the exact values are not to be taken into account, but only the variations of the graph. Charron-Bost’s measure m is also presented on this graph, but the plots are in fact the values $m(\Theta_p)$, Θ being the pattern observed on this execution.

This diagram clearly shows that our measure remains relevant when time flows, whereas the others converge to a single value.

6 Detection of Well-synchronized Executions

6.1 Motivation

Much research is being done on efficient sequential code distribution techniques (see for instance [16]). Our work originates with the practical problem of evaluating automatically distributed programs.

To automatically parallelize a sequential program for a distributed memory parallel computer, compiling directives must be given. For a data-driven distribution technique, the key directive is to specify a data distribution, that is, to indicate how data structures are to be decomposed and mapped onto the network of processes.

A programmer needs tools that help him to select a good distribution of the data structures of a source program. That is, he must be able to evaluate quantitatively and qualitatively the executions of the distributed code that can be generated given a data decomposition. For instance, he needs to determine the fragments of the source code a data decomposition is unsuited to. For this, a tool that only measures the average degree of concurrency of a distributed execution is clearly inadequate.

To be efficient enough, the tools to be designed should be able to produce relevant outcomes without having to entirely run a generated parallel program. In other words, tools are needed that can collect as much relevant information as possible by efficient static analyses [7] of a source code and of an associated data decomposition.

Semi-automatic distribution is used in application fields (scientific computing) where source codes are generally composed of loops operating on arrays [9]. Actually, available compilers are inefficient when a source program is not regular.

Applying a distribution technique does not affect syntactic regularity. First, generated codes are said to be SPMD (Single Program Multiple Data) [3]: the control structure of each generated process is a copy of that of the source program. Second, data distribution rules are regular as well: arrays, for instance, are decomposed into blocks of contiguous rows or blocks of contiguous columns.

6.2 Detection of Regular Well-synchronized Executions

We deal with programs composed mainly of loops operating on arrays. Data distribution is expressed by a distribution function π that associates with each array element the processor on which it is located, called its owner.

From an intuitive point of view, it is clear that the loops of a parallelized program may lead to regular executions, i.e. the finite repetition of the same distributed order. The idea is to consider as the “basic pattern” the distributed order corresponding to the execution of one step of the loop. Then, an execution is actually regular (in the sense of section 4) if *each* step of the loop produces this pattern.

Let us take as example a program with the following nested loop :

```

for i:=a to b
  for j:=a' to b'
    A[f1(i,j)][g1(i,j)] := h(B[f2(i,j)][g2(i,j)]);

```

This sequential code will be translated into a SPMD code, where each assignment is executed only by the owner of the left-hand-side variable, the right-hand-side variables this processor needs being sent to him by their owner⁵.

Let us assume that the events we observe are only the executions of the assignments. There is repetition of a pattern from one step of the a loop to another if the same events are observed on each processor, as well as the same comparabilities between events. To check this, we have to look at each array reference. For all values of the loop index, a reference must correspond to array elements that are owned by the same processor. On the example, it means that the owner $\pi(A[f_1(i,j)][g_1(i,j)])$ does not depend on the loop index i . Two possibilities arise: the index is not *syntactically* used in the reference, we have in fact $A[f'_1(j)][g'_1(j)]$ (for instance the external loop of the Jacobi relaxation algorithm), or it is syntactically used but not *semantically* (its value has no incidence on the result of the determination of the owners, i.e. on the computation of $\pi(A[f_1(i,j)][g_1(i,j)])$).

Checking this statically is not possible in the general case. However, in practice, the distribution function and the expressions in array references are often affine. Therefore compile-time checking is possible in some cases. Indeed, in most compilers, the distribution of the data (arrays) is made by blocks. With a distribution of an array by columns, for instance, the row index will have no effect on the computation of the owner of a variable. In such a case of block distribution, the distribution function π looks like:

$$\pi(A[f_1(i,j)][g_1(i,j)]) = ((f_1(i,j)/v) * (nbc\!ol/h) + g_1(i,j)/h) \bmod nb\!proc$$

where v (resp. h) is the block size in vertical direction – number of rows – (resp. in horizontal direction – number of columns) and $nbc\!ol$ is the number of columns of array A (see [1]). A distribution by columns would be to set $v = nbc\!ol$, therefore, as $f_1(i,j) < nbc\!ol$, we would have $f_1(i,j)/v = 0$, and if i is not in the expression g_1 (for instance $g_1(i,j) = j$), then $\pi(A[f_1(i,j)][g_1(i,j)])$ would be independent on i .

Thanks to this kind of analysis – that can be done statically because parameters h and v depend on the user-defined data distribution – we have a *sufficient* condition to detect events that are repeated on the same processor for all steps of the loop.

⁵We consider here a very “basic” compilation scheme, without any optimization; this has no influence on the principles we present.

Considering a regular execution, it is easy to detect if it is well-synchronized or not: it suffices to execute one step of the loop, build the communication graph, and check its connectivity.

We have not made an exhaustive study of benchmark programs, but we have found programs whose runtime behaviours are regular and well-synchronized. For example, Jacobi-like programs (walking n times through a matrix, updating each time the values by a function of the neighboring values) satisfy this property whatever the data distribution is, and many linear algebra programs also satisfy it but only for particular data distributions.

7 Conclusion

The contribution we have presented in this paper originates with the practical problem of evaluating the synchronizations of a distributed program running on a network of processors. We are faced with such a problem in the field of automatic parallelization of sequential programs for distributed memory computers (high performance computing).

In this field, generated programs are weakly deterministic [2] and are often control static. Therefore studying one particular execution of a program gives information on the exact quantity of parallelism extracted by the compiler/parallelizer. Another salient feature of the run-time behaviours of these programs is their regularity.

This is why we have felt interested in a concurrency measure that would take regularity into account. We have defined a measure that associates a value with each event of an execution. In the case of a regular and well-synchronized execution, this value remains bounded even if the execution is infinite. This is not the case of other measures from the literature, that in the end associate the same value with all events that occur, although some of them could be extended to take the well-synchronization into account as we have done for Charron-Bost's measure. Our measure is therefore relevant whatever the length of the execution is, and can be computed from the basic pattern of the execution by taking at most N repetitions of this pattern into account, where N is the number of processors.

To obtain this result, we used partial order theory. A distributed execution is modeled as the causality partial order between events. The degree of synchronization is captured by counting the number of antichains that contain a given event.

This theory has proved useful, providing us with an adequate framework to describe the regularity of an execution.

The computation of our measure has been integrated in a parallelization environment developed in our research team [1]. Its exploitation is at the planning stage.

References

- [1] F. André, O. Chéron, and J-L. Pazat. Compiling Sequential Programs for Distributed Memory Parallel Computers with Pandore II. In Jack J. Dongarra and Bernard Touran-

- cheau, editors, *Environments and Tools for Parallel Scientific Computing*, pages 293–308, Elsevier Science Publishers B.V., 1993.
- [2] C. Bateau, B. Caillaud, C. Jard, and R. Thoraval. Correctness of automated distribution of sequential programs. In A. Bode, M. Reeve, and G. Wolf, editors, *PARLE'93*, pages 517–528, LNCS 694, Springer Verlag, June 1993.
 - [3] D. Callahan and K. Kennedy. Compiling programs for distributed-memory multiprocessors. *Journal of Supercomputing*, 2:151–169, 1988.
 - [4] B. Charron-Bost. Combinatorics and Geometry of Consistent Cuts : Application to Concurrency Theory. In Bernard and Raynal, editors, *Int. Workshop on Parallel and Distributed Algorithms*, pages 45–56, Springer Verlag, Nice, France, 1989.
 - [5] B.A. Davey and Priestley H.A. *Introduction to Lattices and Order*. Cambridge University Press, 1990.
 - [6] C. Diehl, C. Jard, and J.X. Rampon. Reachability analysis on distributed executions. In JP. Jouannaud MC. Gaudel, editor, *Proc. TAPSOFT,93 LNCS 668*, pages 629–643, Springer-Verlag, Orsay, Paris, April 1993.
 - [7] T. Fahringer, R. Blasko, and H.P. Zima. Automatic Performance Prediction to Support Parallelization of Fortran Programs for Massively Parallel Systems. In *Proc. of the '92 International Conference on Supercomputing*, pages 347–356, ACM press, July 1992.
 - [8] C.J. Fidge. A simple run-time concurrency measure. In *Proceedings of the 3rd Australian Transputer and OCCAM User Group Conference*, pages 92–101, 1990.
 - [9] G.H. Golub and C.F. Van Loan. *Matrix computations*. The Johns Hopkins University Press, second edition, 1990.
 - [10] M. Habib, M. Morvan, and J.X. Rampon. Remarks on some concurrency measures. In *Graph-Theoretic Concepts in Computer Science*, pages 221–238, LNCS 484, june 1990.
 - [11] C. Jard and J.-M. Jézéquel. ECHIDNA, an Estelle-compiler to prototype protocols on distributed computers. *Concurrency Practice and Experience*, 4(5):377–397, 1992.
 - [12] C. Jard, G.V. Jourdan, and J.X. Rampon. *Some On-Lines Computations of the Ideal Lattice of Posets*. Research Report 773, IRISA, December 1993.
 - [13] L. Lamport. Time, clocks and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.
 - [14] F. Mattern. Virtual time and global states of distributed systems. In Cosnard, Quinton, Raynal, and Robert, editors, *Proc. Int. Workshop on Parallel and Distributed Algorithms Bonas, France, Oct. 1988*, North Holland, 1989.

- [15] M. Raynal, M. Mizuno, and M.-L. Neilsen. A synchronization and concurrency measure for distributed computations. In *12th IEEE Int. Conf. on Distributed Computing Systems*, pages 657–664, Yokokama, June 1992.
- [16] H. Sips. *4th Int. Workshop on Compilers for Parallel Computers*. Sips, H. Editor. Delft, 13-16 december. Delft University of Technology, 1993.



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, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 46 avenue Félix Viallet, 38031 GRENOBLE Cedex 1
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)
ISSN 0249-6399