

A Capability Calculus for Concurrency and Determinism ^{*}

Tachio Terauchi¹ and Alex Aiken²

¹ EECS Department, University of California, Berkeley

² Computer Science Department, Stanford University

Abstract. We present a capability calculus for checking partial confluence of channel-communicating concurrent processes. Our approach automatically detects more programs to be partially confluent than previous approaches and is able to handle a mix of different kinds of communication channels, including shared reference cells.

1 Introduction

Deterministic programs are easier to debug and verify than non-deterministic programs, both for testing (or simulation) and for formal methods. However, sometimes programs are written as communicating concurrent processes, for speed or for ease of programming, and therefore are possibly non-deterministic. In this paper, we present a system that can automatically detect more programs to be deterministic than previous methods [7, 10, 8, 9, 5]. Our system is able to handle programs communicating via a mix of different kinds of channels: rendezvous, output buffered, input buffered, and shared reference cells. Section 3.2 shows a few examples that can be checked by our system: producer consumer, token ring, and barrier synchronization. The companion technical report contains the omitted proofs [12].

We cast our system as a *capability calculus* [4]. The capability calculus was originally proposed as a framework for reasoning about resources in sequential computation, but has recently been extended to reason about determinism in concurrent programs [3, 11]. However, these systems can only reason about synchronization at join points, and therefore cannot verify determinism of channel-communicating processes. This paper extends the capability calculus to reason about synchronization due to channel communications. A key insight comes from our previous work [11] which showed that confluence can be ensured in a principled way from ordering dependencies between the side effects; dependencies are enforced by finding a flow assignment (which can be interpreted as *fractional capabilities* [3]) in the dependence graph.

2 Preliminaries

We focus on the simple concurrent language shown in Figure 1. A program, p , is a parallel composition of finitely many processes. A process, s , is a sequential

^{*} This research was supported in part by NSF Grant No. CCR-0326577.

$p ::= s_1 s_2 \dots s_n$	(<i>program</i>)	$s ::= s_1; s_2$	(<i>sequence</i>)
$e ::= c$	(<i>channel</i>)	if e then s_1 else s_2	(<i>branch</i>)
x	(<i>local variable</i>)	while e do s	(<i>loop</i>)
n	(<i>integer constant</i>)	skip	(<i>skip</i>)
$e_1 \text{ op } e_2$	(<i>integer operation</i>)	$x := e$	(<i>assignment</i>)
		$!(e_1, e_2)$	(<i>write channel</i>)
		$?(e, x)$	(<i>read channel</i>)

Fig. 1. The syntax of the small concurrent language.

statement consisting of the usual imperative features as well as channel communication operations. Here, $!(e_1, e_2)$ means writing the value of e_2 to the channel e_1 , and $?(e, x)$ means storing the value read from the channel e to the variable x . The variables are process-local, and so the only means of communication are channel reads and writes. We use meta-variables x, y, z , etc. for variables and c, d , etc. for channels.

The language cannot dynamically create channels or spawn new processes, but these restrictions are imposed only to keep the main presentation to the novel features of the system. Section 3.3 shows that techniques similar to previous work in the capability calculus can be used to handle dynamic channels and processes.

2.1 Channel Kinds

The literature on concurrency includes several forms of channels with distinct semantics. We introduce these channel kinds and show how they affect determinism.

If c and d are *rendezvous* channels, then the following program is deterministic³ because $(x, y) = (1, 2)$ when the process terminates:

$$!(c, 1); !(d, 2) || !(d, 3); ?(c, x) || ?(d, y); ?(d, y)$$

The same program is non-deterministic if c is *output buffered* because $!(c, 1)$ does not need to wait for the reader $?(c, x)$, and therefore (x, y) could be $(1, 2)$ or $(1, 3)$.

While all the processes share one output buffer per channel, each process has its own input buffer per channel. Therefore, $!(c, 1); !(c, 2) || ?(c, x) || ?(c, y)$ is deterministic if c is input buffered but not if c is output buffered or rendezvous. Input buffered channels are the basis of Kahn process networks [7].

We also consider a buffered channel whose buffer is overwritten by every write but never modified by a read. Such a channel is equivalent to a reference cell. If c is a reference cell, $!(c, 1); !(c, 2) || ?(c, x)$ is not deterministic because $!(c, 2)$ may or may-not overwrite 1 in the buffer before $?(c, x)$ reads the buffer. The program is deterministic if c is any other channel kind. On the other hand,

³ Here, we use the term informally. Determinism is formally defined in Section 2.2.

$$\begin{array}{c}
\frac{(S(i), e) \Downarrow n \quad n \neq 0}{(B, S, i.(\text{if } e \text{ then } s_1 \text{ else } s_2); s||p) \rightarrow (B, S, i.s_1; s||p)} \quad \mathbf{IF1} \\
\frac{(S(i), e) \Downarrow 0}{(B, S, i.(\text{if } e \text{ then } s_1 \text{ else } s_2); s||p) \rightarrow (B, S, i.s_2; s||p)} \quad \mathbf{IF2} \\
\frac{(S(i), e) \Downarrow n \quad n \neq 0}{(B, S, i.(\text{while } e \text{ do } s_1); s||p) \rightarrow (B, S, i.s_1; (\text{while } e \text{ do } s_1); s||p)} \quad \mathbf{WHILE1} \\
\frac{(S(i), e) \Downarrow 0}{(B, S, i.(\text{while } e \text{ do } s_1); s||p) \rightarrow (B, S, i.s||p)} \quad \mathbf{WHILE2} \\
\frac{(S(i), e) \Downarrow e' \quad S' = S[i \mapsto S(i) :: (x, e')]}{(B, S, i.x := e; s||p) \rightarrow (B, S', i.s||p)} \quad \mathbf{ASSIGN} \\
\frac{\begin{array}{c} (S(i), e_1) \Downarrow c \quad (S(i), e_2) \Downarrow e'_2 \quad (S(j), e_3) \Downarrow c \\ \neg \text{buffered}(c) \quad S' = S[j \mapsto S(j) :: (x, e'_2)] \end{array}}{(B, S, i.!(e_1, e_2); s_1||j.?(e_3, x); s_2||p) \rightarrow (B, S', i.s_1||j.s_2||p)} \quad \mathbf{UNBUF} \\
\frac{(S(i), e_1) \Downarrow c \quad (S(i), e_2) \Downarrow e'_2 \quad \text{buffered}(c) \quad B' = B.\text{write}(c, e'_2)}{(B, S, i.!(e_1, e_2); s||p) \rightarrow (B', S, i.s||p)} \quad \mathbf{BUF1} \\
\frac{\begin{array}{c} (S(i), e) \Downarrow c \quad \text{buffered}(c) \\ (B', e') = B.\text{read}(c, i) \quad S' = S[i \mapsto S(i) :: (x, e')] \end{array}}{(B, S, i.?(e, x); s||p) \rightarrow (B', S', i.s||p)} \quad \mathbf{BUF2}
\end{array}$$

Fig. 2. The operational semantics of the small concurrent language.

$!(c, 1); !(c, 2); !(d, 3); ?(c, x) \parallel ?(d, x); ?(c, y)$ is deterministic if c is a reference cell and d is rendezvous because both reads of c happen after $!(c, 2)$ overwrites the buffer. But the program is not deterministic if c is output buffered.

2.2 Operational Semantics

The operational semantics of the language is defined as a series of reductions from states to states. A state is represented by the triple (B, S, p) where B is a buffer, S is a store, and p is a program such that each concurrent process in p is indexed by a process number, i.e., $p ::= 1.s_1 \parallel 2.s_2 \parallel \dots \parallel n.s_n$. Indexes are used to connect a process to its input buffer and its store.

A store is a mapping from process indexes to histories of assignments where a *history* is a sequence of pairs (x, e) , meaning e was assigned to x . We use meta-variables h, h' , etc. for histories. Let $::$ be append. A lookup in a history is defined as: $(h :: (x, e))(x) = e$ and $(h :: (y, e))(x) = h(x)$ if $y \neq x$. We use history instead of memory for the purpose of defining determinism.

Expressions are evaluated entirely locally. The semantics of expressions are defined as: $(h, c) \Downarrow c$, $(h, x) \Downarrow h(x)$, $(h, n) \Downarrow n$, and $(h, e_1 \text{ op } e_2) \Downarrow e'_1 \text{ op } e'_2$ if $(h, e_1) \Downarrow e'_1$ and $(h, e_2) \Downarrow e'_2$.

$$\begin{aligned}
B.write(c, e) &= \begin{cases} B[c \mapsto enq(B(c), e)] & \text{if } c \text{ is output buffered} \\ B[c \mapsto \langle enq(q_1, e), \dots, enq(q_n, e) \rangle] & \text{if } c \text{ is input buffered} \\ \quad \text{where } B(c) = \langle q_1, \dots, q_n \rangle \\ B[c \mapsto e] & \text{if } c \text{ is a reference cell} \end{cases} \\
B.read(c, i) &= \begin{cases} (B[c \mapsto q'], e) & \text{if } c \text{ is output buffered} \\ \quad \text{where } B(c) = q \text{ and } (q', e) = deq(q) \\ (B[c \mapsto \langle q_1, \dots, q'_i, \dots, q_n \rangle], e) & \text{if } c \text{ is input buffered} \\ \quad \text{where } B(c) = \langle q_1, \dots, q_i, \dots, q_n \rangle \\ \quad \quad (q'_i, e) = deq(q_i) \\ (B, B(c)) & \text{if } c \text{ is a reference cell} \end{cases}
\end{aligned}$$

Fig. 3. Buffer operations.

Figure 2 shows the reduction rules. Programs are equivalent up to re-ordering of parallel processes, e.g., $p_1 || p_2 = p_2 || p_1$. If p is an empty program (i.e., p contains 0 processes), then $p' || p = p'$. Also, we let $s = s; \text{skip} = \text{skip}; s$. Note that the rules only reduce the left-most processes, and so we rely on process re-ordering to reduce other processes. The rules **IF1**, **IF2**, **WHILE1**, and **WHILE2** do not involve channel communication and are self-explanatory. **AS-SIGN** is also a process-local reduction because variables are local. Here, $S[i \mapsto h]$ means $\{j \mapsto S(j) \mid j \neq i \wedge j \in \text{dom}(S)\} \cup \{i \mapsto h\}$. We use the same notation for other mappings.

UNBUF handles communication over rendezvous channels. The predicate $\neg buffered(c)$ says c is unbuffered (and therefore rendezvous). Note that the written value e'_2 is immediately transmitted to the reader. **BUF1** and **BUF2** handle communication over buffered channels, which include output buffered channels, input buffered channels, and reference cells. The predicate $buffered(c)$ says that c is a buffered channel. We write $B.write(c, e'_2)$ for the buffer B after e'_2 is written to the channel c , and $B.read(c, i)$ for the pair (B', e') where e' is the value process i read from channel c and B' is the buffer after the read.

Formally, a buffer B is a mapping from channels to buffer contents. If c is a rendezvous channel, then $B(c) = nil$ indicating that c is not buffered. If c is output buffered, then $B(c) = q$ where q is a FIFO queue of values. If c is input buffered, then $B(c) = \langle q_1, q_2, \dots, q_n \rangle$, i.e., a sequence of FIFO queues where each q_i represents the buffer content for process i . If c is a reference cell, then $B(c) = e$ for some value e . Let $enq(q, e)$ be q after e is enqueued. Let $deq(q)$ be the pair (q', e) where q' is q after e is dequeued. Buffer writes and reads are defined as shown in Figure 3. Note that $B.read(c, i)$ and $B.write(c, e)$ are undefined if c is rendezvous.

We write $P \rightarrow^* Q$ for 0 or more reduction steps from P to Q . We define partial confluence and determinism.

Definition 1. Let Y be a set of channels. We say that P is partially confluent with respect to Y if for any $P \rightarrow^* P_1$ communicating only over channels in Y , and for any $P \rightarrow^* P_2$, there exists a state Q such that $P_2 \rightarrow^* Q$ communicating only over channels in Y and $P_1 \rightarrow^* Q$.

Definition 2. Let Y be a set of channels. We say that P is deterministic with respect to Y if for each process index i , there exists a (possibly infinite) sequence h_i such that for any $P \rightarrow^* (B, S, p)$ that communicates only over channels in Y , $S(i)$ is a prefix of h_i .

Determinism implies that for any single process, interaction with the rest of the program is deterministic. Determinism and partial confluence are related in the following way.

Lemma 1. If P is partially confluent with respect to Y then P is deterministic with respect to Y .

Note that the definitions are sufficient for programs interacting with the environment because an environment can be modeled as a process using integer operators with unknown (but deterministic) semantics.

3 Calculus of Capabilities

We now present a capability calculus for ensuring partial confluence. While capability calculi are typically presented as a type system in the literature, we take a different approach and present the capability calculus as a dynamic system. We then construct a type system to statically reason about the dynamic capability calculus. This approach allows us to distinguish approximations due to the type abstraction from approximations inherent in the capability concept. (We have taken a similar approach in previous work [11].)

We informally describe the general idea. To simplify matters, we begin this initial discussion with rendezvous channels and total confluence. Given a program, the goal is to ensure that for each channel c , at most one process can write c and at most one process can read c at any point in time. To this end, we introduce capabilities $r(c)$ and $w(c)$ such that a process needs $r(c)$ to read from c and $w(c)$ to write to c . Capabilities are distributed to the processes at the start of the program and are not allowed be duplicated.

Recall the following confluent program from Section 2:

$$1.!(c, 1); !(d, 2) \parallel 2.!(d, 3); ?(c, x) \parallel 3.?(d, y); ?(d, y)$$

Note that for both c and d , at most one process can read and at most one process can write at any point in time. However, because both process 1 and process 2 write to d , they must somehow share $w(d)$. A novel feature of our capability calculus is the ability to pass capabilities between processes. The idea is to let capabilities be passed when the two processes synchronize, i.e., when the processes communicate over a channel. In our example, we let process 2

have $w(d)$ at the start of the program. Then, when process 1 and process 2 communicate over c , we pass $w(d)$ from process 2 to process 1 so that process 1 can write to d .

An important observation is that capability passing works in this example because $!(d, 3)$ is guaranteed to occur before the communication on c due to c being rendezvous. If c is buffered (recall that the program is not confluent in this case), then $!(c, 1)$ may occur before $!(d, 3)$. Therefore, process 1 cannot obtain $w(d)$ from process 2 when c is written because process 2 may still need $w(d)$ to write on d . In general, for a buffered channel, while the read is guaranteed to occur after the write, there is no ordering dependency in the other direction, i.e., from the read to the write. Therefore, capabilities can be passed from the writer to the reader but not vice versa, whereas capabilities can be passed in both directions when communicating over a rendezvous channel.

Special care is needed for reference cells. If c is a reference cell, the program $1.!(c, 1); !(c, 2) || 2.?(c, x)$ is not deterministic although process 1 is the only writer and process 2 is the only reader. We use *fractional capabilities* [3, 11] such that a read capability is a fraction of the write capability. Capabilities can be split into multiple fractions, which allows concurrent reads on the same reference cell, but must be re-assembled to form the write capability. Fractional capabilities can be passed between processes in the same way as other capabilities. Recall the following confluent program from Section 2 where c is a reference cell and d is rendezvous:

$$1.!(c, 1); !(c, 2); !(d, 3); ?(c, x) || 2.?(d, x); ?(c, y)$$

Process 1 must start with the capability to write c . Because both processes read from c after communicating over d , we split the capability for c such that one half of the capability stays in process 1 and the other half is passed to process 2 via d . As a result, both processes obtain the capability to read from c . We have shown previously that fractional capabilities can be derived in a principled way from ordering dependencies [11].

We now formally present our capability calculus. Let

$$\begin{aligned} \mathbf{Capabilities} = & \{w(c), r(c) \mid c \text{ is rendezvous or output buffered}\} \\ & \cup \{w(c) \mid c \text{ is input buffered}\} \cup \{w(c) \mid c \text{ is a reference cell}\} \end{aligned}$$

A *capability set* C is a function from **Capabilities** to rational numbers in the range $[0, 1]$. If c is rendezvous, output buffered, or input buffered, $C(w(c)) = 1$ (resp. $C(r(c)) = 1$) means that the capability to write (resp. read) c is in C . Read capabilities are not needed for input buffered channels because each process has its own buffer. For reference cells, $C(w(c)) = 1$ means that the capability to write is in C , whereas $C(w(c)) > 0$ means that the capability to read is in C . To summarize, we define the following predicates:

$$\begin{aligned} hasWcap(C, c) & \Leftrightarrow C(w(c)) = 1 \\ hasRcap(C, c) & \Leftrightarrow \begin{cases} C(r(c)) = 1 & \text{if } c \text{ is rendezvous or output buffered} \\ true & \text{if } c \text{ is input buffered} \\ C(w(c)) > 0 & \text{if } c \text{ is reference cell} \end{cases} \end{aligned}$$

$$\begin{array}{c}
\frac{(S(i), e) \Downarrow n \qquad n \neq 0}{(X, B, S, i.C.(\mathbf{if} \ e \ \mathbf{then} \ s_1 \ \mathbf{else} \ s_2); s||p) \rightarrow (X, B, S, i.C.s_1; s||p)} \quad \mathbf{IF1}' \\
\frac{(S(i), e) \Downarrow 0}{(X, B, S, i.C.(\mathbf{if} \ e \ \mathbf{then} \ s_1 \ \mathbf{else} \ s_2); s||p) \rightarrow (X, B, S, i.C.s_2; s||p)} \quad \mathbf{IF2}' \\
\frac{(S(i), e) \Downarrow n \qquad n \neq 0}{(X, B, S, i.C.(\mathbf{while} \ e \ \mathbf{do} \ s_1); s||p) \rightarrow (X, B, S, i.C.s_1; (\mathbf{while} \ e \ \mathbf{do} \ s_1); s||p)} \quad \mathbf{WHILE1}' \\
\frac{(S(i), e) \Downarrow 0}{(X, B, S, i.C.(\mathbf{while} \ e \ \mathbf{do} \ s_1); s||p) \rightarrow (X, B, S, i.C.s||p)} \quad \mathbf{WHILE2}' \\
\frac{(S(i), e) \Downarrow e' \qquad S' = S[i \mapsto S(i) :: (x, e')]}{(X, B, S, i.C.x := e; s||p) \rightarrow (X, B, S', i.C.s||p)} \quad \mathbf{ASSIGN}'
\end{array}$$

Fig. 4. The capability calculus: sequential reductions.

To denote capability merging and splitting, we define:

$$C_1 + C_2 = \{cap \mapsto C_1(cap) + C_2(cap) \mid cap \in \mathbf{Capabilities}\}$$

We define $C_1 - C_2 = C_3$ if $C_1 = C_3 + C_2$. (We avoid negative capabilities.)

Figure 4 and Figure 5 show the reduction rules of the capability calculus. The reduction rules (technically, labeled transition rules) are similar to those of operational semantics with the following differences.

Each concurrent process is prefixed by a capability set C representing the current capabilities held by the process. The rules in Figure 4 do not utilize capabilities (i.e., capabilities are only passed sequentially) and are self-explanatory. Figure 5 shows how capabilities are utilized at communication points. **UNBUF'** sends capabilities C from the writer process to the reader process and sends capabilities C' from the reader process to the writer process. **UNBUF'** checks whether the right capabilities are present by $hasWcap(C_i, c) \wedge hasRcap(C_j, c)$. The label ℓ records whether the check succeeds. Because we are interested in partial confluence with respect to some set Y of channels, we only check the capabilities if $c \in Y$. To this end, the predicate $confch()$ parameterizes the system so that $confch(c)$ iff $c \in Y$.

BUF1' and **BUF2'** handle buffered communication. Recall that the writer can pass capabilities to the reader. **BUF1'** takes capabilities C' from the writer process and stores them in X . **BUF2'** takes capabilities C' from X and gives them to the reader process. The mapping X from channels to capability sets maintains the capabilities stored in each channel.

We now formally state when our capability calculus guarantees partial confluence. Let $erase((X, B, S, 1.C_1.s_1 || \dots || n.C_n.s_n)) = (B, S, 1.s_1 || \dots || n.s_n)$, i.e., $erase()$ erases all capability information from the state. We use meta-variables P, Q, R , etc. for states in the operational semantics and underlined meta-variables $\underline{P}, \underline{Q}, \underline{R}$, etc. for states in the capability calculus.

$$\begin{array}{c}
\frac{(S(i), e_1) \Downarrow c \quad \neg \text{buffered}(c) \quad \ell = (\text{confch}(c) \Rightarrow (\text{hasWcap}(C_i, c) \wedge \text{hasRcap}(C_j, c)))}{(X, B, S, i.C_i.(e_1, e_2); s_1 || j.C_j?(e_3, x); s_2 || p) \xrightarrow{\ell} (X, B, S', i.(C_i - C + C').s_1 || j.(C_j + C - C').s_2 || p)} \quad \text{UNBUF}' \\
\\
\frac{(S(i), e_1) \Downarrow c \quad B' = B.\text{write}(c, e_2) \quad (S(i), e_2) \Downarrow e_2' \quad \ell = (\text{confch}(c) \Rightarrow \text{hasWcap}(C, c)) \quad \text{buffered}(c)}{(X, B, S, i.C.(e_1, e_2); s || p) \xrightarrow{\ell} (X[c \mapsto X(c) + C'], B', S, i.(C - C').s || p)} \quad \text{BUF1}' \\
\\
\frac{(S(i), e) \Downarrow c \quad S' = S[i \mapsto S(i) :: (x, e')] \quad \text{buffered}(c) \quad \ell = (\text{confch}(c) \Rightarrow \neg \text{hasRcap}(C, c)) \quad (B', e') = B.\text{read}(c, i)}{(X, B, S, i.C?(e, x); s || p) \xrightarrow{\ell} (X[c \mapsto X(c) - C'], B', S', i.(C + C').s || p)} \quad \text{BUF2}'
\end{array}$$

Fig. 5. The capability calculus: communication reductions.

A *well-formed state* is a state in the capability calculus that does not carry duplicated capabilities. More formally,

Definition 3. Let $\underline{P} = (X, B, S, 1.C_1.s_1 || \dots || n.C_n.s_n)$. Let $C = \sum_{i=1}^n C_i + \sum_{c \in \text{dom}(X)} X(c)$. We say \underline{P} is *well-formed* if for all $\text{cap} \in \text{dom}(C)$, $C(\text{cap}) = 1$.

We define *capability-respecting states*. Informally, \underline{P} is capability respecting with respect to a set of channels Y if for any sequence of reductions from $\text{erase}(\underline{P})$, there exists a strategy to pass capabilities between the processes such that every communication over the channels in Y occurs under the appropriate capabilities. More formally,

Definition 4. Let Y be a set of channels and let $\text{confch}(c) \Leftrightarrow c \in Y$. Let M be a set of states in the capability calculus. M is said to be *capability-respecting with respect to Y* if for any $\underline{P} \in M$,

- \underline{P} is well-formed, and
- for any state Q such that $\text{erase}(\underline{P}) \rightarrow Q$, there exists $\underline{Q} \in M$ such that $\text{erase}(\underline{Q}) = Q$, $\underline{P} \xrightarrow{\ell} \underline{Q}$, and if ℓ is not empty then $\ell = \text{true}$.

We now state the main claim of this section.

Theorem 1. Let P be a state. Suppose there exists M such that M is capability-respecting with respect to Y and there exists $\underline{P} \in M$ such that $\text{erase}(\underline{P}) = P$. Then P is partially confluent with respect to Y .

3.1 Static Checking of Capabilities

Theorem 1 tells us that to ensure that P is partially confluent, it is sufficient to find a capability-respecting set containing some \underline{P} such that $\text{erase}(\underline{P}) = P$.⁴

⁴ It is not a necessary condition, however. For example, $!(c, 1) || !(c, 1) || ?(c, x) || ?(c, x)$ is confluent but does not satisfy the condition.

Ideally, we would like to use the largest capability-respecting set, but such a set is not recursive (because it is reducible from the halting problem). Instead, we use a type system to compute a safe approximation of the set.

We define four kinds of channel types, one for each channel kind.

$$\begin{array}{l|l} \tau ::= ch(\rho, \tau, \Psi_1, \Psi_2) & (\textit{rendezvous}) \\ | ch(\rho, \tau, \Psi) & (\textit{output buffered}) \\ | ch(\rho, \tau, \langle \Psi_1, \dots, \Psi_n \rangle) & (\textit{input buffered}) \\ | ch(\rho, \tau) & (\textit{reference cell}) \\ | int & (\textit{integers}) \end{array}$$

Meta-variables ρ, ρ' , etc. are *channel handles*. Let **Handles** be the set of channel handles. Let **StaticCapabilities** = $\{w(\rho), r(\rho) \mid \rho \in \mathbf{Handles}\}$. Meta-variables Ψ, Ψ' , etc. are mappings from **StaticCapabilities** to $[0, 1]$. We call such a mapping a *static capability set*. The rendezvous channel type can be read as follows: the channel communicates values of type τ , any writer of the channel sends capabilities Ψ_1 , and any reader of the channel sends capabilities Ψ_2 . For an output buffered channel, because readers cannot send capabilities, only one static capability set, Ψ , is present in its type. For an input buffered channel, the sequence $\langle \Psi_1, \dots, \Psi_n \rangle$ lists capabilities such that each process i gets Ψ_i from a read. Because a value stored in a reference cell may be read arbitrarily many times, our type system cannot statically reason about processes passing capabilities through reference cells, and so a reference cell type does not carry any static capability set.

Additions and subtractions of static capabilities are analogous to those of (actual) capabilities:

$$\begin{aligned} \Psi_1 + \Psi_2 &= \{cap \mapsto \Psi_1(cap) + \Psi_2(cap) \mid cap \in \mathbf{StaticCapabilities}\} \\ \Psi_1 - \Psi_2 &= \Psi_3 \quad \text{if } \Psi_1 = \Psi_3 + \Psi_2 \end{aligned}$$

We say $\Psi_1 \geq \Psi_2$ if there exists Ψ_3 such that $\Psi_1 = \Psi_2 + \Psi_3$.

For channel type τ , $hdl(\tau)$ is the handle of the channel, and $valtype(\tau)$ is the type of the communicated value. That is, $hdl(ch(\rho, \dots)) = \rho$ and $valtype(ch(\rho, \tau, \dots)) = \tau$. Also, $writeSend(\tau)$ (resp. $readSend(\tau)$) is the set of capabilities sent by a writer (resp. reader) of the channel. More formally,

$$\begin{aligned} writeSend(ch(\rho, \tau, \Psi_1, \Psi_2)) &= \Psi_1 \\ writeSend(ch(\rho, \tau, \Psi)) &= \Psi \\ writeSend(ch(\rho, \tau, \langle \Psi_1, \dots, \Psi_n \rangle)) &= \sum_{i=1}^n \Psi_i \\ writeSend(ch(\rho, \tau)) &= 0 \\ readSend(\tau) &= \begin{cases} \Psi_2 & \text{if } \tau = ch(\rho, \tau', \Psi_1, \Psi_2) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

(0 is the constant zero function $\lambda x.0$.) Similarly, $writeRecv(\tau)$ (resp. $readRecv(\tau, i)$) is the set of capabilities received by the writer (resp. the reader process i):

$$\begin{aligned} writeRecv(\tau) &= readSend(\tau) \\ readRecv(\tau, i) &= \begin{cases} \Psi_i & \text{if } \tau = ch(\rho, \tau, \langle \Psi_1, \dots, \Psi_n \rangle) \\ writeSend(\tau) & \text{otherwise} \end{cases} \end{aligned}$$

$$\begin{array}{c}
\frac{\Gamma, i, \Psi \vdash s_1 : \Psi_1 \quad \Gamma, i, \Psi_1 \vdash s_2 : \Psi_2}{\Gamma, i, \Psi \vdash s_1; s_2 : \Psi_2} \text{ SEQ} \\
\\
\frac{\Gamma \vdash e : \text{int} \quad \Gamma, i, \Psi \vdash s_1 : \Psi_1 \quad \Gamma, i, \Psi \vdash s_2 : \Psi_2 \quad \Psi_1 \geq \Psi_3 \quad \Psi_2 \geq \Psi_3}{\Gamma, i, \Psi \vdash \text{if } e \text{ then } s_1 \text{ else } s_2 : \Psi_3} \text{ IF} \\
\\
\frac{\Gamma \vdash e : \text{int} \quad \Gamma, i, \Psi_1 \vdash s : \Psi_2 \quad \Psi_2 \geq \Psi_1 \quad \Psi \geq \Psi_1}{\Gamma, i, \Psi \vdash \text{while } e \text{ do } s : \Psi_2} \text{ WHILE} \\
\\
\frac{}{\Gamma, i, \Psi \vdash \text{skip} : \Psi} \text{ SKIP} \quad \frac{\Gamma \vdash e : \Gamma(x)}{\Gamma, i, \Psi \vdash x := e : \Psi} \text{ ASSIGN} \\
\\
\frac{\Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \text{valtype}(\tau) \quad \text{confch}(\tau, \Gamma) \Rightarrow \text{hasWcap}(\Psi, \tau)}{\Gamma, i, \Psi \vdash !(e_1, e_2) : \Psi - \text{writeSend}(\tau) + \text{writeRecv}(\tau)} \text{ WRITE} \\
\\
\frac{\Gamma \vdash e : \tau \quad \Gamma(x) = \text{valtype}(\tau) \quad \text{confch}(\tau, \Gamma) \Rightarrow \text{hasRcap}(\Psi, \tau)}{\Gamma, i, \Psi \vdash ?(e, x) : \Psi - \text{readSend}(\tau) + \text{readRecv}(\tau, i)} \text{ READ}
\end{array}$$

Fig. 6. Type checking rules.

Note that the writer of the input buffered channel $ch(\rho, \tau, \langle \Psi_1, \dots, \Psi_n \rangle)$ must be able to send the sum of all capabilities to be received by each process (i.e., $\sum_{i=1}^n \Psi_i$), whereas the reader receives only its own share (i.e., Ψ_i).

For channel type τ , $\text{hasWcap}(\Psi, \tau)$ and $\text{hasRcap}(\Psi, \tau)$ are the static analog of $\text{hasWcap}(C, c)$ and $\text{hasRcap}(C, c)$. More formally,

$$\begin{array}{l}
\text{hasWcap}(\Psi, \tau) \Leftrightarrow \Psi(w(\text{hdl}(\tau))) = 1 \\
\text{hasRcap}(\Psi, \tau) \Leftrightarrow \begin{cases} \Psi(r(\text{hdl}(\tau))) = 1 & \text{if } \tau \text{ is rendezvous or output buffered} \\ \text{true} & \text{if } \tau \text{ is input buffered} \\ \Psi(w(\text{hdl}(\tau))) > 0 & \text{if } \tau \text{ is reference cell} \end{cases}
\end{array}$$

A *type environment* Γ is a mapping from channels and variables to types such that for each channel c and d ,

- the channel type kind of $\Gamma(c)$ coincides with the channel kind of c , and
- if $c \neq d$ then $\text{hdl}(\Gamma(c)) \neq \text{hdl}(\Gamma(d))$, i.e., each handle ρ uniquely identifies a channel. (Section 3.3 discusses a way to relax this restriction.)

We sometimes write $\Gamma[c]$ to mean $\text{hdl}(\Gamma(c))$.

Expressions are type-checked as follows:

$$\frac{}{\Gamma \vdash c : \Gamma(c)} \quad \frac{}{\Gamma \vdash x : \Gamma(x)} \quad \frac{}{\Gamma \vdash n : \text{int}} \quad \frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 \text{ op } e_2 : \text{int}}$$

Figure 6 shows the type checking rules for statements. The judgments are of the form $\Gamma, i, \Psi \vdash s : \Psi'$ where i is the index of the process where s appears in, Ψ the capabilities before s , and Ψ' the capabilities after s . **SEQ**, **IF**, **WHILE**,

SKIP, and **ASSIGN** are self-explanatory. **WRITE** handles channel writes and **READ** handles channel reads. Here, $\text{confch}(\tau, \Gamma)$ is defined as:

$$\text{confch}(\tau, \Gamma) \Leftrightarrow \exists c. \Gamma[c] = \text{hdl}(\tau) \wedge \text{confch}(c)$$

We write $\Gamma \vdash B(c)$ if the buffer $B(c)$ is well-typed, i.e., $\Gamma \vdash e : \text{valtype}(\Gamma(c))$ for each value e stored in the buffer $B(c)$. We write $\Gamma \vdash h$ if the history h is well-typed, i.e., $\Gamma \vdash h(x) : \Gamma(x)$ for each $x \in \text{dom}(h)$. We write $\Gamma \vdash C : \Psi$ if Ψ represents C , i.e., for each $w(c) \in \text{dom}(C)$, $\Psi(w(\Gamma[c])) = C(w(c))$ and for each $r(c) \in \text{dom}(C)$, $\Psi(r(\Gamma[c])) = C(r(c))$.

Let $\underline{P} = (B, X, S, 1.C_1.s_1 \parallel \dots \parallel n.C_n.s_n)$. An *environment* for \underline{P} consists of a type environment Γ for typing the channels, a type environment Γ_i for typing each process i , the starting static capability Ψ_i for each process i , and the mapping W from handles to static capabilities that represents X . We say \underline{P} is well-typed under the environment $(\Gamma, \Gamma_1, \dots, \Gamma_n, \Psi_1, \dots, \Psi_n, W)$, written $(\Gamma, \Gamma_1, \dots, \Gamma_n, \Psi_1, \dots, \Psi_n, W) \vdash \underline{P}$, if

- For each c , $\Gamma \vdash B(c)$.
- For each i , $\Gamma_i \supseteq \Gamma$, $\Gamma_i \vdash S(i)$, $\Gamma \vdash C_i : \Psi_i$, and $\Gamma_i, i, \Psi_i \vdash s_i : \Psi'_i$ for some Ψ'_i .
- For each c , $\Gamma \vdash X(c) : W(\Gamma[c])$, i.e., W is a static representation of X .
- Let $\Psi_{\text{total}} = \sum_{i=1}^n \Psi_i + \sum_{\rho \in \text{dom}(W)} W(\rho)$. Then for each $\text{cap} \in \text{dom}(\Psi_{\text{total}})$, $\Psi_{\text{total}}(\text{cap}) = 1$, i.e., there are no duplicated capabilities.
- For all output buffered channels c , $W(\Gamma[c]) = |B(c)| \times \text{writeSend}(\Gamma(c))$. For all input buffered channels c , $W(\Gamma[c]) = \sum_{i=1}^n |B(c).i| \times \text{readRecv}(\Gamma(c), i)$.

In the last condition, $|B(c)|$ denotes the length of the queue $B(c)$, and $|B(c).i|$ denotes the length of the queue for process i (for input buffered channels). The condition ensures that there are enough capabilities in X for buffered reads. We now state the main claim of this section.

Theorem 2. *Let Y be a set of channels and let $\text{confch}(c) \Leftrightarrow c \in Y$. Let $M = \{\underline{P} \mid \exists \text{Env}. \text{Env} \vdash \underline{P}\}$. Then M is capability-respecting with respect to Y .*

Theorem 2 together with Theorem 1 implies that to check if P is confluent, it suffices to find a well-typed \underline{P} such that $P = \text{erase}(\underline{P})$. More formally,

Corollary 1. *Let Y be a set of channels and let $\text{confch}(c) \Leftrightarrow c \in Y$. P is partially-confluent and deterministic with respect to Y if there exists \underline{P} and Env such that $P = \text{erase}(\underline{P})$ and $\text{Env} \vdash \underline{P}$.*

The problem of finding \underline{P} and Env such that $P = \text{erase}(\underline{P})$ and $\text{Env} \vdash \underline{P}$ can be formulated as linear inequality constraints satisfaction problem. The details are similar to the type inference algorithm from our previous work [11]. The constraints can be generated in time polynomial in the size of P , which can then be solved efficiently by a linear programming algorithm.

3.2 Examples

Producer Consumer: Let c be an output buffered channel. The program $1.\text{while } 1 \text{ do } !(c, 1) \parallel 2.\text{while } 1 \text{ do } ?(c, x)$ is a simple but common communication pattern of sender and receiver processes being fixed for each channel; no capabilities need to be passed between processes. The type system can prove confluence by assigning the starting capabilities $\theta[w(\rho) \mapsto 1]$ to process 1 and $\theta[r(\rho) \mapsto 1]$ to process 2 where $c : ch(\rho, int, \theta)$.

Token Ring: Let c_1, c_2, c_3 be rendezvous and d be output buffered. The program below models a token ring where processes 1, 2, and 3 take turns writing to d :

```

1.while 1 do (?( $c_3, x$ ); !( $d, 1$ ); !( $c_1, 0$ ))
|| 2.while 1 do (?( $c_1, x$ ); !( $d, 2$ ); !( $c_2, 0$ ))
|| 3.!( $c_3, 0$ ); while 1 do (?( $c_2, x$ ); !( $d, 3$ ); !( $c_3, 0$ ))
|| 4.while 1 do ?( $d, y$ )

```

Recall that variables x and y are process local. The type system can prove confluence by assigning the channel d the type $ch(\rho_d, int, \theta)$ and each c_i the type $ch(\rho_{c_i}, int, \theta[w(\rho_d) \mapsto 1], \theta)$, which says that a write to c_i sends $w(d)$ to the reader. The starting capabilities are $\theta[r(\rho_{c_3}) \mapsto 1, w(\rho_{c_1}) \mapsto 1]$ for process 1, $\theta[r(\rho_{c_1}) \mapsto 1, w(\rho_{c_2}) \mapsto 1]$ for process 2, $\theta[r(\rho_{c_2}) \mapsto 1, w(\rho_{c_3}) \mapsto 1, w(\rho_d) \mapsto 1]$ for process 3, and $\theta[r(\rho_d) \mapsto 1]$ for process 4.

Barrier Synchronization: Let c_1, c_2, c_3 be reference cells. Let $d_1, d_2, d_3, d'_1, d'_2, d'_3$ be input buffered channels. Consider the following program:

```

1.while 1 do (!( $c_1, e_1$ ); !( $d_1, 0$ ); BR; ?( $c_1, y$ ); ?( $c_2, z$ ); ?( $c_3, w$ ); !( $d'_1, 0$ ); BR')
|| 2.while 1 do (!( $c_2, e_2$ ); !( $d_2, 0$ ); BR; ?( $c_1, y$ ); ?( $c_2, z$ ); ?( $c_3, w$ ); !( $d'_2, 0$ ); BR')
|| 3.while 1 do (!( $c_3, e_3$ ); !( $d_3, 0$ ); BR; ?( $c_1, y$ ); ?( $c_2, z$ ); ?( $c_3, w$ ); !( $d'_3, 0$ ); BR')

```

Here, $\mathbf{BR} = ?(d_1, x); ?(d_2, x); ?(d_3, x)$ and $\mathbf{BR}' = ?(d'_1, x); ?(d'_2, x); ?(d'_3, x)$. The program is an example of barrier-style synchronization. Process 1 writes to c_1 , process 2 writes to c_2 , process 3 writes to c_3 , and then the three processes synchronize via a barrier so that none of the processes can proceed until all are done with their writes. Note that $!(d_i, 0); \mathbf{BR}$ models the barrier for each process i . After the barrier synchronization, each process reads from all three reference cells before synchronizing themselves via another barrier, this time modeled by $!(d'_i, 0); \mathbf{BR}'$, before the next iteration of the loop.

The type system can prove confluence by assigning the following types (assume e_1, e_2 , and e_3 are of type int): $c_1 : ch(\rho_{c_1}, int)$, $c_2 : ch(\rho_{c_2}, int)$, $c_3 : ch(\rho_{c_3}, int)$, and for each $i \in \{1, 2, 3\}$,

$$\begin{aligned}
d_i &: ch(\rho_{d_i}, int, \langle \theta[w(\rho_{c_i}) \mapsto \frac{1}{3}], \theta[w(\rho_{c_i}) \mapsto \frac{1}{3}], \theta[w(\rho_{c_i}) \mapsto \frac{1}{3}] \rangle) \\
d'_i &: ch(\rho_{d'_i}, int, \langle \theta[w(\rho_{c_1}) \mapsto \frac{1}{3}], \theta[w(\rho_{c_2}) \mapsto \frac{1}{3}], \theta[w(\rho_{c_3}) \mapsto \frac{1}{3}] \rangle)
\end{aligned}$$

The initial static capability set for each process i is $\theta[w(\rho_{c_i}) \mapsto 1, w(\rho_{d_i}) \mapsto 1, w(\rho_{d'_i}) \mapsto 1]$. Note that fractional capabilities are passed at barrier synchronization points to enable reads and writes on c_1, c_2 , and c_3 .

Type inference fails if the program is changed so that d_1, d_2, d_3 are also used for the second barrier (in place of d'_1, d'_2, d'_3) because while the first write to d_i must send the capability to read c_i , the second write to d_i must send to each process j the capability to access c_j , and there is no single type for d_i to express this behavior. This demonstrates the *flow-insensitivity* limitation of our type system, i.e., a channel must send and receive the same capabilities every time it is used.

However, if synchronization points are syntactically identifiable (as in this example) then the program is easily modified so that flow-insensitivity becomes sufficient by using distinct channels at each syntactic synchronization point.⁵ In our example, the first barrier in each process matches the other, and the second barrier in each process matches the other. Synchronizations that are not syntactically identifiable are often considered as a sign of potential bugs [1]. Note that reference cells c_1 and c_2 are not used for synchronization and therefore need no syntactic restriction.

3.3 Extensions

We discuss extensions to our system.

Regions: Aliasing becomes an issue when channels are used as values, e.g., like in a π calculus program. For example, our type system does not allow two different channels c and d to be passed to the same channel because two different channels cannot be given the same handle. One way to resolve aliasing is to use *regions* so that each ρ represents a set of channels. Then, we may give both c and d the same type $ch(\rho, \dots)$ at the cost of sharing $w(\rho)$ (and $r(\rho)$) for all the channels in the region ρ .

Existential Abstraction and Linear Types: Another way to resolve aliasing is to existentially abstract capabilities as in $\exists \rho. \tau \otimes \Psi$. Any type containing a capability set must be handled linearly⁶ to prevent the duplication of capabilities. The capabilities are recovered by opening the existential package. Existential abstraction can encode linearly typed channels [10, 8] (for rendezvous channels) as: $\exists \rho. ch(\rho, \tau, \theta, \theta) \otimes \theta[w(\rho) \mapsto 1, r(\rho) \mapsto 1]$. Note that the type encapsulates both a channel and the capability to access the channel. This encoding allows transitions to and from linearly typed channels to the capabilities world, e.g., it is possible to use once a linearly-typed channel multiple times. An analogous approach has been applied to express updatable recursive data structures in the capability calculus [13].

Dynamically Created Channels: Dynamically created channels can be handled in much the same way heap allocated objects are handled in the capability

⁵ This can be done without changing the implementation. See *named barriers* in [1].

⁶ Actually, a more relaxed sub-structural type system is preferred for handling fractional capabilities [11].

calculus [4] (we only show the rule for the case where c is rendezvous):

$$\frac{\begin{array}{c} \rho \text{ is not free in the conclusion} \\ \Gamma \cup \{c \mapsto ch(\rho, \tau, \Psi_1, \Psi_2)\}, i, \Psi + \theta[w(\rho) \mapsto 1][r(\rho) \mapsto 1] \vdash s : \Psi' \end{array}}{\Gamma, i, \Psi \vdash \nu c.s : \Psi'}$$

Existential abstraction allows dynamically created channels to leave their lexical scope. An alternative approach is to place the newly created channel in an existing region. In this case, we can remove the hypothesis “ ρ is not free in the conclusion”, but we also must remove the capabilities $\theta[w(\rho) \mapsto 1][r(\rho) \mapsto 1]$.

Dynamically Spawned Processes: Dynamic spawning of processes can be typed as follows:

$$\frac{\Gamma, i, \Psi_2 \vdash s : \Psi'}{\Gamma, i, \Psi_1 + \Psi_2 \vdash \text{spawn}(s) : \Psi_1}$$

(For simplicity, we assume that the local store of the parent process is copied for the spawned process. Details for handling input buffered channels are omitted.) Note that the spawned process may take capabilities from the parent process.

4 Related Work

We discuss previous approaches to inferring partial confluence. Kahn process networks [7] restrict communication to input buffered channels with a unique sender process to guarantee determinism. Edwards et al. [5] restricts communication to rendezvous channels with a unique sender process and a unique receiver process to model deterministic behavior of embedded systems. These models are the easy cases for our system where capabilities are not passed between processes.

Linear type systems can infer partial confluence by checking that each channel is used at most once [10, 8].⁷ Section 3.3 discusses how to express linearly typed channels in our system. König presents a type system that can be parameterized to check partial confluence in the π -calculus [9]. Her system corresponds to the restricted case of our system where each (rendezvous) channel is given a type of the form $ch(\rho, \tau, \theta[w(\rho) \mapsto 1], \theta[r(\rho) \mapsto 1])$, i.e., each channel sends its own write capability at writes and sends its own read capability at reads. Therefore, for example, while her system can check the confluence of $!(c, 1); ?(c, x) || ?(c, x); !(c, 2)$, it cannot check the confluence of $!(c, 1); !(c, 2) || ?(c, x); ?(c, x)$.

A more exhaustive approach for checking partial confluence has been proposed in which the confluence for every state of the program is individually checked by following the transitions from that state [6, 2]. These methods are designed specifically to drive state space reduction, and hence have somewhat a different aim from our work. They have been shown effective for programs with a small number of states.

This work was motivated by our previous work on inferring confluence in functional languages with side effects [11] (see also [3]). These systems can only reason about synchronization at join points, and therefore cannot infer confluence of channel-communicating processes.

⁷ [8] uses asynchronous π calculus, and so is not entirely comparable with our work.

5 Conclusions

We have presented a system for inferring partial confluence of concurrent programs communicating via a mix of different kinds of communication channels. We casted our system as a capability calculus where fractional capabilities can be passed at channel communications, and presented a type system for statically inferring partial confluence by finding an appropriate capability passing strategy in the calculus.

References

1. A. Aiken and D. Gay. Barrier inference. In *Proceedings of the 25th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 342–354, San Diego, California, Jan. 1998.
2. S. Blom and J. van de Pol. State space reduction by proving confluence. In *Proceedings of the 14th International Conference on Computer Aided Verification*, pages 596–609, Copenhagen, Denmark, July 2002.
3. J. Boyland. Checking interference with fractional permissions. In R. Cousot, editor, *Static Analysis, Tenth International Symposium*, volume 2694 of *Lecture Notes in Computer Science*, pages 55–72, San Diego, CA, June 2003. Springer-Verlag.
4. K. Crary, D. Walker, and G. Morrisett. Typed memory management in a calculus of capabilities. In *Proceedings of the 26th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 262–275, San Antonio, Texas, Jan. 1999.
5. S. A. Edwards and O. Tardieu. Shim: a deterministic model for heterogeneous embedded systems. In *Proceedings of the 5th ACM International Conference On Embedded Software*, pages 264–272, Jersey City, NJ, Sept. 2005.
6. J. F. Groote and J. van de Pol. State space reduction using partial tau-confluence. In *Proceedings of 25th International Symposium on the Mathematical Foundations of Computer Science 2000*, pages 383–393, Bratislava, Slovakia, Aug. 2000.
7. G. Kahn. The semantics of a simple language for parallel programming. In *Information processing*, pages 471–475, Stockholm, Sweden, Aug 1974.
8. N. Kobayashi, B. C. Pierce, and D. N. Turner. Linearity and the pi-calculus. *ACM Transactions on Programming Languages and Systems*, 21(5):914–947, Sept. 1999.
9. B. König. Analysing input/output-capabilities of mobile processes with a generic type system. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming*, pages 403–414, Geneva, Switzerland, July 2000.
10. U. Nestmann and M. Steffen. Typing confluence. In *Proceedings of FMICS '97*, pages 77–101, 1997.
11. T. Terauchi and A. Aiken. Witnessing side-effects. In *Proceedings of the 10th ACM SIGPLAN International Conference on Functional Programming*, pages 105–115, Tallinn, Estonia, Sept. 2005. ACM.
12. T. Terauchi and A. Aiken. A capability calculus for concurrency and determinism. Technical Report UCB/EECS-2006-84, University of California, Berkeley, 2006.
13. D. Walker and G. Morrisett. Alias types for recursive data structures. In *International Workshop on Types in Compilation*, Montreal, Canada, Sept. 2000.