

CHOCOLATE: Calculi of Higher Order COmmunication and LAMBDA TErms (Preliminary Report)

Bard Bloom*
Cornell University, Ithaca, NY
bard@cs.cornell.edu

October 29, 1993

Abstract

We propose a general definition of higher-order process calculi, generalizing CHOCS [Tho89] and related calculi, and investigate its basic properties. We give sufficient conditions under which a calculus is finitely-branching and effective. We show that a suitable notion of higher-order bisimulation is a congruence for a subclass of higher-order calculi. We illustrate our definitions with a sample calculus strictly stronger than CHOCS.

1 Introduction

Higher-order process calculi provide have captured considerable interest recently. These are calculi for concurrency in which the basic elements of the system, processes or channels, are presented as first-class objects. As in sequential programming, higher-order features grant programmers a wide variety of programming methods.

For example, in first-order calculi (CCS, CSP, ACP, and so forth), servers must use fixed communication channels. Requests come to the server on channel a , replies return on channel b , and woe betide the system if anyone else uses these channels. Clients of the server will have to have some kind of mutual exclusion or other access control protocol, requiring some careful programming and introducing more ways to make errors: a bad thing in a system intended for specifications.

*Supported by NSF grants CCR-9003441
and CCR-9223183.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

POPL 94- 1/94, Portland Oregon, USA

© 1994 ACM 0-89791-636-0/94/001..\$3.50

Higher-order calculi such as CHOCS [Tho89] and the related calculus of [Hen93], and the π -calculus [MPW92] address these issues. In these calculi, processes and channels may be communicated along channels. Thus, in CHOCS, a client of a server could send the server a small process capable of sending messages back to the client, thereby reducing and encapsulating the possible misuse of the server. In the π -calculus, the interface is even cleaner: the client sends the server a channel to use for the reply.

There is a distinction between two kinds of higher-order calculi. Some calculi, CHOCS and its relatives, allow transmission of *processes* along channels. This is extremely powerful, and the implications of this programming paradigm remain to be fully explored. The other school, the π -calculus and its relatives, simply allow transmission of *channels* (or channel names) along channels. Programming in this setting is somewhat harder (*e.g.*, the implementation of the λ -calculus in the π -calculus is rather more subtle than in CHOCS), and the classical operational semantics is considerably trickier. Implementation of π -calculus-like systems is much easier; *e.g.*, the communications fragment of Concurrent ML [Rep91, Rep89] resembles the π -calculus in flavor.

In this study, we make preliminary investigations of the general theory of CHOCS-like languages. (The π -calculus is rather more delicate, and will be the subject of later studies.) The operations provided with CHOCS and its relatives are the basic concurrent operations: *e.g.*, a CCS-like parallel composition $p|q$ which allows p and q to either execute independently or communicate.

One useful operation not definable in CHOCS is *broadcast*, ∇ . This associative, commutative binary operation has the property that, if one combines several processes:

$$p \nabla q \nabla r \nabla s$$

if one process, say p , sends a message m on channel

a , then all of q , r , and s that are currently capable of reading from channel a receive m . Unlike anything programmable with simple parallel composition, this guarantees that *all* possible interested receivers get the message, simultaneously. Broadcast in various forms is an essential programming operation in many distributed systems [BCG91, BC90].

Another operation which cannot even be simulated easily in CHOCS is the LOTOS operation of *disabling*, $[>$. The process $p [> q$ behaves like p , until such time as q takes its first step; thereafter, p is killed and q executes. This is vital in the higher-order setting; we may receive a process, start to execute it, discover that it is misbehaving, and wish to kill it before it can do any further damage. The process $a?.x.((x\{b\}) [> (b?.y.O))$ receives a process p along a , and runs it — but provides an “off” switch, as a signal along channel b will cancel it.

Any number of other operations have been proposed and used in first-order process calculi. There are *synchronous products* which run processes in lock-step parallel. There are *polling operations*, which allow one process to ask if another is ready to communicate without actually performing the communication. Some operations are simply for programming convenience: sequential composition $p; q$ and **while** loops are useful programming constructs which can be simulated to some extent in most process calculi, but might reasonably be made part of the language.

1.1 Developing a Metatheory

Rather than rebuilding the basic theory whenever a new operation is needed, we investigate the metatheory of CHOCS-like process calculi. In particular, we define a type system and form of rules, called the *chocolate* rules in honor of their derivation from CHOCS. CHOCS and Hennessy’s calculus are chocolate, *mutatis mutandis*, as are their extensions by all of the operations mentioned above and much more.

Chocolate languages have two kinds of computation. The primary form is process-algebraic: processes produce signals (carrying other processes as data) and evolve into other processes, much as in CCS and CHOCS. The other form of computation is reduction in a typed λ -calculus.¹ For example, a term that is prepared to receive a value is written as a function taking values as arguments. Giving it a value corresponds to function application. This application is evaluated, eventually yielding a process capable of engaging in further communication. The evaluation part of computation is **hidden inside the two or three relevant rules in CHOCS**; it is made explicit (though instantaneous) in

¹Typing is not essential to the definition of computation. It is essential to many finiteness and computability theorems.

chocolate.

Some basic properties hold for all chocolate languages: in particular, the evaluation part of computation is strongly normalizing (which is no great surprise, as the evaluation part of the calculus is little more than simply typed λ -calculus.) We use this fact to give sufficient conditions on chocolate languages guaranteeing some important properties; *e.g.*, explaining when the transition relation is computable and finitely branching. We illustrate with NESTLÉ, an extension of CHOCS including broadcast and channel creation and transmission.

We then investigate notions of bisimulation. The most natural and useful definition of bisimulation at all types is dangerous; technically, it is not a monotone functional and hence does not necessarily have a greatest fixed point. The theory may work out appropriately, but this remains to be shown.

Instead, we provide a generalization of Thomsen’s notion of higher-order bisimulation. Thomsen’s notion is *not* an instance of the natural notion, as we explain in Section 4. Nonetheless, in a restricted second-order case, Thomsen’s notion can be generalized to the concept of *white chocolate bisimulation*.² We show that white chocolate bisimulation is a congruence for a respectable class of operations, including those of CHOCS and Hennessy’s calculus. This suffices for most of the operations we used to motivate this study. Unfortunately, it fails for the most interesting example, the channel transmission part of the calculus NESTLÉ.

1.2 Types and Channels

A fully-fledged higher-order process calculus should allow the kinds of programming paradigms that make higher-order sequential computation powerful: encapsulating and transmitting programming constructs, such as transmission and reception operations. That is, we should be able to transmit *channels*, rather in the style of the π -calculus. Keeping track of types will prove informative.

Let \mathbf{P} be the type of processes. A receiver is a function that takes an input (of type \mathbf{P}) and then acts like a process. It is thus an *abstraction*, of type

$$\mathbf{Ab} = \mathbf{P} \rightarrow \mathbf{P}$$

A sender produces a datum (of type \mathbf{P}) as output, and then continues running its main thread of communication. That is, it simply is a pair of processes $\langle d, p \rangle$: d to be used as data, p for computation. It is thus a *concretion* [Mil91], of type

$$\mathbf{Co} = \mathbf{P} \times \mathbf{P}$$

²It’s good, but it’s not real chocolate.

Most higher-order process algebras have send and receive operations, written something like $a!v.p$ and $a?x.p$. We follow [Mil91] in considering $a!$ and $a?$ to be the operations. $a!$ takes a value and a process, and returns a process. $a?$ takes a function from values to processes, returning a process. Thus, we have

$$\begin{aligned} a! & : \mathbf{Snd} = \mathbf{Co} \rightarrow \mathbf{P} \\ a? & : \mathbf{Rcv} = \mathbf{Ab} \rightarrow \mathbf{P} \end{aligned}$$

Given $r : \mathbf{Rcv}$, we read a value from r and call it v by a construct of the form $r(\lambda v.\square)$. Given $s : \mathbf{Snd}$, we transmit n along it by $s(\langle n, \square \rangle)$. In both cases, the \square is the continuation: the process to be executed after the communication.

A communication channel consists of a way to send, and a way to receive: a pair $\langle a!, a? \rangle$, of matching sending and receiving capabilities. They thus have type:

$$\mathbf{Chan} = \mathbf{Snd} \times \mathbf{Rcv}$$

If, in our process algebra, we are able to communicate values of type \mathbf{Chan} , then we can do much π -calculus-like programming straightforwardly. For example, consider a server: a process which repeatedly accepts a \mathbf{Chan} , *viz.* a pair $\langle s, r \rangle : \mathbf{Snd} \times \mathbf{Rcv}$, reads a value v from the channel, performs some computation F , and sends its result $F(v)$ back along the channel. The server should be reentrant: once it has gotten a channel, it can fork off a process $r(\lambda v.s(F(v), \mathbf{0}))$, abbreviated as $rv.s(F(v)).\mathbf{0}$, to handle the request along that channel; the server itself would wait for more requests. Each request can use a different channel.

$$S = a?\langle s, r \rangle . (S|(rv.s(F(v)).\mathbf{0})) \quad (1)$$

Or, without abbreviations:

$$S = a?\left(\lambda h.S\left(\left(\pi_1 h\right)\left(\lambda v.\left(\pi_0 h\right)\langle F(v), \mathbf{0} \rangle\right)\right)\right)$$

A first draft of a client (sending first the channel b , as the value $\langle b!, b? \rangle$; then sending the value 3 coded in some unspecified way, and sending the answer along the channel *print*) would look like:

$$C_0 = a!\langle b!, b? \rangle . b!3 . b?x . \text{print!}x . \mathbf{0}$$

To make full use of this, we should have a process which generates new channels on request. Suppose that there are actions c_0, c_1, \dots . We'd like a process with roughly the following behavior, repeatedly offering new channels as values on some fixed channel c :

$$NC(n) = c!\langle c_n!, c_n? \rangle . NC(n+1)$$

(A somewhat more powerful version of the new-channel process will be provided as a primitive.) A better client

would first get a channel from NC , and use that channel to communicate with the server.

$$C_1 = c?\langle s, r \rangle . a!\langle s, r \rangle . s3 . rx . \text{print!}x . \mathbf{0} \quad (2)$$

The whole system will simply consist of the server, the client, and the channel server in parallel:

$$\text{System} = (S|C_1|NC(1)) \setminus \{a, b, c, c_1, c_2, \dots\}.$$

Note that the $c!$ and $c?$ used in NC and C_1 are higher-order communication operations, as they are being used to communicate values of type \mathbf{Chan} .

$$\begin{aligned} c? & : (\mathbf{Chan} \rightarrow \mathbf{P}) \rightarrow \mathbf{P} \\ & = \left(\left(\begin{array}{c} (\mathbf{P} \times \mathbf{P} \rightarrow \mathbf{P}) \\ \times \\ ((\mathbf{P} \rightarrow \mathbf{P}) \rightarrow \mathbf{P}) \end{array} \right) \rightarrow \mathbf{P} \right) \rightarrow \mathbf{P} \end{aligned}$$

This development of channel transmission was quite natural and straightforward. It is somewhat surprising that it requires constants of such high types. We are now beyond the scope of existing process algebras, even higher-order ones, and must check to see that the foundations are solid.

1.2.1 Typed Transition Relations

Hennessy's calculus has two actions $a?$ and $a!$ for each channel, and rules

$$a?\lambda x.p \xrightarrow{a?} \lambda x.p \qquad a!\langle q, p \rangle \xrightarrow{a!} \langle q, p \rangle$$

We have transitions $p \xrightarrow{a, \xi} q$, where $p : \mathbf{P}$, a is an action, ξ is a type, and q has type ξ . (The definitions are sketched in Section 2.4.) The type information renders the $?$ and $!$ redundant. In our system, those transitions are

$$a \mathbf{Ab} \lambda x.p \xrightarrow{a, \mathbf{Ab}} \lambda x.p \qquad a \mathbf{Co} \langle q, p \rangle \xrightarrow{a, \mathbf{Co}} \langle q, p \rangle$$

Type information distinguishes between sending and receiving, so we use the same operation and action symbols.

This notion of transition is sensible in practice. It has the somewhat disturbing property that it doesn't respect types in any sense: a process – that is, a closed term of type \mathbf{P} – can take a -transitions to terms of two different types. This presents no difficulties, though it may seem odd.

2 Calculi

We generalize the notion of a *GSOS language* [BIM88, Blo89]. GSOS languages are a first-order generalization of the rules used to define CCS.³ We add a type system and higher-typed terms.

³The acronym GSOS currently stands for "Grand Structured Operational Semantics".

2.1 Types and Terms

The type system is straight from the simply typed λ -calculus with products and one base type, which we call $\lambda^{\rightarrow, \times}$. Types ξ have the form:

$$\xi ::= \mathbf{P} \mid \xi \rightarrow \xi \mid \xi \times \xi$$

Definition 2.1 A Higher Order GSOS Signature is a set Act of actions, a set of operation symbols, and a vector of k_f types $\langle \xi_{f1}, \dots, \xi_{fk_f} \rangle$ for each operation symbol f , signifying that f takes arguments of types $\xi_{f1}, \dots, \xi_{fk_f}$ and returns a value of type \mathbf{P} .

We assume infinitely many variables x^ξ of each type. We define terms and their types simultaneously:

$$\frac{}{x^\xi : \xi} \quad \frac{t_i : \xi_{fi}, 1 \leq i \leq k_f}{f(t) : \mathbf{P}}$$

$$\frac{t : \xi_0 \times \xi_1}{\pi_i(t) : \xi_i} \quad \frac{t_i : \xi_i, i \in \{0, 1\}}{\langle t_0, t_1 \rangle : \xi_0 \times \xi_1}$$

$$\frac{t : \xi_0 \rightarrow \xi_1, t' : \xi_0}{tt' : \xi_1} \quad \frac{t : \xi_1}{\lambda x^{\xi_0}. t : \xi_0 \rightarrow \xi_1}$$

We use a number of fairly obvious abbreviations to make processes easier to read. As we are combining process algebra and $\lambda^{\rightarrow, \times}$, we have in effect two computation systems, plus a rule connecting them.

2.2 The languages NESTLÉ and NESTLÉ/2

Based on the considerations above, we define two higher-order process calculi, NESTLÉ and its second-order subset NESTLÉ/2. NESTLÉ includes all the operations of CHOCS, plus full higher-order communication, broadcast, disabling, and channel creation; it illustrates all the features of the rules given in our study. Our action alphabet Act is infinite; a, b, c range over it. It includes an infinite, coinfinite subset $\{c_0, c_1, \dots\}$, and a distinguished symbol τ . The operation symbols are:

$\mathbf{0} : \langle \rangle$	Null process
$a^\xi : \langle \xi \rangle$	Prefixing at all types, for each action a
$+$: $\langle \mathbf{P}, \mathbf{P} \rangle$	Nondeterministic choice
$ $: $\langle \mathbf{P}, \mathbf{P} \rangle$	Parallel Composition
$[>$: $\langle \mathbf{P}, \mathbf{P} \rangle$	Disabling
$\backslash S$: $\langle \mathbf{P} \rangle$	Restriction: forbid actions in $S \subseteq \text{Act}$
$[R]$: $\langle \mathbf{P} \rangle$	Renaming a 's to $R(a)$'s,
:	where $R : \text{Act} \rightarrow \text{Act}$
∇_C : $\langle \mathbf{P}, \mathbf{P} \rangle$	Broadcasting on channels in $C \subseteq \text{Act}$
$[\cdot]S$: $\langle \mathbf{P} \rangle$	Delimit broadcasting
NC_n : $\langle \rangle$	New channel creator:
:	absent from NESTLÉ/2

We write operations in infix in fairly obvious ways. We use the abbreviations $a!p.q = a\langle p, q \rangle$ and $a?x.p =$

$a(\lambda x.p)$. We have one rule for each choice of a, ξ, C, S , and R . The rules for prefixing and choice are quite standard:

$$\frac{}{a^\xi x \xrightarrow{a, \xi} x} \quad \frac{x \xrightarrow{a, \xi} y}{x + x' \xrightarrow{a, \xi} y} \quad \frac{x \xrightarrow{a, \xi} y}{x' + x \xrightarrow{a, \xi} y}$$

For parallel composition, we define $y|_\xi x'$ (where $y : \xi$ and $x' : \mathbf{P}$) to be an abbreviation for the following term, which has the same communication effect as y does, but has x' still running in parallel with it. That is, if y wants to receive a value (i.e., has type $\xi_0 \rightarrow \xi_1$), then $y|_{\xi_0 \rightarrow \xi_1} x'$ will receive a value and route it to y . Similarly, if y wants to transmit a value (viz, has type $\xi_0 \times \xi_1$), $y|_{\xi_0 \times \xi_1} x'$ will transmit the same value $\pi_0 y$, and continue running the rest of y in parallel with x' . This is the natural generalization of [Hen93].

$$y|_{\mathbf{P}} x' = y|x'$$

$$y|_{\xi_0 \times \xi_1} x' = \langle (\pi_0 y), (\pi_1 y)|_{\xi_1} x' \rangle$$

$$y|_{\xi_0 \rightarrow \xi_1} x' = \lambda z^{\xi_0}. (y.z|_{\xi_1} x')$$

Note that $y|_\xi x : \xi$. We omit the symmetric rules, and use informal pattern matching notation for clarity.

$$\frac{x \xrightarrow{a, \xi} y}{x|x' \xrightarrow{a, \xi} y|_\xi x'} \quad \frac{x \xrightarrow{a, \xi \rightarrow \mathbf{P}} y, \quad x' \xrightarrow{a, \xi \times \mathbf{P}} \langle y'_0, y'_1 \rangle}{x|x' \xrightarrow{\tau, \mathbf{P}} (y y'_0)|y'_1}$$

For example,

$$p_0 = (a?x.p)|q \xrightarrow{a, \mathbf{P} \rightarrow \mathbf{P}} \lambda z.(pz|q)$$

$$p_1 = (a!s.t)|u \xrightarrow{a, \mathbf{P} \times \mathbf{P}} \langle s, t|u \rangle$$

then s is sent to p as desired in $p_0|p_1$:

$$p_0|p_1 \xrightarrow{\tau, \mathbf{P}} (ps|q)|(t|u)$$

Disabling $p [> q$ runs p until q takes a step, then kills p and runs q . The type extension $[>_\xi$ is defined in the same way as $|_\xi$.

$$\frac{x \xrightarrow{a, \xi} y}{x [> x' \xrightarrow{a, \xi} y [>_\xi x'} \quad \frac{x' \xrightarrow{x, \xi} y'}{x [> x' \xrightarrow{a, \xi} y'}$$

Restriction and renaming must be extended to higher types in much the same way that parallel composition is. Let $t : \xi$; we define $f_\xi(t)$ as follows:

$$f_{\mathbf{P}}(t) = t$$

$$f_{\xi_0 \rightarrow \xi_1}(t) = \lambda z^{\xi_0}. f_{\xi_1}(tz)$$

$$f_{\xi_0 \times \xi_1}(t) = \langle \pi_0 t, f_{\xi_1}(\pi_1 t) \rangle$$

The restriction and renaming rules are fairly standard:

$$\frac{x \xrightarrow{a, \xi} y \text{ for } a \notin S}{x \setminus S \xrightarrow{a, \xi} y \setminus S} \quad \frac{x \xrightarrow{a, \xi} y}{x[R] \xrightarrow{R(a), \xi} y[R]_\xi}$$

Broadcasting $x \Downarrow_C x'$ is similar to parallel composition. It has the parallel composition rules for $a \notin C$, plus the following rules and their symmetric versions. If x broadcasts y_0 , then x' must receive it if it can; either way, $x \Downarrow_C x'$ will broadcast y_0 .

$$\frac{x \xrightarrow{a, \xi \times \mathbf{P}} \langle y_0, y_1 \rangle, \quad x' \xrightarrow{a, \xi \rightarrow \mathbf{P}} y'}{x \Downarrow_C x' \xrightarrow{a, \xi \times \mathbf{P}} \langle y_0, (y_1 \Downarrow_C y' y_0) \rangle}$$

$$\frac{x \xrightarrow{a, \xi \times \mathbf{P}} \langle y_0, y_1 \rangle, \quad x' \xrightarrow{a, \xi \rightarrow \mathbf{P}}}{x \Downarrow_C x' \xrightarrow{a, \xi \times \mathbf{P}} \langle y_0, y_1 \Downarrow_C x' \rangle}$$

Note that a broadcast always emits a value y_0 . If $p \Downarrow_C q \xrightarrow{a, \xi} r$ for some $a \in C$, then r is always a pair $\langle r_0, r_1 \rangle$ where r_0 was the value broadcast and r_1 the system after the broadcast occurred. Recall that if $t \xrightarrow{a, \xi} u$, then t must be a process. We thus need a broadcast-delimiting operation $[x]C$ which strips off the value being broadcast and lets the remaining process continue.

$$\frac{x \xrightarrow{a, \xi \times \mathbf{P}} y, \quad a \in C}{[x]C \xrightarrow{\tau, \mathbf{P}} \pi_1 y} \quad \frac{x \xrightarrow{a, \xi} y, \quad a \notin C}{[x]C \xrightarrow{a, \xi} [x]C, \xi}$$

$$\frac{x \xrightarrow{a, \mathbf{P}} y}{[x]C \xrightarrow{a, \mathbf{P}} [y]C}$$

So, the idiom for broadcast is

$$[p \Downarrow_C q \Downarrow_C \cdots \Downarrow_C r \Downarrow_C s]C$$

much as the idiom for parallel composition with point-to-point communication is

$$(p|q|\cdots|r|s)\setminus S.$$

The new channel creator NC_n simply offers new communication channels (on channel c_0). Let $\mathbf{Chan}(\xi) = (\xi \rightarrow \mathbf{P}) \times (\xi \times \mathbf{P})$.

$$\text{NC}_n \xrightarrow{c_0, \mathbf{Chan}(\xi) \times \mathbf{P}} \langle \langle \lambda x. c_n. \xi \rightarrow \mathbf{P}(x), \lambda x. c_n. \xi \times \mathbf{P}(x) \rangle, \text{NC}_{n+1} \rangle$$

Recursive definition of processes is straightforward as a consequence of our general theory: we allow constants

defined by an arbitrary set of axioms. For example, the server of (1) would be defined as a new constant S , with a single rule

$$S \xrightarrow{a, \mathbf{Chan} \rightarrow \mathbf{P}} \lambda \langle s, \tau \rangle. (S[\tau(\lambda v. s(F(v), \mathbf{0}))]).$$

NESTLÉ is a remarkably powerful calculus, allowing CHOCS-like and π -calculus-like programming. With a bit of care (having a single NC operator in parallel, and avoiding use of c_n for $n > 0$), programming with dynamically-allocated channels is straightforward.

The only π -calculus operation missing from NESTLÉ is matching: $[a = b]q$ behaves like q if the channel names a and b are the same, and is stopped otherwise. This is programmable (with some extra τ -moves) in NESTLÉ.

Suppose that we have two channel values of the same type ξ , $\chi_1 = \langle s_1, r_1 \rangle$ and $\chi_2 = \langle s_2, r_2 \rangle$, produced by NC. The two are equal if a signal sent on s_1 is received on r_2 . We use broadcasting to force the communication.

Let $\mathbf{0}^\xi$ be any closed term of type ξ . Choose channels a, b_y , and b_n . The following code will send a signal on b_y if $\chi_1 = \chi_2$, and on b_n otherwise. Let $G = \{a, c_0, c_1, \dots\}$;

$$p_1 = (r_2?x. b_y! \mathbf{0} . \mathbf{0}) + (a?x. b_n! \mathbf{0} . \mathbf{0})$$

$$p = \left[(s_1! \mathbf{0}^\xi . a! \mathbf{0} . \mathbf{0}) \Downarrow_G p_1 \right] G \setminus G$$

If $\chi_1 = \chi_2$, this will evolve as follows:

$$p \xrightarrow{\tau, \mathbf{P}} [a! \mathbf{0} . \mathbf{0} \Downarrow_G b_y! \mathbf{0} . \mathbf{0}] G \setminus G \xrightarrow{b_y, \mathbf{C}\mathbf{0}} \langle \mathbf{0}, \square \rangle$$

where \square is $\mathbf{0}$ inside some communication delimiters. (There might be another τ -move at one end or the other.) The use of \Downarrow_G forces the communication between s_1 and r_2 to occur if it is possible.

On the other hand, if $\chi_1 \neq \chi_2$, then the left-hand process broadcasts into the void; we have only a single computation:

$$p \xrightarrow{\tau, \mathbf{P}} [a! \mathbf{0} . \mathbf{0} \Downarrow_G p_1] G \setminus G$$

$$\xrightarrow{\tau, \mathbf{P}} [\mathbf{0} \Downarrow_G b_n! \mathbf{0} . \mathbf{0}] G \setminus G$$

$$\xrightarrow{b_n, \mathbf{C}\mathbf{0}} \langle \mathbf{0}, \square \rangle$$

Thus, p can tell if two channels are the same or not.

NESTLÉ/2 is NESTLÉ restricted to second-order types: no channel creation, and only tuples of processes can be transmitted. It is strictly between CHOCS and NESTLÉ in expressive power, and our stronger theory applies to it.

2.3 Operational Semantics

A chocolate language consists of a signature and some structured operational rules defining the behavior of the

operators of the signature. The λ -calculus fragment is standard.

Definition 2.2 *Higher Order GSOS rules have the form*

$$\frac{x_1 \xrightarrow{a_{11}, \xi_{11}} y_{11}, \dots, x_1 \xrightarrow{b_{11}, \xi'_{11}}, \dots}{f(\vec{x}) \xrightarrow{c} u} \quad (3)$$

There may be infinitely many antecedents, both positive and negative; u may be any term. The variables x_i and y_i , must all be distinct, and the only variables appearing in the term u are x_i 's and y_i 's.

Definition 2.3 *A chocolate language \mathcal{L} over a signature is a set of higher-order GSOS rules over that signature.*

The evaluation rules for $\lambda^{\rightarrow, \times}$ are standard, and not left up to the language designer.⁴

$$\frac{}{(\lambda x.t)u \Longrightarrow t[x := u]} \quad \frac{}{\pi_i(\langle t_0, t_1 \rangle) \Longrightarrow t_i}$$

$$\frac{t \Longrightarrow t'}{tu \Longrightarrow t'u} \quad \frac{t \Longrightarrow t'}{\pi_i(t) \Longrightarrow \pi_i(t')}$$

The transition relation \Longrightarrow is deterministic. By standard theories of typed λ -calculus, it is strongly normalizing; that is, there are no infinite reductions sequences $p_0 \Longrightarrow p_1 \Longrightarrow \dots$.

Chocolate calculi all use the *Instantaneous Evaluation Rule*: local computation is instantaneous, and only communication takes time. This is actually moderately realistic, in software systems at least: local processing time are generally insignificant compared to communication delays.

$$\frac{x \Longrightarrow x', \quad x' \xrightarrow{a, \xi} y}{x \xrightarrow{a, \xi} y} \quad (4)$$

In the full paper we explain why having evaluation steps take τ -moves is undesirable.

For example, let p be a process which receives a channel along a , sends some d along that channel, and stops; and q be a process which sends a channel b along a , reads a value from b , and thereafter behaves like that value.

$$p = a?(\lambda c.\pi_0(c)\langle d, \mathbf{0} \rangle)$$

$$q = a!\langle\langle b!, b? \rangle, b?(\lambda x.x)\rangle$$

⁴It would probably make good sense to use a more powerful calculus, though the finiteness results are unlikely to hold if the base calculus is not strongly normalizing.

We have transitions:

$$p \xrightarrow{\tau, \mathbf{Chan} \rightarrow \mathbf{P}} \lambda c.\pi_0(c)\langle d, \mathbf{0} \rangle$$

$$q \xrightarrow{a, \mathbf{Chan} \times \mathbf{P}} \langle\langle b!, b? \rangle, b?(\lambda x.x)\rangle$$

and hence

$$p|q \xrightarrow{a, \mathbf{P}} ((\lambda c.\pi_0(c)\langle d, \mathbf{0} \rangle)\langle b!, b? \rangle)|b?(\lambda x.x) = p'|q'$$

Only the evaluation rule applies to p' . We have

$$((\lambda c.\pi_0(c)\langle d, \mathbf{0} \rangle)\langle b!, b? \rangle) \Longrightarrow \pi_0(\langle\langle b!, b? \rangle\rangle\langle d, \mathbf{0} \rangle)$$

$$\Longrightarrow b!\langle d, \mathbf{0} \rangle \xrightarrow{b, \mathbf{Co}} \langle d, \mathbf{0} \rangle$$

which matches q' 's transition $q' \xrightarrow{b, \mathbf{Ab}} \lambda x.x$ to give us a transition

$$p'|q' \xrightarrow{b, \mathbf{P}} \mathbf{0}|(\lambda x.x)d$$

That is, p has sent d to q , which is what we intended.

2.4 Definition of the Transition Relation

A transition $p \xrightarrow{a, \xi} q$ occurs in \mathcal{L} iff there is a proof of this formula from the rules of \mathcal{L} , plus the evaluation rules and the connection rule. The formal definition in the full paper is nontrivial, especially as we must define proofs for $p \xrightarrow{a, \xi}$ as well. The following lemma shows that the definitions fit together right:

Lemma 2.4 *Let \mathcal{L} be any chocolate language, and $p : \mathbf{P}$ any closed term. Then there is a \mathcal{L} -proof of $p \xrightarrow{a, \xi} q$ for some q iff there is no \mathcal{L} -proof of $p \xrightarrow{a, \xi}$.*

3 Computability and Finite Branching

There are some fundamental concerns about this calculus. Clearly, as we have allowed arbitrary sets of operation symbols, actions, and rules, we cannot expect anything to be computable, or even countable. It is often desirable to keep the transition relation computable and finitely branching.

3.1 Finite Branching

It is frequently important to have some kind of finiteness and computability properties. We discuss two important finite-branching properties.

Suppose that Π is a proof about p . The immediate subproofs of Π are either a proof about the arguments of p (if $p = f(\vec{p})$), or about a term p' with $p \Longrightarrow p'$. This suggests the following definition:

Definition 3.1 $p \triangleright\!\!\!\circ q$ iff either $p \implies q$, or $p = f(p_1, \dots, p_n)$ and $q = p_i$ for some i such that $p_i : \mathbf{P}$.

$\triangleright\!\!\!\circ$ is a lousy notion of reduction: it is nondeterministic, non-confluent (though finitely branching), and not respecting of meaning in any sense. The main excuse for $\triangleright\!\!\!\circ$ is the following:

Lemma 3.2 Let Π be a proof about p . If Π contains a subproof about a term p' , then $p \triangleright\!\!\!\circ^* p'$.

The usual proofs for strong normalization of the simply typed λ -calculus can be adapted to show that $\triangleright\!\!\!\circ$ is strongly normalizing and finitely branching. Thus, the following function is well-defined:

$$\text{dur}(p) = \max \left\{ n \mid \exists p'. p \triangleright\!\!\!\circ^n p' \right\}$$

The heart of the finiteness and computability proofs are the following lemma, where $\text{depth}(\Pi)$ is the depth of a proof in the obvious sense.

Lemma 3.3 Let Π be a proof about p . Then $\text{depth}(\Pi) \leq \text{dur}(p)$.

Definition 3.4 \mathcal{L} is image finite if, for all processes p , actions a , and types ξ , the set $\{p' \mid p \xrightarrow{a, \xi} p'\}$ is finite. \mathcal{L} is finitely branching if, for all processes p , the set $\{a, \xi, p' \mid p \xrightarrow{a, \xi} p'\}$ is finite.

Image finiteness is the weakest finite-branching property in general use. In many calculi, it is the strongest property that we want. The NESTLÉ operation NC_n is image-finite, but (as it offers channels of all types) is not finitely branching.

Theorem 3.5 Suppose that there are only a finite number of rules for $f(\vec{x}) \xrightarrow{a, \xi}$ for each f , a , and ξ , and that all rules have a finite number of positive antecedents. Then \mathcal{L} is image finite.

The essence of the following method to get finite branching is taken from [Ace92, Vaa93].

Definition 3.6 The trigger of a rule ρ in the form (3) is the vector $\langle S_1, \dots, S_n \rangle$, where

$$S_i = \left\{ \langle a_{ij}, \xi_{ij} \rangle \mid (x_i \xrightarrow{a_{ij}, \xi_{ij}} y_{ij}) \in \text{ante}(\rho) \right\}$$

Definition 3.7 A chocolate language is finitely triggered if, for each operation symbol f and potential trigger T , there are only finitely many rules for f with trigger T ; and each rule has only finitely many positive antecedents.

NESTLÉ is not finitely triggered, as there are infinitely many rules for NC_n with the empty tuple for a trigger. Without NC_n — or even with NC_n restricted to a finite set of types — it would be finitely triggered.

Theorem 3.8 If \mathcal{L} is finitely triggered, then it is finitely branching.

3.2 Computability

There are several possible definitions of computability. For this study, we choose one of the strongest: that the computation tree of a process is computable.

Definition 3.9 Let $D(p) = \left\{ \langle a, q \rangle \mid p \xrightarrow{a, \mathbf{P}} q \right\}$. \mathcal{L} is strongly computable iff $D(p)$ is computable in canonical index as a function of p .

By insisting that appropriate parts of Lemma 3.8 be computable, we can give sufficient conditions for strong computability.

Lemma 3.10 Let \mathcal{L} be a chocolate language, such that:

1. The sets of actions, operations, and rules have recursive codings, such that equality is recursive.
2. All rules have finite numbers of positive antecedents, and recursive sets of negative antecedents; and that the set of antecedents in canonical index and a program deciding the negative antecedents are computable from the code of rule.
3. For any operation symbol and finite trigger, the set of rules for that symbol and trigger is computable in canonical index.

Then \mathcal{L} is strongly computable.

These conditions are not sufficient, for stupid reasons at least: e.g., a thoroughly nonrecursive language without constants would have no closed terms, and thus no processes and an easily computable (*viz.*, empty) transition relation.

4 Bisimulation as a Congruence: White Chocolate

There are technical obstacles to defining bisimulation for chocolate languages in general. The definition that we would like is

Definition 4.1 A relation between closed terms of the same type, \smile , is a \forall bisimulation relation if it is symmetric and whenever $p \smile p'$:

- If $p, p' : \mathbf{P}$, then whenever $p \xrightarrow{a, \xi} q$, there is a $q' \sim q$ such that $p' \xrightarrow{a, \xi} q'$.
- If $p, p' : \xi_0 \times \xi_1$ then $\pi_i p \sim \pi_i p'$ for $i = 0, 1$.
- If $p, p' : \xi_0 \rightarrow \xi_1$, then for all $q \sim q'$ of type ξ_0 , $pq \sim p'q'$.

(This is a kind of late bisimulation.) We would then define \rightleftharpoons to be the largest \forall bisimulation relation. It works for first-order processes; the first-order parts of this definition are monotone, and thus there is indeed a well-defined notion of \rightleftharpoons . However, for higher-order processes, the universal quantifier in the $\xi_0 \rightarrow \xi_1$ case of the definition makes the definition non-monotone; the more processes of type ξ_0 which are related, the harder it is to be related at type $\xi_0 \rightarrow \xi_1$. It is an open question whether or not there is a maximum \forall bisimulation relation.

There are several variant definitions, replacing the universal quantifier at functional type with an existential: if $p \sim p' : \xi_0 \rightarrow \xi_1$, for every q there is a $q' \sim q$ such that $pq \sim p'q'$. This definition is monotone at least, so there is a well-defined \rightleftharpoons relation, but the weaker definition complicates the analysis. For this study, we restrict attention to a much smaller class of languages, scarcely more powerful than CHOCS. Definitions specific to this weaker class are termed “white chocolate.”

We take a restricted type system, the *White Chocolate Type System*:

$$\xi := \mathbf{P} | \mathbf{P}^n | \mathbf{P}^n \rightarrow \mathbf{P}$$

where $n \geq 0$. We do not bother to identify \mathbf{P} , \mathbf{P}^1 , and $\mathbf{P}^0 \rightarrow \mathbf{P}$. We define the natural generalization of Thomsen’s notion of higher-order bisimulation, which does not try to relate functions directly; instead, it applies them arguments in all ways and relates the resulting processes.

Let $\text{nf}(q)$ be the \Rightarrow -normal form of q .

Definition 4.2 A relation \sim between closed terms of type \mathbf{P} is a white bisimulation relation iff it is symmetric, and whenever $p \sim p'$, we have:

\mathbf{P} : If $p \xrightarrow{a, \mathbf{P}} q$, then there exists $q' \sim q$ such that $p' \xrightarrow{a, \mathbf{P}} q'$.

\mathbf{P}^k : If $p \xrightarrow{a, \mathbf{P}^k} q$, then there exists q' such that $p' \xrightarrow{a, \mathbf{P}^k} q'$, such that $\text{nf}(q)$ is componentwise \sim to $\text{nf}(q')$.

$\mathbf{P}^k \rightarrow \mathbf{P}$: If $p \xrightarrow{a, \mathbf{P}^k \rightarrow \mathbf{P}} q$, then for every closed $r : \mathbf{P}^k$, there is a r' componentwise \sim to r and a q' such that $p' \xrightarrow{a, \mathbf{P}^k \rightarrow \mathbf{P}} q'$ and $qr \sim q'r'$.

Processes p and p' are white bisimilar, $p \rightleftharpoons_{\text{white}} p'$, if there exists a white bisimulation relation relating them.

The definition of white bisimulation is monotone, and thus the usual basic bisimulation theory applies; e.g., $\rightleftharpoons_{\text{white}}$ is itself a white bisimulation. Note that $\rightleftharpoons_{\text{white}}$ is a rather peculiar bisimulation. Suppose that we have the following operation, where b is a fixed action, but a and ξ ranges over types.

$$\frac{x \xrightarrow{a, \xi} y, \quad x \xrightarrow{b, \mathbf{P}} y'}{f(x) \xrightarrow{a, \xi} y}$$

That is, $f(p)$ behaves like p if p has a b -transition at type \mathbf{P} on its first move, and $\mathbf{0}$ otherwise. Consider the processes

$$\begin{aligned} p &= (a?x.\mathbf{0}) + (a?x.x) = a(\lambda x.\mathbf{0}) + a(\lambda x.x) \\ p' &= p + (a?x.f(x)) = p + a(\lambda x.f(x)) \end{aligned}$$

Clearly, p' can evolve via an a -transition into the state $\lambda x.f(x)$, which is clearly different from both $\mathbf{0}$ and $\lambda x.x$. That is, p and p' are intuitively not bisimilar.

However, we have $p \rightleftharpoons_{\text{white}} p'$. Clearly both of p ’s moves match moves of p' , and two of p' ’s moves match moves of p . For any $r : \mathbf{P}$, we have $qr =_{\beta\pi} f(r)$, which is either bisimilar to r or to $\mathbf{0}$ depending on whether or not r can take a b action at \mathbf{P} . In particular, the usual intuitions of bisimulation, of “staying in related states on all actions,” have been lost in white bisimulation.⁵

Nonetheless, white bisimulation is a worthwhile notion. It enjoys many of the properties of bisimulation in ordinary process algebra; e.g., the method of “bisimulation up to $\rightleftharpoons_{\text{white}}$ ” works. More importantly, $\rightleftharpoons_{\text{white}}$ is a good notion of process equivalence on a suitable class of languages; that is, all suitably defined operations respect $\rightleftharpoons_{\text{white}}$.

Definition 4.3 A White Chocolate Language is a set of operations over a white chocolate type system, and a set of rules, subject to the following condition. For each rule ρ , let u be either the target of ρ (if the target has type \mathbf{P} or \mathbf{P}^k), or the normal form of $\text{target}(\rho)$ applied to a fresh variable z (if the target has type $\mathbf{P}^k \rightarrow \mathbf{P}$). Then, we require that

1. If a target variable y of type $\mathbf{P}^k \rightarrow \mathbf{P}$ appears in u , then it appears only once, and that occurrence is in a subterm of the form $y(p)$ for some term p .

⁵This may also be construed as a flaw in the higher-order transition system.

2. *There are no abstractions in u .*

For example, if we restrict NESTLÉ to white chocolate types, we get a white chocolate language. Consider the restriction rule. Ignoring some tupling and projecting that the actual system requires, we have the rule

$$\frac{x \xrightarrow{a, \mathbf{P}^1 \rightarrow \mathbf{P}} y}{x \setminus S \xrightarrow{a, \mathbf{P}^1 \rightarrow \mathbf{P}} \lambda w.(yw) \setminus S}$$

Then $u = (\lambda w.(yw) \setminus S)z =_{\beta\pi} (yz) \setminus S$, which satisfies the definition.

Theorem 4.4 *Let \mathcal{L} be a white chocolate language. White bisimulation is a congruence for \mathcal{L} .*

This subsumes, *e.g.*, Thomsen’s result that white bisimulation (his higher-order bisimulation) is a congruence for CHOCS; it also implies that white bisimulation is a congruence for Hennessy’s somewhat more elaborate language, and for NESTLÉ/2.

5 Acknowledgements

We would like to acknowledge Paul Taylor for his diagram package, which was helpful for drawing the many sorts of amazing arrows used in this study.

References

- [Ace92] Luca Aceto. Eliminating junk rules from GSOS languages. (Unpublished note; to appear, probably as part of something else.), July 1992.
- [BC90] Kenneth Birman and Robert Cooper. The Isis project: Real experience with a fault tolerant programming system. Technical Report 90-1138, Department of Computer Science, Cornell University, July 1990.
- [BCG91] Kenneth P. Birman, Robert Cooper, and Barry Gleeson. Programming with process groups: Group and multicast semantics. Technical Report 91-1185, Department of Computer Science, Cornell University, January 1991.
- [BIM88] Bard Bloom, Sorin Istrail, and Albert R. Meyer. Bisimulation can’t be traced (preliminary report). In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, pages 229–239, 1988. Also appears as MIT Technical Memo MIT/LCS/TM-345.
- [Blo89] Bard Bloom. *Ready Simulation, Bisimulation, and the Semantics of CCS-Like Languages*. PhD thesis, Massachusetts Institute of Technology, August 1989.
- [Hen93] Matthew Hennessy. A fully abstract denotational model for higher-order processes (extended abstract). In *Proceedings of the Eighth Annual IEEE Symposium on Logic in Computer Science*, pages 397–408. IEEE Computer Society Press, 1993. (Full version in Sussex TR 6/92).
- [Mil91] Robin Milner. The polyadic π -calculus: A tutorial. In *Proc. International Summer School on Logic and Algebra of Specification*. Marktoberdorf, 1991.
- [MPW92] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes I and II. *Information and Computation*, 100:1–40 and 41–77, 1992.
- [Rep89] John H. Reppy. First-class synchronous operations in Standard ML. Technical Report 89-1068, Department of Computer Science, Cornell University, December 1989.
- [Rep91] John H. Reppy. CML: A higher-order concurrent language. *SIGPLAN Notices*, 26(6):293–305, June 1991.
- [Tho89] Bent Thomsen. A calculus of higher order communicating systems. In *Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages*, pages 142–154, 1989.
- [Vaa93] Frits Vaandrager. Expressiveness results for process algebras. In de Bakker, de Roever, and Rozenberg, editors, *Semantics: Foundations and Applications*, pages 609–620. Springer-Verlag, 1993. LNCS 666. Also appears as CWI Tech Report CS-R9301.