The chemical abstract machine

Gérard Berry, Gérard Boudol

► To cite this version:

Gérard Berry, Gérard Boudol. The chemical abstract machine. [Research Report] RR-1133, INRIA. 1989. inria-00075426

HAL Id: inria-00075426

<https://hal.inria.fr/inria-00075426>

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.

IRIA

UNITÉ DE RECHERCHE
INRIA-SOPHIA ANTIPOLIS

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Rocquencourt
BP105
78153 Le Chesnay Cedex
France
Tél. (1) 39 63 55 11

Rapports de Recherche

N° 1133

Programme 1
Programmation, Calcul Symbolique
et Intelligence Artificielle

THE CHEMICAL ABSTRACT MACHINE

Gérard BERRY
Gérard BOUDOL

Décembre 1989

The Chemical Abstract Machine

La Machine Abstraite Chimique

Gérard Berry

Ecole des Mines

Sophia-Antipolis

06565-VALBONNE FRANCE

Gérard Boudol

INRIA

Sophia-Antipolis

06565-VALBONNE FRANCE

Abstract.

We introduce a new kind of abstract machine based on the chemical metaphor used in the Γ language of Banâtre & *al.* States of a machine are chemical solutions where floating molecules can interact according to reaction rules. Solutions can be stratified by encapsulating subsolutions within membranes that force reactions to occur locally. We illustrate the use of this model by describing the operational semantics of the TCCS and CCS process calculi. We also show how to extract a higher-order concurrent λ -calculus out of the basic concepts of the chemical machine.

Résumé.

Nous présentons une nouvelle notion de machine abstraite, inspirée du langage Γ de Banâtre & *al.* L'état d'une telle machine peut être conçu comme une solution chimique, où baignent des molécules qui interagissent selon certaines lois de réaction. La solution peut être stratifiée, en isolant certaines parties grâce à une membrane qui contraint les réactions à avoir lieu localement. Le modèle de la machine chimique est utilisé pour décrire la sémantique opérationnelle de calculs de processus, tels TCCS et CCS. Nous montrons également que les idées essentielles de la machine chimique peuvent être formalisées en un calcul parallèle qui généralise le λ -calcul.

The Chemical Abstract Machine

Grard Berry*

Ecole des Mines
Sophia-Antipolis
06560 Valbonne, France

Grard Boudol

INRIA
Sophia-Antipolis
06560 Valbonne, France

Abstract

We introduce a new kind of abstract machine based on the chemical metaphor used in the Γ language of Bantre & al. States of a machine are chemical solutions where floating molecules can interact according to reaction rules. Solutions can be stratified by encapsulating subsolutions within membranes that force reactions to occur locally. We illustrate the use of this model by describing the operational semantics of the TCCS and CCS process calculi. We also show how to extract a higher-order concurrent λ -calculus out of the basic concepts of the chemical abstract machine.

1 Introduction

We present the notion of a *Chemical Abstract Machine*, suited to model asynchronous concurrent computations. We show that chemical abstract machines can "implement" known models of concurrent computation such as algebraic process calculi [18,6] and a concurrent λ -calculus similar to the one presented in [7].

Abstract machines are widely used in the classical theory of sequential computations. Turing Machines or Random Access Machines are primary tools within the theories of recursive functions and computational complexity. The SECD machine [17] and the Categorical Abstract Machine [10] are used to study and implement the λ -calculus, while the SMC machine [19] may be used to describe the semantics of usual imperative constructs.

*presently at LIX, Ecole Polytechnique, 91 128 Palaiseau, France

Work supported by the french Programme de Recherches Coordonnes C³.

The situation is much less clear in the field of concurrent programming. Models such as Petri Nets, Communicating Automata, or Data Flow Networks can be considered as abstract machines, but certainly they lack expressive power. More expressive models such as Algebraic Process Calculi [18,6] are intended to be specification formalisms for distributed systems rather than abstract machines. Implementation models of Concurrent Programming Languages such as CSP [16] are conceptually based on standard sequential machine models augmented with scheduling facilities, not on specific abstract machines.

Most available concurrency models are based on architectural concepts, e.g. networks of processes communicating by means of ports or channels. Such concepts convey a rigid geometrical vision of concurrency. Our chemical abstract machine model is based on a radically different paradigm, which originated in the Γ language of Bantre and Le Metayer [2,3]. These authors pointed out that parallel programming with control threads is more difficult to manage than sequential programming, a fact that contrasts with the common expectation that parallelism should ease program design. They argued that a high-level parallel programming methodology should be liberated from control management. A similar idea motivates the UNITY model of Chandy and Misra [9]. Then they proposed a model where the concurrent components are freely "moving" in the system and communicate when they come in contact.

Intuitively, the state of a system is like a *chemical solution* in which floating molecules can interact with each other according to *reaction rules*; a *magical mechanism* stirs the solution, allowing for possible contacts between molecules. In chemistry, this is the result of Brownian motion, but we don't insist on any particular mechanism, this being an implementation matter not studied here, see [2,9]. The solution transformation process is obviously truly parallel: any number of reactions can be performed in

parallel, provided that they involve disjoint sets of molecules. Notice that the tuple space model of Linda [8] is based on very similar concepts and bears the same degree of potential parallelism, as well as the sets of assignments used in UNITY [9].

Let us give a simple but striking example from [2,3]. Assume the solution is originally made of all integers from 2 to n , along with the rule that any integer destroys its multiples. Then the solution will end up containing the prime numbers between 2 and n . See [2,3] for more examples and for implementation techniques.

Technically, a Γ program is defined by the structure of the molecules it handles and by a set of reaction rules. Solutions are represented by *multisets* of molecules: this accounts for the associativity and commutativity of parallel composition, that is the implicit stirring mechanism. The reaction rules are multiset rewritings.

We keep the same basic notions for chemical abstract machines. We elaborate on the original Γ language by presenting molecules in a systematic way as terms of algebras and refining the classification of rules. Some molecules do not exhibit interaction capabilities; those which are ready to interact are called *ions*. A solution can be *heated* to break complex molecules into smaller ones up to ions. Conversely, a solution can be *cooled* to rebuild heavy molecules from components. Furthermore, to deal with abstraction and hierarchical programming, we allow a molecule to contain a subsolution enclosed in a *membrane*, which can be somewhat *porous* to allow communication between the encapsulated solution and its environment.

The chemical abstract machines all obey a simple set of structural laws. Each particular machine is given by adding a set of simple rules that specify how to produce new molecules from old ones. Unlike the inference rules classically used in structural operational semantics, the specific rules have no premisses and are purely local.

In this paper, we concentrate on the descriptive power of chemical abstract machines. The strength of the model lies in the membrane notion. Membranes make it possible to build chemical abstract machines that have the power of classical process calculi or that behave as concurrent generalizations of the lambda-calculus.

To familiarize the reader with our concepts, the next section presents a simple machine for a subset of CCS. Section 3 gives some formal definitions. In Section 4, we treat the full TCCS [13] calculus and indicate how to handle other process calculi. Section 5 is devoted to a concurrent lambda-calculus similar to that of [7]. We conclude in section 6.

2 Handling a Subset of CCS

Our first illustrative example is a fragment CCS^- of Milner's process calculus CCS [18], containing the most basic operators 0 (inaction), '.' (prefixing), and '|' (parallel), as well as the restriction operator '\ ' to make the example non-trivial.

Let $\mathcal{N} = \{a, b, \dots\}$ be a set of *names* and $\mathcal{L} = \{a, \bar{a} \mid a \in \mathcal{N}\}$ be the set of *labels* built on $\mathcal{N}$. We use the symbols α, β , etc., to range over labels, with $\bar{\alpha} = \alpha$. The CCS^- agents p, q , etc., are given by the syntax:

$$p ::= 0 \mid \alpha.p \mid (p \mid q) \mid p \setminus a$$

2.1 The Semantics of TCCS

Process calculi semantics are usually defined by inference rules in Plotkin's structural operational semantics style, called SOS for short. Milner's original rules involve an additional τ label representing internal communication. This happens to be quite unnatural with respect to abstract machine executions, where internal transitions should not be visible to the user. We prefer to use the De Nicola - Hennessy TCCS rules [13] that define two kinds of transitions between agents: the *internal* transitions $p \rightarrow p'$ and the *labelled* transitions $p \xrightarrow{\alpha} p'$. Intuitively, $p \rightarrow p'$ means that p can become p' by executing an internal action, and $p \xrightarrow{\alpha} p'$ means that p can offer its environment to accept the action α and then become p' .

Both transitions systems are defined in a structural way: the behavior of an agent is deduced from the behaviors of its components. Since internal communications generate internal transitions, the inference systems for $\rightarrow$ invokes the one for $\xrightarrow{\alpha}$:

$$\begin{array}{c} \frac{\alpha.p \xrightarrow{\alpha} p}{p \rightarrow p'} \\ \frac{p \rightarrow p' \quad q \rightarrow q' \quad \text{and} \quad q \mid p \rightarrow q \mid p'}{p \mid q \rightarrow p' \mid q'} \\ \frac{p \xrightarrow{\alpha} p' \quad q \xrightarrow{\bar{\alpha}} q'}{p \mid q \xrightarrow{\alpha} p' \mid q'} \\ \frac{p \xrightarrow{\alpha} p' \quad q \xrightarrow{\bar{\alpha}} q'}{p \mid q \xrightarrow{\alpha} p' \mid q'} \\ \frac{p \xrightarrow{\alpha} p' \quad \alpha \notin \{a, \bar{a}\}}{p \setminus a \xrightarrow{\alpha} p' \setminus a} \end{array}$$

2.2 Basic Chemistry: Concurrency and Communication

We now take the chemical abstract machine point of view, limiting us to internal transitions of restriction-free agents in this section. Restriction and external communication will be treated in the next section.

Instead of composing their behaviors, we make the agents or *molecules* directly react with each other within a *solution*, that is a multiset $S = \{p, q, \dots\}$. There are only two basic rules:

parallel:

$$p|q \rightleftharpoons p, q$$

reaction:

$$a.p, \bar{a}.q \rightarrow p, q$$

The rules apply to molecules present in the solution; they do not apply inside molecules.

The first rule is reversible. It says that any molecule of the form $p|q$ that floats in the solution can be *heated up* (symbol $\rightarrow$) to decompose it into its components p and q , and conversely that any pair p, q of molecules can be *cooled down* (symbol $\leftarrow$) to rebuild a compound molecule $p|q$. The comma ‘,’ appearing in the right-hand side expresses that the heating and cooling rule respectively yield and take a pair of molecules.

The reaction rule deals with *ions*, i.e. molecules of the form $\alpha.p$. Since α is the ion’s communication capability, we call it its *valence*. Whenever two complementary ions float in the solution, they can react with each other and release their bodies in the solution. The valences simply vanish. Unlike the parallel rule, the reaction rule is irreversible.

To execute an agent p , we start from the solution $S_0 = \{p\}$. Heating the solution exhibits the potential communications, which can then be performed using the reaction rule. Notice that a *hot* solution obtained by heating an agent as much as possible contains only ions. Conversely, any solution obtained by transitions from S_0 can be *frozen* by cooling rules into a solution $\{q\}$ consisting of a single CCS⁻ term, or into the empty solution, analogous to the term 0.

To see the chemical abstract machine at work, let us consider an execution of the agent $a.b.0|\bar{a}.0|\bar{b}.0$.

$$\begin{aligned} & \{a.b.0|\bar{a}.0|\bar{b}.0\} \\ \xrightarrow{*} & \{a.b.0, \bar{a}.0, \bar{b}.0\} \quad (\text{parallel}) \\ \rightarrow & \{b.0, 0, \bar{b}.0\} \quad (\text{reaction}) \\ \rightarrow & \{0, 0, 0\} \quad (\text{reaction}) \end{aligned}$$

The final solution $\{0, 0, 0\}$ only contains the inert molecule 0. It is natural to clean it up by using the following rule, which says that 0 evaporates when heated:

inaction cleanup:

$$0 \rightarrow$$

A last cleaning step yields the empty solution $\{\}$.

Generally speaking, chemical executions are non-deterministic. For example, in the solution $\{a.0, \bar{a}.b.0, \bar{a}.c.0\}$, the $a.0$ ion can react with any of the two others ions, yielding either $\{b.0, \bar{a}.c.0\}$ or $\{\bar{a}.b.0, c.0\}$ after cleanup.

The reader will appreciate the simplicity of the chemical executions compared to the sequence of proofs and simplifications involved in the SOS semantics. The use of the structural rules for ‘|’ is factored throughout an execution by the heating process, since we directly chain reactions by keeping the solution hot. The SOS evaluation involves structural rules at each computation step.

In fact, the simplification comes from the abandon of the rigid algebraic syntax. Chemical concurrency is *naturally* associative and commutative, since multisets are intrinsically unordered. The notion of syntactic position disappears even for the standard syntactic parallel construct ‘|’: it is impossible to know whether $\{p, q\}$ was obtained by heating $\{p|q\}$ or $\{q|p\}$. On the contrary, the SOS semantics need to first introduce behaviors to recover concurrency out of the fixed syntax, then to define what it means for processes to be equivalent, and finally to *prove* equivalences such as $p|q \sim q|p$. SOS also involves inference rules with non-trivial premisses, which are certainly more complex than the naive cham rewrite rules.

Furthermore, we treat structural simplifications in the same way as reactions: to suppress a 0, we simply evaporate it. In SOS semantics, one needs to prove that $p|0$ is equivalent to p , and one performs transitions and simplifications in separate steps and by separate techniques.

However, the notion of behavior that underlies the SOS semantics has advantages. In particular, $p \xrightarrow{\alpha} p'$ also tells that p can perform an α when requested by some *external* observer, thus defining at once how an agent communicates with its environment. We must define the same notion in our setting.

A solution should be able to perform a visible α action whenever it contains an ion $\alpha.p$. This ion should then export the α valence and become p . One could imagine to let it disintegrate into p and emit an α -particle to the environment. However, such a technique would violate Milner’s most useful principle,

which states that there should be no difference in nature between internal and external communications. The right solution is to make the observer *react* with the valence of any molecule of the solution. This requires a richer machinery developed in the next section.

2.3 More Advanced Chemistry: Membranes and Airlocks

Consider a restriction agent $p \backslash a$ floating in a solution. If p is already of the form $\alpha.q$, $\alpha \notin \{a, \bar{a}\}$, we can build a new ion by the following simple rule:

restriction ion:

$$(\alpha.p) \backslash a \Rightarrow \alpha.(p \backslash a) \quad \text{if } \alpha \notin \{a, \bar{a}\}$$

But this does not work if p is compound. In this case, p should be able to freely perform internal reactions and to also propose communications to other ions floating in the main solution, using its own ions of unrestricted valences. To let p evolve on its own, we put it in a new local solution contained within a *membrane* $\{\cdot\}$:

restriction membrane:

$$p \backslash a \Rightarrow \{p\} \backslash a$$

The new subsolution obeys the same rules as the global solution. To realize global communications, we need to make the membrane *porous* to valences. A first simple idea would be to use an heavy ion formation rule such as:

$$\{\alpha.p, p_1, p_2, \dots, p_n\} \rightarrow \alpha.\{p, p_1, p_2, \dots, p_n\}$$

We reject such a rule for two reasons. First, it does not involve only simple molecules as did previous rules; on the contrary, it involves finding an ion within a solution, which is neither simple nor general. Second, it is irreversible, since the information of where α comes from is lost. If a wrong valence is chosen, the heavy ion can stay forever in the main solution, like a *precipitate*. Consider for example $\{a.0, \{\bar{a}.0, b.0\} \backslash c\}$ when choosing b : we are stuck with the inert solution $\{a.0, b.(\{\bar{a}.0, 0\} \backslash c)\}$.

The technique we propose is in two steps. First, we introduce a new mechanism at the general chemical machine level: the *airlock* mechanism. It extracts any molecule from a solution (not necessarily an ion), puts the rest of the solution within a membrane, and isolates the extracted molecule within an airlock attached to the membrane. The airlock construct is written $m \triangleleft S$ where m is the isolated molecule and S

the remaining solution; it is a *single* molecule. The reversible (meta) rule is:

airlock:

$$\{m, m_1, m_2, \dots, m_n\} \Rightarrow \{m \triangleleft \{m_1, m_2, \dots, m_n\}\}$$

The solution $\{m_1, m_2, \dots, m_n\}$ contained in the new molecule is allowed to freely continue internal reactions (see the formal definitions in the next section).

Second, we build an heavy ion from any ion in the airlock, using the rule:

heavy ion:

$$(\alpha.p) \triangleleft S \Rightarrow \alpha.(p \triangleleft S)$$

In this way, we guarantee reversibility by preserving the attachment between α and p . A restriction molecule can propose several valences in succession to its environment until a communication takes place.

Let us give a simple example of communication involving an heavy ion:

$$\begin{array}{ll} \{a.0 \mid (\bar{a}.p \mid q) \backslash b\} & \\ \xrightarrow{*} \{a.0, \{\bar{a}.p, q\} \backslash b\} & (\text{par., restr. membr.}) \\ \rightarrow \{a.0, \{\bar{a}.p \triangleleft \{q\}\} \backslash b\} & (\text{airlock}) \\ \rightarrow \{a.0, \{\bar{a}.(p \triangleleft \{q\})\} \backslash b\} & (\text{heavy ion}) \\ \rightarrow \{a.0, (\bar{a}.(p \triangleleft \{q\})) \backslash b\} & (\text{restr. membrane}) \\ \rightarrow \{a.0, \bar{a}.((p \triangleleft \{q\}) \backslash b)\} & (\text{restriction ion}) \\ \rightarrow \{0, (p \triangleleft \{q\}) \backslash b\} & (\text{reaction}) \\ \rightarrow \{\{p \triangleleft \{q\}\} \backslash c\} & (\text{restr. membrane}) \\ \rightarrow \{\{p, q\} \backslash c\} & (\text{airlock}) \end{array}$$

Notice how we guarantee reversibility by putting or removing membranes. Once the heavy ion $\bar{a}.(p \triangleleft \{q\})$ has been constructed, it is not possible to put back p into q 's solution, since the airlock molecule is not contained within a membrane. It is not possible to build such a membrane, since the membrane rule cannot be applied inside an ion.

The airlock technique makes it now easy to define what it means for an external observer to *observe* a solution. If the solution is reduced to a single ion, then the observer can pick up the ion's valence and release its body. More precisely, let S, S' denote solutions. We set $S \xrightarrow{a} S'$ if there exist a molecule m such that $S \xrightarrow{*} \{\alpha.m\}$ and $\{m\} \xrightarrow{*} S'$. For example, one has

$$\{a.0, b.0\} \xrightarrow{a} \{b.0\}$$

taking $m = 0 \triangleleft \{b.0\}$.

The relation between this new kind of behavior and the standard TCCS one will be precisely stated in section 4.

3 Formal Definitions

3.1 Chemical Abstract Machines

A *chemical abstract machine* or *cham* C is specified by defining *molecules* $m, m', \text{etc.}$, *solutions* $S, S', \text{etc.}$, and *transformation rules*. Molecules are terms of algebras, with specific operations for each cham. Solutions are finite multisets of molecules, written $\{m_1, m_2, \dots, m_k\}$. Furthermore, in each cham, any solution S can itself be considered as a single molecule and can therefore appear as a *subsolution* of another molecule. The corresponding $\{\cdot\}$ operator is called the *membrane* operator. For instance, if 0 and $+$ are the molecule building operations, then 0 , $0 + 0$, $0 + \{0\}$, and $\{0, \{0 + 0, 0\}\}$ are molecules, the latter also being a solution.

The transformation rules have the form

$$m_1, m_2, \dots, m_k \rightarrow m'_1, m'_2, \dots, m'_l$$

where the m_i and m'_j are molecules. As usual, the rules will be presented by means of *rule schemata*, the actual rules being the instances of these schemata. To avoid “multiset matching”, we require the subsolutions appearing in rule schemata to be either a single solution variable S that generates all solutions, or of the form $\{m\}$ where m is a single molecule.

To state how reaction rules apply to solutions, we need some notation. The multiset union of S and S' is written $S \uplus S'$. As in the λ -calculus [4], we use the context notation $C[\cdot]$ to denote a molecule with a hole $[\cdot]$ in which to place other molecule.

The transformation rules determine a *transformation relation* $S \rightarrow S'$ between solutions, according to the following three laws:

- **The Reaction Law.** An instance of the right-hand-side of a rule can replace the corresponding instance of its left-hand-side. If

$$m_1, m_2, \dots, m_k \rightarrow m'_1, m'_2, \dots, m'_l$$

is a rule and $M_1, M_2, \dots, M_k, M'_1, M'_2, \dots, M'_l$ are instances of the m_i 's and the m_j 's, then

$$\{M_1, M_2, \dots, M_k\} \rightarrow \{M'_1, M'_2, \dots, M'_l\}$$

- **The Chemical Law.** Reactions can be performed freely within any solution:

$$\frac{S \rightarrow S'}{S \uplus S'' \rightarrow S' \uplus S''}$$

- **The Membrane Law.** A subsolution can evolve freely in any context:

$$\frac{S \rightarrow S'}{\{C[S]\} \rightarrow \{C[S']\}}$$

The chemical and membrane laws are the only ones to involve premisses. They factor what is usually called “structural rules” in particular calculi. All other laws and rules are purely local.

Some chams, but not all of them, use the additional *airlock* construct. An airlock is a molecule of the form $m \triangleleft S$ where m is a molecule and S is a solution. Airlocks are built and suppressed by the following reversible law:

- **The Airlock Law**

$$\{m\} \uplus S \leftrightarrow \{m \triangleleft S\}$$

A cham is an intrinsically parallel machine: one can simultaneously apply several rules to a solution provided that no molecule is involved in more than one rule; one can also transform subsolutions in parallel. In this paper, we only study the expressive power of chams; it does not depend on using parallel evaluation, since a non-conflicting parallel application of rules is equivalent to any sequence of the individual rules. See [2] for a practical use of parallel reductions.

3.2 A Classification of Rules

We usually distinguish between three kinds of rules: *heating rules* $\rightarrow$, *cooling rules* $\leftarrow$, and *reaction rules* $\rightarrow$. The distinction is merely a matter of taste. As a rule of thumb, we present all structural rules as heating rules, possibly paired with inverse cooling rules. Heating rules decompose a single molecule into simpler ones, and cooling rules recompose a compound molecule from its components. We generally write the heating and cooling rules together, using the symbol $\rightleftharpoons$. We say that a solution is *hot* (resp. *frozen*) if no heating (resp. cooling) rule applies to it. In the sequel, we shall always consider that the transitions given by the airlock law are heating and cooling ones.

The reflexive, symmetric, and transitive closure of $(\rightarrow \cup \leftarrow)$ is written $\rightleftharpoons^*$. According to our conventions, it usually represents structural equivalence.

The reaction rules usually involve several molecules. These molecules are often in a particular form in which they cannot be heated further; we call them *ions*. A solution is *inert* if no reaction rule applies to it, nor to any solution obtained by heating it.

4 Process Calculi Chams

In this section, we finish the treatment of the TCCS process calculus, we establish the semantical equivalence between the cham semantics and the original structural semantics, and we briefly indicate how to handle other process calculi.

4.1 The Full TCCS Calculus

We finish the description of the TCCS calculus [13] and of its SOS semantics. We have already seen the inaction '0', parallel '|', prefixing '.', and restriction '\ ' operators. We now add the remaining operators: the relabelling operator '[.]', the two sum operators '⊕' (internal sum) and '||' (external sum), and the fixpoint definition $\text{fix}_i(\vec{x} = \vec{p})$, which is a shorthand for:

$$\text{letrec } x_1 = p_1 \text{ and } \dots \text{ and } x_n = p_n \text{ in } x_i$$

The final syntax is as follows:

$$p ::= 0 \mid \alpha.p \mid (p \mid q) \mid p \backslash a \mid p[\phi] \\ \mid p \oplus q \mid p \parallel q \mid \text{fix}_i(\vec{x} = \vec{p})$$

We give the semantics of the new operators. A relabelling is a mapping $\phi : \mathcal{M} \mapsto \mathcal{L}$, extended to labels by setting $\phi(\bar{\alpha}) = \phi(\alpha)$. The relabelling operator takes an agent p and a relabelling ϕ and produces a new agent $p[\phi]$ that behaves like p except that all its visible actions are relabelled by ϕ :

$$\frac{p \rightarrow p'}{p[\phi] \rightarrow p'[\phi]} \\ \frac{p \xrightarrow{\alpha} p'}{p[\phi] \xrightarrow{\phi(\alpha)} p'[\phi]}$$

Sums represent non-deterministic choices. There are several possible sums, see [13] for an extensive discussion. The simplest sum is the internal sum $\oplus$, which non-deterministically chooses a component:

$$p \oplus q \rightarrow p \\ p \oplus q \rightarrow q$$

In an external sum $p \parallel q$, the agents p and q can freely perform internal actions and can also propose communications to the environment. The choice is made only when such a communication is performed:

$$\frac{p \rightarrow p' \quad p \parallel q \rightarrow p' \parallel q \text{ and } q \parallel p \rightarrow q \parallel p'}{p \parallel q \xrightarrow{\alpha} p' \text{ and } q \parallel p \xrightarrow{\alpha} p'}$$

Finally, the fixpoint operation is a simple unfolding. Let $p[\vec{q}/\vec{x}]$ denote the result of the simultaneous substitution of the q_i to the x_i in p :

$$\text{fix}_i(\vec{x} = \vec{p}) \rightarrow p_i[\text{fix}_i(\vec{x} = \vec{p})/\vec{x}]$$

4.2 Handling the New Operators

We first explain how to handle the new operators. Then we give the exact syntax of molecules and the complete set of rules of the TCCS cham.

The relabelling operator can be handled just as the restriction operator, by building a membrane and exporting relabelled names as heavy ion valences. Internal sum is handled by the same rules as in the SOS semantics; since the rules are not structural, we call them reaction rules and not heating rules. The fixpoint expansion rule is also as in the SOS semantics, and it is clearly a heating rule. See the exact rules in the rule summary below.

External sum needs more care. Since the summands p and q should each be able to freely perform internal transitions, we open one membrane for each of them. We therefore introduce a new *molecule pairing* operator $\langle ., . \rangle$ and the expansion rule:

$$p \parallel q \rightleftharpoons \langle \{p\}, \{q\} \rangle$$

Assume that the left subsolution S produces an ion $\alpha.m$ and that we want to export α . Then we can give S the form $\{\alpha.n\}$, either by taking $n = m$ if S only contains the given ion, or by building an airlock $(\alpha.m) \triangleleft S_1$ and then an heavy ion $\alpha.(m \triangleleft S_1)$. To export the valence, we can use the rule:

$$\langle \{\alpha.n\}, S' \rangle \rightarrow \alpha. \langle n, S' \rangle$$

However, as discussed in section 2, we must be careful to avoid precipitates and to make the above rule reversible. The reverse cooling rule must recognize that the valence belongs to n and not to S' when it is given a pair $\langle n, S' \rangle$. Furthermore, once the α valence is consumed by some reaction, we are left with a pair $\langle n, S' \rangle$ that we must transform into n to realize the summand selection; we need a cooling rule of the form:

$$\langle n, S' \rangle \rightarrow n$$

Here again, we must recognize which is n and which is S' .

We can use several techniques to solve this problem. The one we choose is to tag the internal pair when the heavy ion is built to directly remember the valence attachment. The rules for the left-hand-side choice are:

$$\langle \{\alpha.m\}, S \rangle \rightleftharpoons \alpha.(l : \langle m, S \rangle) \\ l : \langle m, m' \rangle \rightarrow m$$

The rules for the right-hand-side choice are symmetric with a label r .

An alternative technique would be to always force n to be an airlock, noticing that a solution $\{\alpha.m\}$ containing a single ion can always be heated into $\{\alpha.(m \triangleleft \{\})\}$, and to use the following rules:

$$\begin{aligned} \langle \{\alpha.(m \triangleleft S)\}, S' \rangle &\rightleftharpoons \alpha. \langle m \triangleleft S, S' \rangle \\ \langle m \triangleleft S, S' \rangle &\longrightarrow m \triangleleft S \end{aligned}$$

This technique would decrease the number of molecule constructors, but is rather *ad-hoc* and we don't use it here. This discussion might look a bit tedious, but it shows two things: first, designing a cham has much to do with standard programming; second, the external sum operator is not very natural in concurrent abstract machines.

Let us give a simple external sum evaluation example:

$$\begin{aligned} &\{\{a.\bar{b}.0 \mid ((\bar{a}.0 \mid b.0) \parallel q)\}\} \\ \xrightarrow{*} &\{\{a.\bar{b}.0, \langle \{\bar{a}.0, b.0\}, \{q\} \rangle\}\} \\ \xrightarrow{*} &\{\{a.\bar{b}.0, \langle \{\bar{a}.(0 \triangleleft \{b.0\}), \{q\} \rangle\}\} \\ \xrightarrow{*} &\{\{a.\bar{b}.0, \bar{a}.l : \langle 0 \triangleleft \{b.0\}, \{q\} \rangle\}\} \\ \rightarrow &\{\{\bar{b}.0, l : \langle 0 \triangleleft \{b.0\}, \{q\} \rangle\}\} \\ \rightarrow &\{\{\bar{b}.0, 0 \triangleleft \{b.0\}\}\} \\ \rightarrow &\{\{\bar{b}.0, 0, b.0\}\} \end{aligned}$$

Notice the last step: when an airlock $m \triangleleft S$ floats in a solution, one can cool it down and release m and all the molecules of S in the containing solution. This requires to use both the airlock law and the chemical law.

4.3 The Complete TCCS Cham

We summarize the syntax and rules of the final TCCS cham.

Syntax

Agents:

$$\begin{aligned} p ::= & 0 \mid \alpha.p \mid (p \mid q) \mid p \setminus a \mid p[\phi] \\ & \mid p \oplus q \mid p \parallel q \mid \text{fix}_i(\bar{x} = \bar{p}) \end{aligned}$$

Molecules:

$$\begin{aligned} m ::= & p \mid \alpha.m \mid m \setminus a \mid m[\phi] \mid S \mid m \triangleleft S \\ & \mid \langle m, m \rangle \mid l : \langle m, m \rangle \mid r : \langle m, m \rangle \end{aligned}$$

Notice that parallel and sums are agent operators, not molecules operators.

Rules

parallel:

$$p \mid q \rightleftharpoons p, q$$

reaction:

$$a.m, \bar{a}.n \rightarrow m, n$$

restriction membrane:

$$m \setminus a \rightleftharpoons \{\{m\} \setminus a$$

restriction ion:

$$(\alpha.m) \setminus a \rightleftharpoons \alpha.(m \setminus a) \quad \text{if } \alpha \notin \{a, \bar{a}\}$$

relabelling membrane:

$$m[\phi] \rightleftharpoons \{\{m\}[\phi]$$

relabelling ion:

$$(\alpha.m)[\phi] \rightleftharpoons \phi(\alpha).(m[\phi])$$

$\oplus$ -left:

$$p \oplus q \rightarrow p$$

$\oplus$ -right:

$$p \oplus q \rightarrow q$$

$\parallel$ -expansion:

$$p \parallel q \rightleftharpoons \langle \{p\}, \{q\} \rangle$$

left $\parallel$ -ion:

$$\langle \{\alpha.m\}, S \rangle \longrightarrow \alpha.l : \langle m, S \rangle$$

right $\parallel$ -ion:

$$\langle S, \{\alpha.m\} \rangle \longrightarrow \alpha.r : \langle S, m \rangle$$

left projection:

$$l : \langle m, m' \rangle \longrightarrow m$$

right projection:

$$r : \langle m, m' \rangle \longrightarrow m'$$

fixpoint:

$$\text{fix}_i(\bar{x} = \bar{p}) \rightarrow p_i[\text{fix}(\bar{x} = \bar{p})/\bar{x}]$$

Additional Cleanup Rules

inaction cleanup:

$$0 \longrightarrow$$

restriction cleanup:

$$\{\{ \} \setminus a \longrightarrow$$

relabelling cleanup:

$$\{\{ \}[\phi] \longrightarrow$$

4.4 Comparing the Cham and SOS

We define labeled transitions as explained in section 2.

Definition: Given a solution S , we write $S \xrightarrow{\alpha} S'$ if there exists a molecule m' such that $S \xrightarrow{*} \{\alpha.m'\}$ and $\{\alpha.m'\} \xrightarrow{*} S'$.

Remember that the structural equivalence $\equiv$ between solutions is the reflexive, symmetric, and transitive closure of the heating and cooling relations. In the sequel, we shall neglect the cleanup rules and consider only the reversible heating/cooling rules. Then $S \equiv S'$ if and only if there exists a sequence of heating or cooling steps from S to S' .

The following result shows that the cham differs from the original TCCS calculus only in the number of internal steps involved in computations. As far as labelled transitions are concerned, the solution $\{p\}$ can do whatever the term p can do, and it cannot do more.

Theorem: Let p be a TCCS agent.

1) If $p \rightarrow p'$ in TCCS, then $\{p\} \xrightarrow{*} \{p'\}$ in the TCCS cham. If $p \xrightarrow{\alpha} p'$ in TCCS, then $\{p\} \xrightarrow{\alpha} \{p'\}$; more precisely, there exists a molecule m' such that $\{p\} \rightarrow \{\alpha.m'\}$ and $\{\alpha.m'\} \xrightarrow{*} \{p'\}$.

2) If $\{p\} \rightarrow S'$, then there exists a TCCS agent p' such that $S' \equiv \{p'\}$. If $\{p\} \xrightarrow{\alpha} S'$, then there exists a TCCS agent p' such that $p \xrightarrow{\alpha} p'$ and $S' \equiv \{p'\}$.

Sketch of proof: To prove 1), one shows how to perform given TCCS derivations by chaining cham transitions. The proof is by induction on the size of p and by cases on the form of the given TCCS transition. We show two typical cases.

First if $p = p_1 | p_2 \rightarrow p'_1 | p'_2 = p'$ with $p_1 \xrightarrow{\alpha} p'_1$ and $p_2 \xrightarrow{\beta} p'_2$ for some α, β , by induction, there exist m'_1 and m'_2 such that $\{p_1\} \xrightarrow{*} \{\alpha.m'_1\}$, $\{\alpha.m'_1\} \xrightarrow{*} \{p'_1\}$, $\{p_2\} \xrightarrow{*} \{\beta.m'_2\}$, and $\{\beta.m'_2\} \xrightarrow{*} \{p'_2\}$. We build the following transformation sequence:

$$\begin{aligned}
 & \{p\} \\
 = & \{p_1 | p_2\} \\
 \rightarrow & \{p_1, p_2\} \quad (\text{parallel}) \\
 \xrightarrow{*} & \{\alpha.m'_1, \beta.m'_2\} \quad (\text{cham laws}) \\
 \rightarrow & \{m'_1, m'_2\} \quad (\text{reaction}) \\
 \xrightarrow{*} & \{p'_1, p'_2\} \quad (\text{cham laws}) \\
 \rightarrow & \{p'_1 | p'_2\} \quad (\text{parallel}) \\
 = & \{p'\}
 \end{aligned}$$

which shows the required property of p .

Assume now $p = q \setminus a \xrightarrow{\alpha} q' \setminus a = p'$, with $a \notin \{a, \bar{a}\}$. By induction, there exists n' such that $\{q\} \rightarrow \{\alpha.n'\}$ and $\{n'\} \xrightarrow{*} \{q'\}$. Let $m' = n' \setminus a$. We build the execution sequence:

$$\begin{aligned}
 & \{p\} \\
 = & \{q \setminus a\} \\
 \rightarrow & \{\{q\} \setminus a\} \quad (\text{restriction membrane}) \\
 \xrightarrow{*} & \{\{\alpha.n'\} \setminus a\} \quad (\text{cham laws}) \\
 \rightarrow & \{(\alpha.n') \setminus a\} \quad (\text{restriction membrane}) \\
 \rightarrow & \{\alpha.(n' \setminus a)\} \quad (\text{restriction ion}) \\
 = & \{\alpha.m'\}
 \end{aligned}$$

Furthermore, one has:

$$\begin{aligned}
 & \{m'\} \\
 = & \{n' \setminus a\} \\
 \rightarrow & \{\{n'\} \setminus a\} \quad (\text{restriction membrane}) \\
 \xrightarrow{*} & \{\{q'\} \setminus a\} \quad (\text{cham laws}) \\
 \rightarrow & \{q' \setminus a\} \quad (\text{restriction membrane}) \\
 = & \{p'\}
 \end{aligned}$$

which shows the required property of p .

Proving 2) is harder and we just sketch of the proof architecture. The properties of $\rightarrow$ and $\xrightarrow{\alpha}$ are proved together by induction on the number of irreversible rules applied in the given derivations.

If this number is 0, then the property of $\xrightarrow{\alpha}$ is obvious with $p' = p$. To prove the property of $\xrightarrow{\alpha}$, we use a lemma about ion formation.

The lemma shows how ions $a.m$ can be formed in arbitrary subsolutions using only heating and cooling rules. The valences of such ions always come from label positions in TCCS terms that yield observable transitions. Furthermore, the ion bodies are kept untouched in the heating-cooling process. More formally, assume $\{p\} \xrightarrow{*} C[\alpha.m]$ where the ion $\alpha.m$ floats in some subsolution. Then one can cool down $C[\alpha.m]$ into $C'[\alpha.q]$ by using only ion cooling rules, in such a way that $\alpha.q$ is exactly a subexpression of the original agent $p = C_1[\alpha.q]$, with the additional properties $p \xrightarrow{\alpha} C_1[q]$ in TCCS and $\{C_1[q]\} \xrightarrow{*} \{C'[q]\}$.

Now if $\{p\} \xrightarrow{*} \{\alpha.m'\}$, one can use the lemma to show that p has the form $C[\alpha.q]$ with $C[q] \xrightarrow{*} m'$. This shows the required property of $\xrightarrow{\alpha}$, taking $p' = C[q]$.

Assume now that the number of irreversible transitions in a given derivation is strictly positive. The

derivation can be put in the form $S \xrightarrow{*} S_1 \rightarrow S_2 \xrightarrow{*} S'$ with $S \xrightarrow{*} S_1$ and where $S_1 \rightarrow S_2$ is irreversible. The only difficult case is the one where the transition from S_1 to S_2 is a reaction. By a slight extension of the lemma to two-hole contexts, one can show that $p = C[\alpha.q][\bar{\alpha}.r]$, $p \rightarrow p' = C[q][r]$ in TCCS, and $\{p'\} \xrightarrow{*} C'[q][r]$ with $S_1 \xrightarrow{*} C[\alpha.q][\bar{\alpha}.r]$ and $S_2 \xrightarrow{*} C[q][r]$. The global induction hypothesis applies to p' and gives the final result.

4.5 Handling Other Process Calculi

In Milner's original calculus CCS, there is no notion of an internal unlabelled transition. The special label τ is used to report transitions provoked by internal communications. The sum $p + q$ is defined by the following rule: τ :

$$\frac{p \xrightarrow{\alpha} p'}{p + q \xrightarrow{\alpha} p' \text{ and } q + p \xrightarrow{\alpha} p'}$$

in which one can take $\alpha = \tau$. Therefore, a summand can be chosen either by an external communication or by an internal one.

To simulate CCS by a cham, we abandon the simple reaction rule of TCCS and replace it by the following rule:

τ -reaction:

$$a.m, \bar{a}.n \rightarrow \tau.(m|n)$$

Since a τ -ion can neither be heated nor interact with another molecule, the only thing it can do is to traverse all membranes up to the external observer. An observation by this observer consumes the τ valence, and frees the ion body that can be heated to release the parallel components. With this new definition of reaction, the rules of $+$ are simply the above rules of $\parallel$.

Notice that performing an internal communication is more than just building a τ : the communication is really performed only when the final observer accepts it by *consuming* this τ . Therefore, the machine's behavior can no longer be defined independently of the observation process. Furthermore, the τ -reaction rule reduces the potential parallelism of the execution machine to a bare minimum. The simulation of CCS is rather unsatisfactory. We don't believe that CCS can be "implemented" in a more natural way, which is an indication that τ and $+$ might not be good *programming* primitives. This is quite well-known to CCS simulator implementors.

Given a CCS agent p , we can show that the solution $\{p\}$ is in *weak bisimulation* with the agent p w.r.t. our definition of observation. Strong bisimulation can also be obtained by making the τ -reaction

rule reversible, that is by allowing the machine to *undo* all its internal operations. The details are beyond the scope of this paper.

Handling other process calculi raises no particular problem. For example, the reader can easily write a natural cham for MEIJE [6], which is universal among the labelled process calculi. This tells that the cham formalism have basically the same power as the SOS one w.r.t. process calculi.

5 A Concurrent λ -calculus

5.1 Generalizing the λ -calculus

Algebraic process calculi model concurrency but have a limited expressive power compared to the λ -calculus, where one is able to express all possible combinators and to code many types of data. On the other hand, the λ -calculus is intrinsically sequential [4,5] and cannot handle even the weakest form of concurrency. Building new calculi that combine both abilities is a goal of primary importance [7,21]. In [7], we introduced such a tentative concurrent lambda-calculus called the γ -calculus. We could describe the (lazy) evaluation in this calculus by means of a cham. However, our formalism itself suggests a simpler, and perhaps better calculus of the same kind. To introduce this new calculus, let us first say a few words about the λ - and γ -calculi. Some familiarity with the λ -calculus will be assumed. We just recall the syntax:

$$M ::= x \mid (\lambda x.M) \mid (MM)$$

where x stands for any variable. We are interested here in the *lazy* evaluation of λ -terms (following [1]), that is the reflexive and transitive closure of the relation $M \triangleright M'$ inductively given by

$$(\lambda x.M)N \triangleright M[N/x]$$

$$M \triangleright M' \Rightarrow MN \triangleright M'N$$

Intuitively, a λ -calculus redex $(\lambda x.M)N$ is like a valued CCS communication of the form $\lambda x.M \mid \bar{\lambda}(N)$, since both yield $M[N/x]$ as a result. Hence one could imagine treating the lambda-calculus as a CCS-like process calculus where agents are communicable values, λ becoming a particular label. In such a calculus, functional application should appear as a particular parallel combination of two agents, the function and its argument, and β -reduction should be just a particular case of communication. However, the above simple redex translation would not take care of the non-associative character of application and would not treat double applications correctly. Consider, for

instance, the λ -term $((\lambda x.\lambda y.M)N)P$. The translation would be $\lambda x.\lambda y.M | \bar{\lambda}(N) | \bar{\lambda}(P)$. The associative/commutative character of concurrency would make the arguments N and P interchangeable, which is clearly wrong. Thomsen solved this problem in [21] using the CCS operators of restriction and renaming. However in his higher order calculus, β -reduction is performed in two steps, involving an intermediary state which does not represent a λ -term. Then the λ -calculus is not exactly a sub-calculus of Thomsen's CHOCS calculus.

Another solution was presented in [7] using two concurrency operators: an interleaving operator ' $|$ ' and a binary communication operator ' $\odot$ '. Communications arise as follows: in a term $(M \odot N)$, all ' $|$ ' concurrent components of M can communicate with all concurrent components of N , up to termination of M or N , termination being written as a special symbol 1. Then the $\odot$ operator disappears by application of the simplification rule $(M \odot 1) = (1 \odot M) = M$, and λ -application can be represented by $(M \odot \bar{\lambda}(N))$. For instance, the above double application works in the following way (assuming x, y not free in N):

$$\begin{aligned} & ((\lambda x.\lambda y.M \odot \bar{\lambda}(N)) \odot \bar{\lambda}(P)) \\ \rightarrow & ((\lambda y.M[N/x] \odot 1) \odot \bar{\lambda}(P)) \quad (\text{communication}) \\ = & ((\lambda y.M[N/x]) \odot \bar{\lambda}(P)) \quad (\text{simplification}) \\ \rightarrow & (M[N/x][P/y] \odot 1) \quad (\text{communication}) \\ = & M[N/x][P/y] \quad (\text{simplification}) \end{aligned}$$

A cham describing this calculus would treat the terms $\lambda x.M$ and $\bar{\lambda}(N)$ as ions, but the interpretation of the concurrency operators of this calculus would be somewhat unnatural. In a cham, the parallelism is always commutative and associative and allows for communication, while $(M | N)$ disallows communication and $\odot$ is non-associative. As a matter of fact, the cham framework indicates another possibility for representing properly the λ -application, by means of an encapsulated parallel combination of the function and its argument.

5.2 The γ -calculus

The key idea of our new higher-order concurrent calculus is to *internalize* the concepts of the chemical abstract machine within the syntax. Let us review these concepts:

- *solutions*: these are built by heating a parallel combination of molecules. Therefore the corresponding syntactic construct is parallel composition $(M | N)$. Since solutions are multisets of possibly interacting processes, this operator allows communication.

- *membrane*: encapsulating a subsolution within a membrane forces reactions to occur locally. Here we will introduce a corresponding *localization* construct $\langle M \rangle$.
- *reactions*: basically, these occur when opposite ions float inside the same solution. We shall distinguish two kinds of reactive molecules, the *negative* ones, or receptors, and the *positive* ones, or emitters.

Typically, a receptor in the λ -calculus is an abstraction $\lambda x.M$. To emphasize the ion character, we shall denote such an atomic receptor x^-M , and an atomic emitter sending the value M will be denoted M^+ . Therefore the syntax of our calculus is:

$$M ::= x | x^-M | (M)^+ | (M | M) | \langle M \rangle$$

where x stands for any variable. As usual we shall omit (or add) some parentheses in writing the terms, which will be called processes or sometimes agents. In what follows we shall call this concurrent calculus the γ -calculus, superseding the one proposed in [7].

To formalize the execution mechanism, we need a syntactic notion of *stable* state. Basically, a stable term is made out of ions of the same valence (either positive or negative), and will therefore represent an inert solution. Formally, the syntax for pure emitters or receptors and for stable terms is given by:

$$\begin{aligned} E & ::= M^+ | (E | E) | \langle E \rangle \\ R & ::= x^-M | (R | R) | \langle R \rangle \\ W & ::= E | R \end{aligned}$$

Now we give the γ -cham describing the (lazy) evaluation of terms. The molecules are given by the following grammar:

$$U ::= M | S | (U | U) | \langle U \rangle$$

where M stands for any term and S for any solution (i.e. finite multiset of molecules). The transformation rules are:

solution:

$$U | V \Rightarrow U, V$$

membrane:

$$\langle U \rangle \Rightarrow \{U\}$$

hatching:

$$\langle W \rangle \Rightarrow W$$

β -reaction:

$$x^-M, N^+ \rightarrow M[N/x]$$

where U, V stand for molecules, M, N for terms and W for any stable term. Note that the reaction rule, which is the only irreversible rule, embodies communication. The power of the calculus is essentially due to the rules concerning the membrane construct. This should not be confused with CCS restriction: if a membrane encloses a stable state (i.e. emitter or receptor), then it may vanish. The hatching rule conveniently replaces the termination equations concerning the cooperation operator of [7] (in our calculus, a “cooperation” operator would be $\langle M | N \rangle$). In what follows we shall use the notation $M \xrightarrow{*} N$ as an abbreviation for $\{M\} \xrightarrow{*} \{N\}$.

This γ -calculus contains the λ -calculus, since we may now define the application (MN) as the combination $\langle M | N^+ \rangle$. Let us see this point in some detail; we define a translation θ from the set of λ -terms to the set of terms given by the grammar:

$$M ::= x \mid x^-M \mid \langle M | M^+ \rangle$$

The translation is as follows:

$$\theta(x) = x$$

$$\theta(\lambda x.M) = x^- \theta(M)$$

$$\theta(MN) = \langle \theta(M) | \theta(N^+) \rangle$$

Then we may show that there is a close correspondence between lazy evaluation of λ -terms and evaluation in the γ -cham of their translation. More precisely, it is easy to prove that

$$M \triangleright^* M' \Leftrightarrow \theta(M) \xrightarrow{*} \theta(M')$$

and, moreover, that each intermediate state in the evaluation of $\theta(M)$ cools down to a λ -term. For instance, the above double application works as follow:

$$\begin{aligned} & \langle \langle x^-y^-M | N^+ \rangle | P^+ \rangle \\ \xrightarrow{*} & \{ \{ \{ x^-y^-M, N^+ \}, P^+ \} \} \quad (\text{membrane, solution}) \\ \rightarrow & \{ \{ \{ y^-M[N/x], P^+ \} \} \} \quad (\text{reaction}) \\ \rightarrow & \{ \{ y^-M[N/x], P^+ \} \} \quad (\text{hatching}) \\ \rightarrow & \{ M[N/x] | P/y \} \quad (\text{reaction}) \end{aligned}$$

Since the λ -calculus is embedded in our γ -calculus; one can define arbitrary combinators such as a “replicator”, D , that satisfies $(DM) \xrightarrow{*} M | (DM)$ for all M , or a “killer”, U , that satisfies $(UM) \xrightarrow{*} U$. This is easy using standard fixpoint combinators. Moreover, our concurrent γ -calculus is more powerful than the

λ -calculus. The most important non λ -definable object that can now be constructed is the *internal choice* (or more accurately *join*) operator. To see this, let us denote by K and F the two cancellators, i.e., respectively $\lambda x.\lambda y.x$ and $\lambda x.\lambda y.y$ (in our syntax x^-y^-x and x^-y^-y). Then the choice operator is defined by

$$\oplus =_{\text{def}} \langle K | K^+ | F^+ \rangle$$

This operator may be evaluated either into K , like $(KK)F$, or into F , like $(KF)K$, therefore one easily sees that $\oplus MN \xrightarrow{*} M$ and $\oplus MN \xrightarrow{*} N$. Clearly such a combinator is not λ -definable since it does not preserve the Church-Rosser property.

As in [7], we extend our syntax by defining *concurrent abstractions*, that is sets of negative valences. More precisely, we define receptors of the form $[x_1 \mid \dots \mid x_n]^-M$ where $x_1, \dots, x_n$ are distinct variables. Such a term is able to receive n unordered values, to be substituted for these variables in M . Obviously these terms can be incorporated in our calculus with an additional rule:

choice:

$$[x_1 \mid \dots \mid x_n]^-M \rightarrow x_i^- [\dots \mid x_{i-1} \mid x_{i+1} \mid \dots]^-M$$

Concurrent abstractions do not add power to the original calculus, since we can also define an atomic receptor $[x_1 \mid \dots \mid x_n]^-M$ as a choice among all possible permutations $x_{i_1}^- \dots x_{i_n}^-M$. For instance, using an infix notation for the choice:

$$[x \mid y]^-M =_{\text{def}} x^-y^-M \oplus y^-x^-M$$

The concurrent abstraction feature allows us to define combinators in a very compact way. For instance, the choice operator can be redefined by $\oplus = [x \mid y]^-x$, that is a parallel variant of the usual cancellator K .

We can also define a “parallel or”, which is a parallel variant of the usual “left-sequential or” (cf. [7]). Let us see this point in some detail. It is known (see [4]) that $K = x^-y^-x$ and $F = x^-y^-y$ can be regarded as the truth values, respectively true and false. Then one may define a combinator for disjunction, namely $V = x^-y^-(xK)y$. This combinator is such that VKX reduces to K and VFX reduces to X . However, VXK (that is “ X or true”) cannot be in general reduced to K without evaluating X . For instance if, as usual, Ω denotes the non-terminating term $\Delta\Delta$ (where $\Delta = x^-(xx)$ is the duplicator) then the evaluation of $V\Omega K$ does not terminate. This is why V is “left-sequential”. Moreover from Berry’s sequentiality theorem (see [4]), one can show that there is no λ -definable combinator representing *parallel disjunction*, that is a combinator O such that both OKX

and OXK reduce to K , without evaluating X (and obviously OFF reduces to F). This combinator exists in the γ -calculus and is represented by:

$$O = [x | y]^- (xK)y$$

that is a parallel variant of the left-sequential disjunction (or equivalently a choice between left-sequential disjunction V and right-sequential disjunction $y^- x^- (xK)y$, see [10]).

5.3 Semantics

It seems fair to say that we have not yet established that “parallel disjunction is γ -definable”. This is a semantic statement, so we would have first to define an equivalence relation $\simeq$ on γ -terms such that (using an infix notation for the “parallel disjunction” combinator O):

$$K O \Omega \simeq K \simeq \Omega O K$$

and

$$\Omega O \Omega \simeq \Omega$$

In [7] it was proposed to adapt the notion of observational bisimulation $\approx$ of CCS [18] (see also [21]), to serve as the semantic equivalence. We could define this notion here (with the idea that x^- is an input guard and M^+ an output action), but this does not seem to be a good choice. For instance we would have $(K O \Omega) \not\approx K$ since $(K O \Omega)$ may be reduced to $\Omega K K$, a term without any communication capability, which is certainly different from K .

As a matter of fact, observational bisimulation has often been criticized for being too discriminating, and weaker “extensional” equivalences have been proposed (for a survey, see [12] and [15]). For instance Darondeau in [11] argued that “a semantics which stems from more sophisticated observers [than programs] is not really extensional”. In other words, the semantics of processes should be derived from their observation by means of *program contexts* $C[\cdot]$. These program contexts may be regarded as *tests* over processes, and there is a natural way to define an associated *testing equivalence* (cf. [14]): two process are equivalent if they pass the same tests. This is the kind of semantical equality we propose for our γ -calculus. However, we shall not follow [11] and [14] for what concerns the result of experiments. To report the success of a test we shall use, as in [1], the simplest operational information, namely *convergence*, that is existence of a normal form: the agent M pass the test $C[\cdot]$ if $C[M]$ converges.

Formally, an agent M is said to converge, in notation $M\Downarrow$, if and only if there exists an *inert* solution

S such that

$$\langle\langle M \rangle\rangle \xrightarrow{*} S$$

Then the definition of the testing preorder (on closed terms) is exactly the one of Morris’ preorder (cf. [4], exercise 16.5.5, and [1]), that is:

$$M \sqsubseteq N \Leftrightarrow_{\text{def}} \forall C. C[M]\Downarrow \Rightarrow C[N]\Downarrow$$

As usual the associated equivalence $\simeq$ is given by

$$M \simeq N \Leftrightarrow_{\text{def}} M \sqsubseteq N \ \& \ N \sqsubseteq M$$

Let us see an example, showing that testing allows to distinguish divergent terms in the γ -calculus (unlike the lazy λ -calculus). We still use MN to abbreviate application, that is $\langle M | N^+ \rangle$. As we saw, the typically divergent λ -term is $\Omega = \Delta\Delta$ where Δ is the duplicator $x^- (xx)$. It might be observed that $P = (\Delta | \Delta^+)$ is also a divergent term, since it can only be evaluated into Ω . Similarly, we can define a “triplicator” $T = x^- ((xx)x)$, and it is easy to see that $Q = (T | T^+)$ is again a divergent term. Now there is a test separating P and Q , namely

$$C = \langle\langle [\cdot] | z^- (F)^+ \rangle\rangle | \Omega^+$$

(recall that $F = x^- y^- y$, hence $FM \xrightarrow{*} I = y^- y$). It is not difficult to see that $C[P]$ diverges, whereas $C[Q] \xrightarrow{*} I$, therefore $P \not\sqsubseteq Q$.

We shall not investigate here the properties of the testing preorder. A first step would be to prove a generalization of the well-known “context lemma” (cf. [10]), showing that observers of the form

$$\langle\langle \dots \langle [\cdot] | R_1 \rangle \dots | R_k \rangle\rangle$$

are enough to test the agents, that is

$$M \sqsubseteq N \Leftrightarrow \forall k \forall R_1, \dots, R_k.$$

$$\langle\langle \dots \langle M | R_1 \rangle \dots | R_k \rangle\rangle\Downarrow \Rightarrow \langle\langle \dots \langle N | R_1 \rangle \dots | R_k \rangle\rangle\Downarrow$$

Such a result would allow us to give a simple proof of the desired properties of the parallel-or combinator.

6 Conclusion

Unlike other models, the Γ [2] and cham models handle (true) concurrency as *the* primitive built-in notion. What the cham model adds to Γ is the structure of molecules as terms and the notion of a subsolution. The CCS and TCCS implementations give a simple operational vision of these calculi. Inference rules are replaced by standard rewrite rules. The difference between internal and external transitions is made obvious and so are the well-known difficulties with sums

considered as programming primitives. More powerful "universal" process calculi such as MEIJE [6] can be handled as well. The concurrent λ -calculus fully exploits the ability of going back and forth between terms and solutions. It can be viewed as a direct extension of the lazy λ -calculus of [1].

Of course, this is still a preliminary work. Other concurrent computation paradigms should also be modelled; we think in particular of modelling process handling in operating systems, and of providing a cham for the new calculus of mobile processes proposed by Milner & al. in [20]. The theory of machine executions should also be fully developed.

Acknowledgments

We are indebted to Ilaria Castellani, Philippe Darondeau, Matthew Hennessy, Robin Milner and Serge Yoccoz for helpful discussions about this work and previous versions of the paper.

References

- [1] Samson Abramsky. The lazy λ -calculus. In D. Turner, editor, *Declarative Programming*, Addison Wesley, 1988.
- [2] Jean-Pierre Banâtre, Anne Coutant, and Daniel Le Metayer. A parallel machine for multi-set transformation and its programming style. In *Future Generation Computer Systems 4*, pages 133–144, North-Holland, 1988.
- [3] Jean-Pierre Banâtre and Daniel Le Metayer. *A New Computational Model and Its Discipline of Programming*. Technical Report INRIA Report 566, 1986.
- [4] Henk Barendregt. *The Type-Free Lambda-Calculus. Studies in Logic Volume 103*, North-Holland, 1981.
- [5] Gérard Berry. Séquentialité de l'évaluation formelle des λ -expressions. In B. Robinet, editor, *Program Transformations 3rd International Colloquium on Programming*, pages 67–80, DUNOD, Paris, 1978.
- [6] Gérard Boudol. Notes on algebraic calculi of processes. In K. Apt, editor, *Logic and Models of Concurrent Systems*, NATO ASI Series F13, 1985.
- [7] Gérard Boudol. Towards a lambda-calculus for concurrent and communicating systems. In *TAPSOFT 1989, Lecture Notes in Computer Science 351*, pages 149–161, Springer-Verlag, 1989.
- [8] Nicholas Carriero and David Gelerntner. Linda in context. *Communications of the ACM*, 32(4):444–458, 1989.
- [9] Mani Chandy and Jayadev Misra. *Parallel Program Design, a Foundation*. Addison-Wesley, 1988.
- [10] Pierre-Louis Curien. *Categorical Combinators, Sequential Algorithms, and Functional Programming. Research Notes in Theoretical Computer Science*, Pitman, London, John Wiley & Sons, New York, Toronto, 1986.
- [11] Philippe Darondeau. About fair asynchrony. *Theoretical Computer Science*, 37:305–336, 1985.
- [12] Rocco De Nicola. Extensional equivalences for transition systems. *Acta Informatica*, 24:211–237, 1987.

- [13] Rocco De Nicola and Matthew Hennessy. CCS without τ 's. In *TAPSOFT 87, Lecture Notes in Computer Science 249*, pages 138–152, Springer-Verlag, 1987.
- [14] Rocco De Nicola and Matthew Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1984.
- [15] Matthew Hennessy. Observing processes. In *Linear Time, Branching Time and Partial Orders in Logics and Models for Concurrency, Lecture Notes in Computer Science 354*, pages 173–200, Springer-Verlag, 1989.
- [16] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
- [17] Peter Landin. The mechanical evaluation of expressions. *Computer Journal*, 6:308–320, 1964.
- [18] Robin Milner. *Communication and Concurrency. International Series in Computer Science*, Prentice Hall, 1989.
- [19] Robin Milner. *Program Semantics and Mechanized Proofs*, pages 3–44. Mathematical Center Tracts 82, Amsterdam, 1976.
- [20] Robin Milner, Joachim Parrow, and David Walker. *A Calculus of Mobile Processes*. Technical Report ECS-LFCS-89-85, LFCS, Edinburgh University, 1989.
- [21] Bengt Thomsen. A calculus of higher-order communicating systems. In *Proc. 16th ACM Annual Symposium on Principles of Programming Languages*, pages 143–154, 1989.

Imprimé en France
par
l'Institut National de Recherche en Informatique et en Automatique

13

14

15

16

17

18