

Parallel Complexity Analysis with Temporal Session Types

ANKUSH DAS, Carnegie Mellon University, USA
JAN HOFFMANN, Carnegie Mellon University, USA
FRANK PFENNING, Carnegie Mellon University, USA

We study the problem of parametric parallel complexity analysis of concurrent, message-passing programs. To make the analysis local and compositional, it is based on a conservative extension of binary session types, which structure the type and direction of communication between processes and stand in a Curry-Howard correspondence with intuitionistic linear logic. The main innovation is to enrich session types with the *temporal modalities next* ($\odot A$), *always* ($\square A$), and *eventually* ($\diamond A$), to additionally prescribe the timing of the exchanged messages in a way that is precise yet flexible. The resulting *temporal session types* uniformly express properties such as the message rate of a stream, the latency of a pipeline, the response time of a concurrent queue, or the span of a fork/join parallel program. The analysis is parametric in the cost model and the presentation focuses on communication cost as a concrete example. The soundness of the analysis is established by proofs of progress and type preservation using a timed multiset rewriting semantics. Representative examples illustrate the scope and usability of the approach.

CCS Concepts: • **Theory of computation** → **Concurrency; Modal and temporal logics; Linear logic;**

Additional Key Words and Phrases: Session Types, Linear logic, Concurrency, Resource analysis

ACM Reference Format:

Ankush Das, Jan Hoffmann, and Frank Pfenning. 2018. Parallel Complexity Analysis with Temporal Session Types. *Proc. ACM Program. Lang.* 2, ICFP, Article 91 (September 2018), 30 pages. <https://doi.org/10.1145/3236786>

1 INTRODUCTION

For sequential programs, several type systems and program analyses have been proposed to structure, formalize [Danner et al. 2015; Lago and Gaboardi 2011; Çiçek et al. 2017], and automate [Avanzini et al. 2015; Gulwani et al. 2009; Hoffmann et al. 2017] complexity analysis. Analyzing the complexity of concurrent, message-passing processes poses additional challenges that these systems do not address. To begin with, we need information about the possible interactions between processes to enable compositional and local reasoning about concurrent cost.

Session types [Honda et al. 1998] provide a structured way to prescribe communication behavior between message-passing processes and are a natural foundation for compositional, concurrent complexity analysis. In particular, we use a system of binary session types that stands in a Curry-Howard correspondence with intuitionistic linear logic [Caires and Pfenning 2010; Caires et al. 2016]. Our communication model is *asynchronous* in the sense of the asynchronous π -calculus: sending always succeeds immediately, while receiving blocks until a message arrives.

In addition to the structure of communication, the timing of messages is of central interest for analyzing concurrent cost. With information on message timing we may analyze not only properties such as the rate or latency with which a stream of messages can proceed through a

Authors' addresses: Ankush Das, Carnegie Mellon University, USA, ankushd@cs.cmu.edu; Jan Hoffmann, Carnegie Mellon University, USA, jhoffmann@cmu.edu; Frank Pfenning, Carnegie Mellon University, USA, fp@cs.cmu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2018 Copyright held by the owner/author(s).

2475-1421/2018/9-ART91

<https://doi.org/10.1145/3236786>

pipeline, but also the span of a parallel computation, which can be defined as the time of the final response message assuming maximal parallelism.

There are several possible ways to enrich session types with timing information. A challenge is to find a balance between precision and flexibility. We would like to express precise times according to a global clock as in synchronous data flow languages whenever that is possible. However, sometimes this will be too restrictive. For example, we may want to characterize the response time of a concurrent queue where enqueue and dequeue operations arrive at unpredictable intervals.

In this paper, we develop a type system that captures the parallel complexity of session-typed message-passing programs by adding *temporal modalities next* ($\circ A$), *always* ($\square A$), and *eventually* ($\diamond A$), interpreted over a linear model of time. When considered as types, the temporal modalities allow us to express properties of concurrent programs such as the *message rate* of a stream, the *latency* of a pipeline, the *response time* of concurrent data structure, or the *span* of a fork/join parallel program, all in the same uniform manner. Our results complement prior work on expressing the *work* of session-typed processes in the same base language [Das et al. 2017]. Together, they form a foundation for analyzing the parallel implementation complexity of session-typed processes.

The type system is constructed conservatively over the base language of session types, which makes it quite general and easily able to accommodate various concrete cost models. Our language contains standard session types and process expressions, and their typing rules remain unchanged. They correspond to processes that do not induce cost and send all messages at the constant time 0.

To model computation cost we introduce a new syntactic form **delay**, which advances time by one step. To specify a particular cost semantics we take an ordinary, non-temporal program and add delays capturing the intended cost. For example, if we decide only the blocking operations should cost one unit of time, we add a delay before the continuation of every receiving construct. If we want sends to have unit cost as well, we also add a delay immediately after each send operation. Processes that contain delays cannot be typed using standard session types.

To type processes with non-zero cost, we first introduce the type $\circ A$, which is inhabited only by the process expression (**delay** ; P). This forces time to advance on all channels that P can communicate along. The resulting types prescribe the *exact* time a message is sent or received and sender and receiver are precisely synchronized.

As an example, consider a stream of bits terminated by \$, expressed as the recursive session type

$$\text{bits} = \oplus\{\text{b0} : \text{bits}, \text{b1} : \text{bits}, \$: \mathbf{1}\}$$

where \oplus stands for *internal choice* and $\mathbf{1}$ for *termination*, ending the session. A simple cost model for asynchronous communication prescribes a cost of one unit of time for every receive operation. A stream of bits then needs to delay every continuation to give the recipient time to receive the message, expressing a *rate* of one. This can be captured precisely with the temporal modality $\circ A$:

$$\text{bits} = \oplus\{\text{b0} : \circ\text{bits}, \text{b1} : \circ\text{bits}, \$: \circ\mathbf{1}\}$$

A transducer *neg* that negates each bit it receives along channel x and passes it on along channel y would be typed as

$$x : \text{bits} \vdash \text{neg} :: (y : \circ\text{bits})$$

expressing a *latency* of one. A process *negneg* that puts two negations in sequence has a latency of two, compared with *copy* which passes on each bit, and *id* which terminates and identifies the channel y with the channel x , short-circuiting the communication.

$$x : \text{bits} \vdash \text{negneg} :: (y : \circ\circ\text{bits}) \quad x : \text{bits} \vdash \text{copy} :: (y : \circ\text{bits}) \quad x : \text{bits} \vdash \text{id} :: (y : \text{bits})$$

All these processes have the same extensional behavior, but different latencies. They also have the same rate since after the pipelining delay, the bits are sent at the same rate they are received, as expressed in the common type *bits* used in the context and the result.

While precise and minimalistic, the resulting system is often too precise for typical concurrent programs such as pipelines or servers. We therefore introduce the dual type formers $\diamond A$ and $\square A$ to talk about varying time points in the future. Remarkably, even if part of a program is typed using these constructs, we can still make precise and useful statements about other aspects.

For example, consider a transducer *compress* that shortens a stream by combining consecutive 1 bits so that, for example, 00110111 becomes 00101. For such a transducer, we cannot bound the latency statically, even if the bits are received at a constant rate like in the type bits. So we have to express that after seeing a 1 bit we will *eventually* see either another bit or the end of the stream. For this purpose, we introduce a new type *sbits* with the same message alternatives as bits, but different timing. In particular, after sending b1 we have to send either the next bit or end-of-stream *eventually* (\diamond sbits), rather than immediately.

$$\begin{aligned} \text{sbits} &= \oplus\{\text{b0} : \circ\text{sbits}, \text{b1} : \circ\diamond\text{sbits}, \$: \circ\mathbf{1}\} \\ x : \text{bits} \vdash \text{compress} &:: (y : \circ\text{sbits}) \end{aligned}$$

We write $\circ\diamond\text{sbits}$ instead of $\diamond\text{sbits}$ for the continuation type after b1 to express that there will always be a delay of at least one; to account for the unit cost of receive in this particular cost model.

The dual modality, $\square A$, is useful to express, for example, that a server providing A is *always* ready, starting from “now”. As an example, consider the following temporal type of an interface to a process of type $\square\text{queue}_A$ with elements of type $\square A$. It expresses that there must be at least four time units between successive enqueue operations and that the response to a dequeue request is immediate, only one time unit later ($\&$ stands for external choice, the dual to internal choice).

$$\begin{aligned} \text{queue}_A &= \&\{\text{enq} : \circ(\square A \multimap \circ^3\square\text{queue}_A), \\ &\quad \text{deq} : \circ\oplus\{\text{none} : \circ\mathbf{1}, \text{some} : \circ(\square A \otimes \circ\square\text{queue}_A)\} \} \end{aligned}$$

As an example of a *parametric* cost analysis, we can give the following type to a process that appends inputs l_1 and l_2 to yield l , where the message rate on all three lists is $r + 2$ units of time (that is, the interval between consecutive list elements needs to be at least 2).

$$l_1 : \text{list}_A[n], l_2 : \circ^{(r+4)n+2} \text{list}_A[k] \vdash \text{append} :: (l : \circ\circ\text{list}_A[n+k])$$

It expresses that *append* has a latency of two units of time and that it inputs the first message from l_2 after $(r + 4)n + 2$ units of time, where n is the number of elements sent along l_1 .

To analyze the span of a fork/join parallel program, we capture the time at which the (final) answer is sent. For example, the type $\text{tree}[h]$ describes the span of a process that computes the parity of a binary tree of height h with boolean values at the leaves. The session type expresses that the result of the computation is a single boolean that arrives at time $5h + 3$ after the parity request.

$$\text{tree}[h] = \&\{\text{parity} : \circ^{5h+3} \text{bool}\}$$

In summary, the main contributions of the paper are (1) a generic framework for parallel cost analysis of asynchronously communicating session-typed processes rooted in a novel combination of temporal and linear logic, (2) a soundness proof of the type system with respect to a timed operational semantics, showing progress and type preservation (3) instantiations of the framework with different cost models, e.g. where either just receives, or receives and sends, cost one time unit each, and (4) examples illustrating the scope of our method. Our technique for proving progress and preservation does not require dependency graphs and may be of independent interest. We further provide decidable systems for *time reconstruction* and *subtyping* that greatly simplify the programmer’s task. They also enhance modularity by allowing the same program to be assigned temporally different types, depending on the context of use.

Related is work on space and time complexity analysis of interaction nets by [Gimenez and Moser \[2016\]](#), which is a parallel execution model for functional programs. While also inspired by linear logic and, in particular, proof nets, it treats only special cases of the additive connectives and

Type	Provider Action	Session Continuation
$\oplus\{\ell : A_\ell\}_{\ell \in L}$	send label $k \in L$	A_k
$\&\{\ell : A_\ell\}_{\ell \in L}$	receive and branch on label $k \in L$	A_k
1	send token close	<i>none</i>
$A \otimes B$	send channel $c : A$	B
$A \multimap B$	receive channel $c : A$	B

Fig. 1. Basic Session Types. Every provider action has a matching client action.

recursive types and does not have analogues of the \square and \diamond modalities. It also does not provide a general source-level programming notation with a syntax-directed type system. On the other hand they incorporate sharing and space bounds, which are beyond the scope of this paper.

Another related thread is the research on timed multiparty session types [Bocchi et al. 2014] for modular verification of real-time choreographic interactions. Their system is based on explicit global timing interval constraints, capturing a new class of communicating timed automata, in contrast to our system based on binary session types in a general concurrent language. Therefore, their system has no need for general \square and \diamond modalities, the ability to pass channels along channels, or the ability to identify channels via forwarding. Their work is complemented by an expressive dynamic verification framework in real-time distributed systems [Neykova et al. 2014], which we do not consider. Semantics counting communication costs for work and span in session-typed programs were given by Silva et al. [2016], but no techniques for analyzing them were provided.

The remainder of the paper is organized as follows. We review our basic system of session types in Section 2, then introduce the next-time modality $\circ A$ in Section 3 followed by $\diamond A$ and $\square A$ in Section 4. We establish fundamental metatheoretic type safety properties in Section 5 and time reconstruction in Section 6. Additional examples in Section 7 are followed by a theorem in Section 8 connecting the semantics presented in Figure 4 to the standard semantics of session-typed programs. Section 9 discusses further related work followed by a brief conclusion.

2 THE BASE SYSTEM OF SESSION TYPES

The underlying base system of session types is derived from a Curry-Howard interpretation of intuitionistic linear logic [Caires and Pfenning 2010; Caires et al. 2016]. We present it here to fix our particular formulation, which can be considered the purely linear fragment of SILL [Pfenning and Griffith 2015; Toninho et al. 2013]. Remarkably, the rules remain exactly the same when we consider temporal extensions in the next section. The key idea is that an intuitionistic linear sequent

$$A_1, A_2, \dots, A_n \vdash C$$

is interpreted as the interface to a *process expression* P . We label each of the antecedents with a channel name x_i and the succedent with channel name z . The x_i 's are *channels used by P* and z is the *channel provided by P* .

$$x_1 : A_1, x_2 : A_2, \dots, x_n : A_n \vdash P :: (z : C)$$

The resulting judgment formally states that process P provides a service of session type C along channel z , while using the services of session types A_1, \dots, A_n provided along channels x_1, \dots, x_n respectively. All these channels must be distinct, and we sometimes implicitly rename them to preserve this presupposition. We abbreviate the antecedent of the sequent by Ω .

Figure 1 summarizes the basic session types and their actions. The process expression for these actions are shown in Figure 2; the process typing rules in Figure 3. The first few examples (well into Section 4) only use internal choice, termination, and recursive types, together with process

	Expression	Action	Continuation	Rules
$P, Q ::=$	$x \leftarrow f \leftarrow \bar{e} ; Q$	spawn process named f	$[a/x]Q$	def
	$x:A \leftarrow P ; Q$	spawn $[a/x]P$	$[a/x]Q$	cut
	$c \leftarrow d$	identify c and d	$none$	id
	$c.k ; P$	send label k along c	P	$\oplus R, \& L$
	case $c (\ell \Rightarrow P_\ell)_{\ell \in L}$	receive label k along c	P_k	$\oplus L, \& R$
	close c	close c	$none$	$1R$
	wait $c ; P$	wait for c to close	P	$1L$
	send $c d ; P$	send d along c	P	$\otimes R, \multimap L$
	$x \leftarrow \text{recv } c ; P$	receive d along c	$[d/x]P$	$\otimes L, \multimap R$

Fig. 2. Basic Process Expressions

definitions and forwarding, so we explain these in some detail together with their formal operational semantics. A summary of all the operational semantics rules can be found in Figure 4.

2.1 Internal Choice

A type A is said to describe a *session*, which is a particular sequence of interactions. As a first type construct we consider *internal choice* $\oplus\{\ell : A_\ell\}_{\ell \in L}$, an n -ary labeled generalization of the linear logic connective $A \oplus B$. A process that provides $x : \oplus\{\ell : A_\ell\}_{\ell \in L}$ can send any label $k \in L$ along x and then continue by providing $x : A_k$. We write the corresponding process as $(x.k ; P)$, where P is the continuation. This typing is formalized by the *right rule* $\oplus R$ in our sequent calculus. The corresponding client branches on the label received along x as specified by the *left rule* $\oplus L$.

$$\frac{(k \in L) \quad \Omega \vdash P :: (x : A_k)}{\Omega \vdash (x.k ; P) :: (x : \oplus\{\ell : A_\ell\}_{\ell \in L})} \oplus R \quad \frac{(\forall \ell \in L) \quad \Omega, x:A_\ell \vdash Q_\ell :: (z : C)}{\Omega, x:\oplus\{\ell : A_\ell\}_{\ell \in L} \vdash \text{case } x (\ell \Rightarrow Q_\ell)_{\ell \in L} :: (z : C)} \oplus L$$

We formalize the operational semantics as a system of *multiset rewriting rules* [Cervesato and Scedrov 2009]. We introduce semantic objects $\text{proc}(c, t, P)$ and $\text{msg}(c, t, M)$ which mean that process P or message M provide along channel c and are at an integral time t . A *process configuration* is a multiset of such objects, where any two offered channels are distinct. Communication is asynchronous, so that a process $(c.k ; P)$ sends a message k along c and continues as P without waiting for it to be received. As a technical device to ensure that consecutive messages on a channel arrive in order, the sender also creates a fresh continuation channel c' so that the message k is actually represented as $(c.k ; c \leftarrow c')$ (read: send k along c and continue as c').

$$(\oplus S) \quad \text{proc}(c, t, c.k ; P) \mapsto \text{proc}(c', t, [c'/c]P), \text{msg}(c, t, c.k ; c \leftarrow c') \quad (c' \text{ fresh})$$

When the message k is received along c , we select branch k and also substitute the continuation channel c' for c .

$$(\oplus C) \quad \text{msg}(c, t, c.k ; c \leftarrow c'), \text{proc}(d, t, \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}) \mapsto \text{proc}(d, t, [c'/c]Q_k)$$

The *message* $(c.k ; c \leftarrow c')$ is just a particular form of process, where $c \leftarrow c'$ is *identity* or *forwarding*, explained in Section 2.3. Therefore no separate typing rules for messages are needed; they can be typed as processes [Balzer and Pfenning 2017].

In the receiving rule we require the time t of the message and receiver process to match. Until we introduce temporal types, this is trivially satisfied since all actions are considered instantaneous and processes will always remain at time $t = 0$.

$$\begin{array}{c}
\frac{\Omega' \vdash P :: (x : A) \quad \Omega, x : A \vdash Q :: (z : C)}{\Omega, \Omega' \vdash (x:A \leftarrow P ; Q) :: (z : C)} \text{ cut} \quad \frac{}{y : A \vdash (x \leftarrow y) :: (x : A)} \text{ id} \\
\\
\frac{(k \in L) \quad \Omega \vdash P :: (x : A_k)}{\Omega \vdash (x.k ; P) :: (x : \oplus\{\ell : A_\ell\}_{\ell \in L})} \oplus R \quad \frac{(\forall \ell \in L) \quad \Omega, x:A_\ell \vdash Q_\ell :: (z : C)}{\Omega, x:\oplus\{\ell : A_\ell\}_{\ell \in L} \vdash \text{case } x (\ell \Rightarrow Q_\ell)_{\ell \in L} :: (z : C)} \oplus L \\
\\
\frac{(\forall \ell \in L) \quad \Omega \vdash P_\ell :: (x : A_\ell)}{\Omega \vdash \text{case } x (\ell \Rightarrow P_\ell)_{\ell \in L} :: (x : \&\{\ell : A_\ell\}_{\ell \in L})} \& R \quad \frac{\Omega, x:A_k \vdash Q :: (z : C)}{\Omega, x:\&\{\ell : A_\ell\}_{\ell \in L} \vdash (x.k ; Q) :: (z : C)} \& L \\
\\
\frac{}{\cdot \vdash (\text{close } x) :: (x : 1)} \mathbf{1}R \quad \frac{\Omega \vdash Q :: (z : C)}{\Omega, x:1 \vdash (\text{wait } x ; Q) :: (z : C)} \mathbf{1}L \\
\\
\frac{\Omega \vdash P :: (x : B)}{\Omega, y:A \vdash (\text{send } x y ; P) :: (x : A \otimes B)} \otimes R \quad \frac{\Omega, y:A, x:B \vdash Q :: (z : C)}{\Omega, x:A \otimes B \vdash (y \leftarrow \text{recv } x ; Q) :: (z : C)} \otimes L \\
\\
\frac{\Omega, y:A \vdash P :: (x : B)}{\Omega \vdash (y \leftarrow \text{recv } x ; P) :: (x : A \multimap B)} \multimap R \quad \frac{\Omega, x:B \vdash Q :: (z : C)}{\Omega, x:A \multimap B, y:A \vdash (\text{send } x y ; Q) :: (z : C)} \multimap L \\
\\
\frac{(\Omega' \vdash f = P_f :: (x : A)) \in \Sigma \quad \Omega, x:A \vdash Q :: (z : C)}{\Omega, \Omega' \vdash (x \leftarrow f \leftarrow \Omega' ; Q) :: (z : C)} \text{ def}
\end{array}$$

Fig. 3. Basic Typing Rules

The dual of internal choice is *external choice* $\&\{\ell : A_\ell\}_{\ell \in L}$, which just reverses the role of provider and client and reuses the same process notation. It is the n -ary labeled generalization of the linear logic connective $A \& B$.

2.2 Termination

The type **1**, the multiplicative unit of linear logic, represents termination of a process, which (due to linearity) is not allowed to use any channels.

$$\frac{}{\cdot \vdash \text{close } x :: (x : 1)} \mathbf{1}R \quad \frac{\Omega \vdash Q :: (z : C)}{\Omega, x:1 \vdash (\text{wait } x ; Q) :: (z : C)} \mathbf{1}L$$

Operationally, a client has to wait for the corresponding closing message, which has no continuation since the provider terminates.

$$\begin{array}{ll}
(1S) \quad \text{proc}(c, t, \text{close } c) & \mapsto \text{msg}(c, t, \text{close } c) \\
(1C) \quad \text{msg}(c, t, \text{close } c), \text{proc}(d, t, \text{wait } c ; Q) & \mapsto \text{proc}(d, t, Q)
\end{array}$$

2.3 Forwarding

A process $x \leftarrow y$ *identifies* the channels x and y so that any further communication along either x or y will be along the unified channel. Its typing rule corresponds to the logical rule of *identity*.

$$\frac{}{y : A \vdash (x \leftarrow y) :: (x : A)} \text{ id}$$

We have already seen this form in the continuations of message objects. Operationally, the intuition is realized by *forwarding*: a process $c \leftarrow d$ *forwards* any message M that arrives along d to c and

(cutC)	$\text{proc}(c, t, x:A \leftarrow P ; Q) \mapsto \text{proc}(a, t, [a/x]P), \text{proc}(c, t, [a/x]Q)$	(<i>a</i> fresh)
(defC)	$\text{proc}(c, t, x \leftarrow f \leftarrow \bar{e} ; Q) \mapsto \text{proc}(a, t, [a/x, \bar{e}/\Omega_f]P_f), \text{proc}(c, t, [a/x]Q)$	(<i>a</i> fresh)
(id ⁺ C)	$\text{msg}(d, t, M), \text{proc}(c, s, c \leftarrow d) \mapsto \text{msg}(c, t, [c/d]M)$	($t \geq s$)
(id ⁻ C)	$\text{proc}(c, s, c \leftarrow d), \text{msg}(e, t, M(c)) \mapsto \text{msg}(e, t, [d/c]M(c))$	($s \leq t$)
(\oplus S)	$\text{proc}(c, t, c.k ; P) \mapsto \text{proc}(c', t, [c'/c]P), \text{msg}(c, t, c.k ; c \leftarrow c')$	(<i>c'</i> fresh)
(\oplus C)	$\text{msg}(c, t, c.k ; c \leftarrow c'), \text{proc}(d, t, \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}) \mapsto \text{proc}(d, t, [c'/c]Q_k)$	
($\&$ S)	$\text{proc}(d, t, c.k ; Q) \mapsto \text{msg}(c', t, c.k ; c' \leftarrow c), \text{proc}(d, t, [c'/c]Q)$	(<i>c'</i> fresh)
($\&$ C)	$\text{proc}(c, t, \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}), \text{msg}(c', t, c.k ; c' \leftarrow c) \mapsto \text{proc}(c', t, [c'/c]Q_k)$	
(1S)	$\text{proc}(c, t, \text{close } c) \mapsto \text{msg}(c, t, \text{close } c)$	
(1C)	$\text{msg}(c, t, \text{close } c), \text{proc}(d, t, \text{wait } c ; Q) \mapsto \text{proc}(d, t, Q)$	
(\otimes S)	$\text{proc}(c, t, \text{send } c d ; P) \mapsto \text{proc}(c', t, [c'/c]P), \text{msg}(c, t, \text{send } c d ; c \leftarrow c')$	(<i>c'</i> fresh)
(\otimes C)	$\text{msg}(c, t, \text{send } c d ; c \leftarrow c'), \text{proc}(e, t, x \leftarrow \text{recv } c ; Q) \mapsto \text{proc}(e, t, [c', d/c, x]Q)$	
(\rightarrow S)	$\text{proc}(e, t, \text{send } c d ; Q) \mapsto \text{msg}(c', t, \text{send } c d ; c' \leftarrow c), \text{proc}(e, t, [c'/c]Q)$	(<i>c'</i> fresh)
(\rightarrow C)	$\text{proc}(c, t, x \leftarrow \text{recv } x ; P), \text{msg}(c', t, \text{send } c d ; c' \leftarrow c) \mapsto \text{proc}(c', t, [c', d/c, x]P)$	

Fig. 4. Basic Operational Semantics

vice versa. Because channels are used linearly the forwarding process can then terminate, making sure to apply the proper renaming. The corresponding rules of operational semantics are as follows.

(id ⁺ C)	$\text{msg}(d, t, M), \text{proc}(c, s, c \leftarrow d) \mapsto \text{msg}(c, t, [c/d]M)$	($t \geq s$)
(id ⁻ C)	$\text{proc}(c, s, c \leftarrow d), \text{msg}(e, t, M(c)) \mapsto \text{msg}(e, t, [d/c]M(c))$	($s \leq t$)

In the last transition, we write $M(c)$ to indicate that c must occur in M , which implies that this message is the sole client of c . In anticipation of the extension by temporal operators, we do not require the time of the message and the forwarding process to be identical, but just that the forwarding process is ready *before* the message arrives.

2.4 Process Definitions

Process definitions have the form $\Omega \vdash f = P :: (x : A)$ where f is the name of the process and P its definition. All definitions are collected in a fixed global signature Σ . We require that $\Omega \vdash P :: (x : A)$ for every definition, which allows the definitions to be mutually recursive. For readability of the examples, we break a definition into two declarations, one providing the type and the other the process definition binding the variables x and those in Ω (generally omitting their types):

$$\begin{aligned} \Omega \vdash f &:: (x : A) \\ x \leftarrow f \leftarrow \Omega &= P \end{aligned}$$

A new instance of a defined process f can be spawned with the expression

$$x \leftarrow f \leftarrow \bar{y} ; Q$$

where \bar{y} is a sequence of variables matching the antecedents Ω . The newly spawned process will use all variables in \bar{y} and provide x to the continuation Q . The operational semantics is defined by

$$(\text{defC}) \quad \text{proc}(c, t, x \leftarrow f \leftarrow \bar{e} ; Q) \mapsto \text{proc}(a, t, [a/x, \bar{e}/\Omega]P), \text{proc}(c, t, [a/x]Q) \quad (a \text{ fresh})$$

Here we write \bar{e}/Ω to denote substitution of the channels in \bar{e} for the corresponding variables in Ω .

Sometimes a process invocation is a *tail call*, written without a continuation as $x \leftarrow f \leftarrow \bar{y}$. This is a short-hand for $x' \leftarrow f \leftarrow \bar{y} ; x \leftarrow x'$ for a fresh variable x' , that is, we create a fresh channel and immediately identify it with x (although it is generally implemented more efficiently).

2.5 Recursive Types

Session types can be naturally extended to include recursive types. For this purpose we allow (possibly mutually recursive) type definitions $X = A$ in the signature, where we require A to be *contractive* [Gay and Hole 2005]. This means here that A should not itself be a type name. Our type definitions are *equi-recursive* so we can silently replace X by A during type checking, and no explicit rules for recursive types are needed.

As a first example, consider a stream of bits (introduced in Section 1) defined recursively as

$$\text{bits} = \oplus \{b0 : \text{bits}, b1 : \text{bits}, \$: 1\}$$

When considering bits as representing natural numbers, we think of the least significant bit being sent first. For example, a process *six* sending the number $6 = (110)_2$ would be

$$\begin{aligned} & \cdot \vdash \text{six} :: (x : \text{bits}) \\ & x \leftarrow \text{six} = x.b0 ; x.b1 ; x.b1 ; x.\$; \text{close } x \end{aligned}$$

Executing $\text{proc}(c_0, 0, c_0 \leftarrow \text{six})$ yields (with some fresh channels c_1, \dots, c_4)

$$\begin{aligned} \text{proc}(c_0, 0, c_0 \leftarrow \text{six}) \mapsto^* & \text{msg}(c_4, 0, \text{close } c_4), \\ & \text{msg}(c_3, 0, c_3.\$; c_3 \leftarrow c_4), \\ & \text{msg}(c_2, 0, c_2.b1 ; c_2 \leftarrow c_3), \\ & \text{msg}(c_1, 0, c_1.b1 ; c_1 \leftarrow c_2), \\ & \text{msg}(c_0, 0, c_0.b0 ; c_0 \leftarrow c_1) \end{aligned}$$

As a first example of a recursive process definition, consider one that just copies the incoming bits.

$$\begin{aligned} & y : \text{bits} \vdash \text{copy} :: (x : \text{bits}) \\ & x \leftarrow \text{copy} \leftarrow y = \\ & \quad \text{case } y \text{ (} b0 \Rightarrow x.b0 ; x \leftarrow \text{copy} \leftarrow y \quad \% \text{ received } b0 \text{ on } y, \text{ send } b0 \text{ on } x, \text{ recurse} \\ & \quad \quad | b1 \Rightarrow x.b1 ; x \leftarrow \text{copy} \leftarrow y \quad \% \text{ received } b1 \text{ on } y, \text{ send } b1 \text{ on } x, \text{ recurse} \\ & \quad \quad | \$ \Rightarrow x.\$; \text{wait } y ; \text{close } x \text{) } \% \text{ received } \$ \text{ on } y, \text{ send } \$ \text{ on } x, \text{ wait on } y, \text{ close } x \end{aligned}$$

The process *neg* mentioned in the introduction would just swap the occurrences of $x.b0$ and $x.b1$. We see here an occurrence of a (recursive) *tail call* to *copy*.

A last example in this section: to increment a bit stream we turn $b0$ to $b1$ but then forward the remaining bits unchanged ($x \leftarrow y$), or we turn $b1$ to $b0$ but then increment the remaining stream ($x \leftarrow \text{plus1} \leftarrow y$) to capture the effect of the carry bit.

$$\begin{aligned} & y : \text{bits} \vdash \text{plus1} :: (x : \text{bits}) \\ & x \leftarrow \text{plus1} \leftarrow y = \\ & \quad \text{case } y \text{ (} b0 \Rightarrow x.b1 ; x \leftarrow y \\ & \quad \quad | b1 \Rightarrow x.b0 ; x \leftarrow \text{plus1} \leftarrow y \\ & \quad \quad | \$ \Rightarrow x.\$; \text{wait } y ; \text{close } x \text{) } \end{aligned}$$

3 THE TEMPORAL MODALITY NEXT ($\circ A$)

In this section we introduce *actual cost* by explicitly advancing time. Remarkably, all the rules we have presented so far remain literally unchanged. As mentioned, they correspond to the cost-free fragment of the language in which time never advances. In addition, we have a new type construct $\circ A$ (read: *next A*) with a corresponding process construct (**delay** ; P), which advances time by one unit. In the corresponding typing rule

$$\frac{\Omega \vdash P :: (x : A)}{\circ \Omega \vdash (\text{delay} ; P) :: (x : \circ A)} \circ LR$$

we abbreviate $y_1:\circ B_1, \dots, y_m:\circ B_m$ by $\circ(y_1:B_1, \dots, y_m:B_m)$. Intuitively, when $(\text{delay}; P)$ idles, time advances on *all* channels connected to P . Computationally, we delay the process for one time unit without any external interactions.

$$(\circ C) \quad \text{proc}(c, t, \text{delay}; P) \mapsto \text{proc}(c, t + 1, P)$$

There is a subtle point about forwarding: A process $\text{proc}(c, t, c \leftarrow d)$ may be ready to forward a message *before* a client reaches time t while in all other rules the times must match exactly. We can avoid this mismatch by transforming uses of forwarding $x \leftarrow y$ at type $\circ^n S$ where $S \neq \circ(-)$ to $(\text{delay}^n; x \leftarrow y)$. In this discussion we have used the following notation which will be useful later:

$$\begin{array}{ll} \circ^0 A & = A & \text{delay}^0; P & = P \\ \circ^{n+1} A & = \circ \circ^n A & \text{delay}^{n+1}; P & = \text{delay}; \text{delay}^n; P \end{array}$$

3.1 Modeling a Cost Semantics

Our system allows us to represent a variety of different abstract cost models in a straightforward way. We will mostly use two different abstract cost models. In the first, called \mathcal{R} , we assign unit cost to every receive (or wait) action while all other operations remain cost-free. We may be interested in this since receiving a message is the only blocking operation in the asynchronous semantics. A second one, called \mathcal{RS} and considered in Section 7, assigns unit cost to both send and receive actions.

To capture \mathcal{R} we take a source program and insert a delay operation before the continuation of every receive. We write this delay as **tick** in order to remind the reader that it arises systematically from the cost model and is never written by the programmer. In all other respects, **tick** is just a synonym for **delay**.

For example, the earlier copy process would become

```
bits = ⊕{b0 : bits, b1 : bits, $ : 1}
y : bits ⊢ copy :: (x : bits)           % No longer correct!
x ← copy ← y =
  case y (b0 ⇒ tick; x.b0; x ← copy ← y
         | b1 ⇒ tick; x.b1; x ← copy ← y
         | $ ⇒ tick; x.$; wait y; tick; close x)
```

As indicated in the comment, the type of *copy* is now no longer correct because the bits that arrive along y are delayed by one unit before they are sent along x . We can observe this concretely by starting to type-check the first branch

```
y : bits ⊢ copy :: (x : bits)
x ← copy ← y =
  case y (b0 ⇒                                     % y : bits ⊢ x : bits
         tick; ...)
```

We see that the delay **tick** does not type-check, because neither x nor y have a type of the form $\circ(-)$. We need to redefine the type *bits* so that the continuation type after every label is delayed by one, anticipating the time it takes to receive the label b_0 , b_1 , or $\$$. Similarly, we capture in the type of *copy* that its *latency* is one unit of time.

```
bits = ⊕{b0 : ◦bits, b1 : ◦bits, $ : ◦1}
y : bits ⊢ copy :: (x : ◦bits)
```

With these declarations, we can now type-check the definition of *copy*. We show the intermediate type of the used and provided channels after each interaction.

```

x ← copy ← y =
  case y ( b0 ⇒
            tick ;
            x.b0 ;
            x ← copy ← y
          | b1 ⇒
            tick ;
            x.b1 ;
            x ← copy ← y
          | $ ⇒
            tick ;
            x.$ ;
            wait y ;
            tick ;
            close x )

```

$\% y : \circ\text{bits} \vdash x : \circ\text{bits}$
 $\% y : \text{bits} \vdash x : \text{bits}$
 $\% y : \text{bits} \vdash x : \circ\text{bits}$
 $\% \text{ well-typed by type of copy}$
 $\% y : \circ\text{bits} \vdash x : \circ\text{bits}$
 $\% y : \text{bits} \vdash x : \text{bits}$
 $\% y : \text{bits} \vdash x : \circ\text{bits}$
 $\% y : \circ 1 \vdash x : \circ\text{bits}$
 $\% y : 1 \vdash x : \text{bits}$
 $\% y : 1 \vdash x : \circ 1$
 $\% \cdot \vdash x : \circ 1$
 $\% \cdot \vdash x : 1$

Armed with this experience, we now consider the increment process *plus1*. Again, we expect the latency of the increment to be one unit of time. Since we are interested in detailed type-checking, we show the transformed program, with a delay *tick* after each receive.

```

bits = ⊕{b0 : ◯bits, b1 : ◯bits, $ : ◯1}
y : bits ⊢ plus1 :: (x : ◯bits)
x ← plus1 ← y =
  case y ( b0 ⇒ tick ; x.b1 ; x ← y
          | b1 ⇒ tick ; x.b0 ; x ← plus1 ← y
          | $ ⇒ tick ; x.$ ; wait y ; tick ; close x )

```

$\% \text{ type error here!}$

The branches for *b1* and *\$* type-check as before, but the branch for *b0* does not. We make the types at the crucial point explicit:

```

x ← plus1 ← y =
  case y ( b0 ⇒ tick ; x.b1 ;
          x ← y
          | ... )

```

$\% y : \text{bits} \vdash x : \circ\text{bits}$
 $\% \text{ ill-typed, since bits} \neq \circ\text{bits}$

The problem here is that identifying *x* and *y* removes the delay mandated by the type of *plus1*. A solution is to call *copy* to reintroduce the latency of one time unit.

```

y : bits ⊢ plus1 :: (x : ◯bits)
x ← plus1 ← y =
  case y ( b0 ⇒ tick ; x.b1 ; x ← copy ← y
          | b1 ⇒ tick ; x.b0 ; x ← plus1 ← y
          | $ ⇒ tick ; x.$ ; wait y ; tick ; close x )

```

In order to write *plus2* as a pipeline of two increments we need to delay the second increment explicitly in the program and stipulate, in the type, that there is a latency of two.

```

y : bits ⊢ plus2 :: (x : ◯◯bits)
x ← plus2 ← y =
  z ← plus1 ← y ;
  delay ;
  x ← plus1 ← z

```

$\% z : \circ\text{bits} \vdash x : \circ\circ\text{bits}$
 $\% z : \text{bits} \vdash x : \circ\text{bits}$

Programming with so many explicit delays is tedious, but fortunately we can transform a source program without all these delay operations (but explicitly temporal session types) automatically

in two steps: (1) we insert the delays mandated by the cost model (here: a **tick** after each receive), and (2) we perform *time reconstruction* to insert the additional delays so the result is temporally well-typed or issue an error message if this is impossible (see Section 6).

3.2 The Interpretation of a Configuration

We reconsider the program to produce the number $6 = (110)_2$ under the cost model from the previous section where each receive action costs one unit of time. There are no receive operations in this program, but time reconstruction must insert a delay after each send in order to match the delays mandated by the type bits.

$$\text{bits} = \oplus\{\text{b0} : \circ\text{bits}, \text{b1} : \circ\text{bits}, \$: \circ\mathbf{1}\}$$

$$\cdot \vdash \text{six} :: (x : \text{bits})$$

$$x \leftarrow \text{six} = x.\text{b0} ; \text{delay} ; x.\text{b1} ; \text{delay} ; x.\text{b1} ; \text{delay} ; x.\$; \text{delay} ; \text{close } x$$

Executing $\text{proc}(c_0, 0, c_0 \leftarrow \text{six})$ then leads to the following configuration

$$\begin{aligned} & \text{msg}(c_4, 4, \text{close } c_4), \\ & \text{msg}(c_3, 3, c_3.\$; c_3 \leftarrow c_4), \\ & \text{msg}(c_2, 2, c_2.\text{b1} ; c_2 \leftarrow c_3), \\ & \text{msg}(c_1, 1, c_1.\text{b1} ; c_1 \leftarrow c_2), \\ & \text{msg}(c_0, 0, c_0.\text{b0} ; c_0 \leftarrow c_1) \end{aligned}$$

These messages are at increasing times, which means any client of c_0 will have to immediately (at time 0) receive b0 , then (at time 1) b1 , then (at time 2) b1 , etc. In other words, the time stamps on messages predict *exactly* when the message will be received. Of course, if there is a client in parallel we may never reach this state because, for example, the first b0 message along channel c_0 may be received before the continuation of the sender produces the message b1 . So different configurations may be reached depending on the *scheduler* for the concurrent processes. It is also possible to give a time-synchronous semantics in which all processes proceed *in parallel* from time 0 to time 1, then from time 1 to time 2, etc.

4 THE TEMPORAL MODALITIES ALWAYS ($\square A$) AND EVENTUALLY ($\diamond A$)

The strength and also the weakness of the system so far is that its timing is very precise. Now consider a process *compress* that combines runs of consecutive 1's to a single 1. For example, compressing 11011100 should yield 10100. First, in the cost-free setting we might write

$$\text{bits} = \oplus\{\text{b0} : \text{bits}, \text{b1} : \text{bits}, \$: \mathbf{1}\}$$

$$y : \text{bits} \vdash \text{compress} :: (x : \text{bits})$$

$$y : \text{bits} \vdash \text{skip1s} :: (x : \text{bits})$$

$$x \leftarrow \text{compress} \leftarrow y =$$

$$\begin{aligned} & \text{case } y \text{ (b0} \Rightarrow x.\text{b0} ; x \leftarrow \text{compress} \leftarrow y \\ & \quad | \text{b1} \Rightarrow x.\text{b1} ; x \leftarrow \text{skip1s} \leftarrow y \\ & \quad | \$ \Rightarrow x.\$; \text{wait } y ; \text{close } x) \end{aligned}$$

$$x \leftarrow \text{skip1s} \leftarrow y =$$

$$\begin{aligned} & \text{case } y \text{ (b0} \Rightarrow x.\text{b0} ; x \leftarrow \text{compress} \leftarrow y \\ & \quad | \text{b1} \Rightarrow x \leftarrow \text{skip1s} \leftarrow y \\ & \quad | \$ \Rightarrow x.\$; \text{wait } y ; \text{close } x) \end{aligned}$$

The problem is that if we adopt the cost model \mathcal{R} where every receive takes one unit of time, then this program cannot be typed. Actually worse: there is no way to insert next-time modalities into the type and additional delays into the program so that the result is well-typed. This is because if

the input stream is unknown we cannot predict how long a run of 1's will be, but the length of such a run will determine the delay between sending a bit 1 and the following bit 0.

The best we can say is that after a bit 1 we will *eventually* send either a bit 0 or the end-of-stream token \$. This is the purpose of the type $\diamond A$. We capture this timing in the type sbits (for *slow bits*).

$$\begin{aligned} \text{bits} &= \oplus\{\text{b0} : \circ\text{bits}, \text{b1} : \circ\text{bits}, \$: \circ\mathbf{1}\} \\ \text{sbits} &= \oplus\{\text{b0} : \circ\text{sbits}, \text{b1} : \circ\diamond\text{sbits}, \$: \circ\mathbf{1}\} \\ y : \text{bits} &\vdash \text{compress} :: (x : \circ\text{sbits}) \\ y : \text{bits} &\vdash \text{skip1s} :: (x : \circ\diamond\text{sbits}) \end{aligned}$$

In the next section we introduce the process constructs and typing rules so we can revise our *compress* and *skip1s* programs so they have the right temporal semantics.

4.1 Eventually A

A process providing $\diamond A$ promises only that it will eventually provide A . There is a somewhat subtle point here: since not every action may require time and because we do not check termination separately, $x : \diamond A$ expresses only that *if the process providing x terminates* it will eventually provide A . Thus, it expresses non-determinism regarding the (abstract) *time* at which A is provided, rather than a strict liveness property. Therefore, $\diamond A$ is somewhat weaker than one might be used to from LTL [Pnueli 1977]. When restricted to a purely logical fragment, without unrestricted recursion, the usual meaning is fully restored so we feel our terminology is justified. Imposing termination, for example along the lines of Fortier and Santocanale [2013] or Toninho et al. [2014] is an interesting item for future work but not necessary for our present purposes.

When a process offering $c : \diamond A$ is ready, it will send a **now!** message along c and then continue at type A . Conversely, the client of $c : \diamond A$ will have to be ready and waiting for the **now!** message to arrive along c and then continue at type A . We use (**when?** $c ; Q$) for the corresponding client. These explicit constructs are a conceptual device and may not need to be part of an implementation. They also make type-checking processes entirely syntax-directed and trivially decidable.

The typing rules for **now!** and **when?** are somewhat subtle.

$$\frac{\Omega \vdash P :: (x : A)}{\Omega \vdash (\text{now! } x ; P) :: (x : \diamond A)} \diamond R \qquad \frac{\circ^*\square\Omega' = \Omega \quad \Omega, x:A \vdash Q :: (z : C) \quad C = \circ^*\diamond C'}{\Omega, x:\diamond A \vdash (\text{when? } x ; Q) :: (z : C)} \diamond L$$

The $\diamond R$ rule just states that, without constraints, we can at any time decide to communicate along $x : \diamond A$ and then continue the session at type A . The $\diamond L$ rule expresses that the process must be ready to receive a **now!** message along $x : \diamond A$, but there are two further constraints. Because the process (**when?** $x ; Q$) may need to wait an indefinite period of time, the rule must make sure that communication along z and any channel in Ω can also be postponed an indefinite period of time. We write $C = \circ^*\diamond C'$ to require that C may be delayed a fixed finite number of time steps and then must be allowed to communicate at an arbitrary time in the future. Similarly, for every channel $y : B$ in Ω , B must have the form $\circ^*\square B$, where \square (as the dual of \diamond) is introduced in Section 4.3.

In the operational semantics, the central restriction is that **when?** is ready *before* the **now!** message arrives so that the continuation can proceed immediately as promised by the type.

$$\begin{aligned} (\diamond S) \quad \text{proc}(c, t, \text{now! } c ; P) &\mapsto \text{proc}(c', t, [c'/c]P), \text{msg}(c, t, \text{now! } c ; c \leftarrow c') \quad (c' \text{ fresh}) \\ (\diamond C) \quad \text{msg}(c, t, \text{now! } c ; c \leftarrow c'), \text{proc}(d, s, \text{when? } c ; Q) &\mapsto \text{proc}(d, t, [c'/c]Q) \quad (t \geq s) \end{aligned}$$

We are now almost ready to rewrite the *compress* process in our cost model \mathcal{R} . First, we insert **tick** before all the actions that must be delayed according to our cost model. Then we insert appropriate additional **delay**, **when?**, and **now!** actions. While *compress* turns out to be straightforward, *skip1s* creates a difficulty after it receives a **b1**:

```

bits =  $\oplus\{b0 : \circ\text{bits}, b1 : \circ\text{bits}, \$ : \circ\mathbf{1}\}$ 
sbits =  $\oplus\{b0 : \circ\text{sbits}, b1 : \circ\Diamond\text{sbits}, \$ : \circ\mathbf{1}\}$ 

y : bits  $\vdash$  compress :: (x :  $\circ\text{sbits}$ )
y : bits  $\vdash$  skip1s :: (x :  $\circ\Diamond\text{sbits}$ )

x  $\leftarrow$  compress  $\leftarrow$  y =
  case y ( b0  $\Rightarrow$  tick ; x.b0 ; x  $\leftarrow$  compress  $\leftarrow$  y
          | b1  $\Rightarrow$  tick ; x.b1 ; x  $\leftarrow$  skip1s  $\leftarrow$  y
          | $  $\Rightarrow$  tick ; x.$ ; wait y ; tick ; close x )

x  $\leftarrow$  skip1s  $\leftarrow$  y =
  case y ( b0  $\Rightarrow$  tick ; now! x ; x.b0 ; x  $\leftarrow$  compress  $\leftarrow$  y
          | b1  $\Rightarrow$  tick ;
            x'  $\leftarrow$  skip1s  $\leftarrow$  y ;           % y : bits  $\vdash$  x :  $\Diamond\text{sbits}$ 
            x  $\leftarrow$  idle  $\leftarrow$  x'           % x' :  $\circ\Diamond\text{sbits}$   $\vdash$  x :  $\Diamond\text{sbits}$ 
          | $  $\Rightarrow$  tick ; now! x ; x.$ ; wait y ; tick ; close x )
            % with x' :  $\circ\Diamond\text{sbits}$   $\vdash$  idle :: (x :  $\Diamond\text{sbits}$ )

```

At the point where we would like to call *skip1s* recursively, we have

```

y : bits  $\vdash$  x :  $\Diamond\text{sbits}$ 
but y : bits  $\vdash$  skip1s :: (x :  $\circ\Diamond\text{sbits}$ )

```

which prevents a tail call since $\circ\Diamond\text{sbits} \neq \Diamond\text{sbits}$. Instead we call *skip1s* to obtain a new channel x' and then use another process called *idle* to go from $x' : \circ\Diamond\text{sbits}$ to $x : \Diamond\text{sbits}$. Intuitively, it should be possible to implement such an idling process: $x : \Diamond\text{sbits}$ expresses *at some time in the future, including possibly right now* while $x' : \circ\Diamond\text{sbits}$ says *at some time in the future, but not right now*.

To type the idling process, we need to generalize the $\circ\text{LR}$ rule to account for the interactions of $\circ A$ with $\square A$ and $\Diamond A$. After all, they speak about the same underlying model of time.

4.2 Interactions of $\circ A$ and $\Diamond A$

Recall the left/right rule for \circ :

$$\frac{\Omega \vdash P :: (x : A)}{\circ\Omega \vdash (\text{delay} ; P) :: (x : \circ A)} \circ\text{LR}$$

If the succedent were $x : \Diamond A$ instead of $x : \circ A$, we should still be able to delay since we can freely choose when to interact along x . We could capture this in the following rule (superseded later by a more general form of $\circ\text{LR}$):

$$\frac{\Omega \vdash P :: (x : \Diamond A)}{\circ\Omega \vdash (\text{delay} ; P) :: (x : \Diamond A)} \circ\Diamond$$

We keep $\Diamond A$ as the type of x since we retain the full flexibility of using x at any time in the future after the initial delay. We will generalize the rule once more in the next section to account for interactions with $\square A$.

With this, we can define and type the idling process parametrically over A :

```

x' :  $\circ\Diamond A$   $\vdash$  idle :: (x :  $\Diamond A$ )
x  $\leftarrow$  idle  $\leftarrow$  x' = delay ; x  $\leftarrow$  x'

```

This turns out to be an example of subtyping (see Section 6.1), which means that the programmer actually will not have to explicitly define or even reference an idling process. The programmer simply writes the original *skip1s* process (without referencing the *idle* process) and our subtyping algorithm will use the appropriate rule to typecheck it successfully.

$$\begin{array}{l}
d : \circ\text{ctr} \vdash \text{bit0} :: (c : \text{ctr}) \\
d : \text{ctr} \vdash \text{bit1} :: (c : \text{ctr}) \\
\cdot \vdash \text{empty} :: (c : \text{ctr}) \\
c \leftarrow \text{bit0} \leftarrow d = \\
\quad \text{when? } c ; \quad \quad \quad \% d : \circ\text{ctr} \vdash c : \&\{\dots\} \\
\quad \text{case } c \text{ (inc} \Rightarrow \text{tick ;} \quad \quad \quad \% d : \text{ctr} \vdash c : \text{ctr} \\
\quad \quad \quad c \leftarrow \text{bit1} \leftarrow d \\
\quad \quad | \text{val} \Rightarrow \text{tick ;} \quad \quad \quad \% d : \text{ctr} \vdash c : \text{bits} \\
\quad \quad \quad c.\text{b0 ;} \quad \quad \quad \% d : \text{ctr} \vdash c : \circ\text{bits} \\
\quad \quad \quad \text{now! } d ; d.\text{val ;} \quad \quad \quad \% d : \circ\text{bits} \vdash c : \circ\text{bits} \\
\quad \quad \quad c \leftarrow d) \\
c \leftarrow \text{bit1} \leftarrow d = \\
\quad \text{when? } c ; \quad \quad \quad \% d : \text{ctr} \vdash c : \&\{\dots\} \\
\quad \text{case } c \text{ (inc} \Rightarrow \text{tick ;} \quad \quad \quad \% d : \text{ctr} \vdash c : \text{ctr} \quad (\text{see Section 4.4}) \\
\quad \quad \quad \text{now! } d ; d.\text{inc ;} \quad \quad \quad \% d : \circ\text{ctr} \vdash c : \text{ctr} \\
\quad \quad \quad c \leftarrow \text{bit0} \leftarrow d \\
\quad \quad | \text{val} \Rightarrow \text{tick ;} \quad \quad \quad \% d : \text{ctr} \vdash c : \text{bit} \quad (\text{see Section 4.4}) \\
\quad \quad \quad c.\text{b1 ;} \quad \quad \quad \% d : \text{ctr} \vdash c : \circ\text{bits} \\
\quad \quad \quad \text{now! } d ; d.\text{val ;} \quad \quad \quad \% d : \circ\text{bits} \vdash c : \circ\text{bits} \\
\quad \quad \quad c \leftarrow d) \\
c \leftarrow \text{empty} = \\
\quad \text{when? } c ; \quad \quad \quad \% \cdot \vdash c : \&\{\dots\} \\
\quad \text{case } c \text{ (inc} \Rightarrow \text{tick ;} \quad \quad \quad \% \cdot \vdash c : \text{ctr} \\
\quad \quad \quad e \leftarrow \text{empty ;} \quad \quad \quad \% e : \text{ctr} \vdash c : \text{ctr} \\
\quad \quad \quad c \leftarrow \text{bit1} \leftarrow e \\
\quad \quad | \text{val} \Rightarrow \text{tick ; } c.\$; \quad \quad \quad \% \cdot \vdash c : \circ 1 \\
\quad \quad \quad \text{delay ; close } c)
\end{array}$$

4.4 Interactions Between Temporal Modalities

Just as $\circ A$ and $\diamond A$ interacted in the rules since their semantics is based on the same underlying notion of time, so do $\circ A$ and $\square A$. If we execute a delay, we can allow any channel of type $\square A$ that we use and leave its type unchanged because we are not obligated to communicate along it at any particular time. It is a little awkward to formulate this because among the channels used there may be some of type $\circ B$ and some of type $\square B$.

$$\frac{\square\Omega, \Omega' \vdash P :: (x : A)}{\square\Omega, \circ\Omega' \vdash (\text{delay} ; P) :: (x : \circ A)} \circ$$

In the example of *bit1* at the end of the previous section, we have already seen two lines where this generalization was crucial, observing that $\text{ctr} = \square\&\{\dots\}$.

But even this rule does not cover all possibilities, because the channel x could be of type $\diamond A$. We introduce a new notation, writing $[A]_L^{-1}$ and $[A]_R^{-1}$ on types and extend it to contexts. Depending on one's point of view, this can be seen as stepping forward or backward by one unit of time.

$$\begin{array}{lll}
[\circ A]_L^{-1} = A & [\circ A]_R^{-1} = A & [x : A]_L^{-1} = x : [A]_L^{-1} \\
[\square A]_L^{-1} = \square A & [\square A]_R^{-1} = \text{undefined} & [x : A]_R^{-1} = x : [A]_R^{-1} \\
[\diamond A]_L^{-1} = \text{undefined} & [\diamond A]_R^{-1} = \diamond A & [\cdot]_L^{-1} = \cdot \\
[S]_L^{-1} = \text{undefined} & [S]_R^{-1} = \text{undefined} & [\Omega, \Omega']_L^{-1} = [\Omega]_L^{-1}, [\Omega']_L^{-1}
\end{array}$$

$$\begin{array}{c}
\frac{[\Omega]_L^{-1} \vdash P :: [x : A]_R^{-1}}{\Omega \vdash (\text{delay } P) :: (x : A)} \text{ } \circ LR \quad \frac{}{\circ^* \Box A \text{ delayed}^\square} \quad \frac{}{\circ^* \Diamond A \text{ delayed}^\diamond} \\
\frac{\Omega \vdash P :: (x : A)}{\Omega \vdash (\text{now! } x ; P) :: (x : \Diamond A)} \text{ } \diamond R \quad \frac{\Omega \text{ delayed}^\square \quad \Omega, x:A \vdash Q :: (z : C) \quad C \text{ delayed}^\diamond}{\Omega, x:\Diamond A \vdash (\text{when? } x ; Q) :: (z : C)} \text{ } \diamond L \\
\frac{\Omega \text{ delayed}^\square \quad \Omega \vdash P :: (x : A)}{\Omega \vdash (\text{when? } x ; P) :: (x : \Box A)} \text{ } \Box R \quad \frac{\Omega, x:A \vdash Q :: (z : C)}{\Omega, x:\Box A \vdash (\text{now! } x ; Q) :: (z : C)} \text{ } \Box L
\end{array}$$

Fig. 5. Explicit Temporal Typing Rules

Here, S stands for any basic session type constructor as in Figure 1. We use this notation in the general rule $\circ LR$ which can be found in Figure 5 together with the final set of rules for $\Box A$ and $\Diamond A$. In conjunction with the rules in Figure 3 this completes the system of temporal session types where all temporal actions are explicit. The rule $\circ LR$ only applies if both $[\Omega]_L^{-1}$ and $[x : A]_R^{-1}$ are defined.

We call a type A *patient* if it does not force communication along a channel $x : A$ at any particular point in time. Because the direction of communication is reversed between the two sides of a sequent, a type A is patient if it has the form $\circ^* \Box A'$ if it is among the antecedents, and $\circ^* \Diamond A'$ if it is in the succedent. We write $A \text{ delayed}^\square$ and $A \text{ delayed}^\diamond$ and extend it to contexts $\Omega \text{ delayed}^\square$ if for every declaration $(x : A) \in \Omega$, we have $A \text{ delayed}^\square$.

5 PRESERVATION AND PROGRESS

The main theorems that exhibit the deep connection between our type system and the timed operational semantics are the usual *type preservation* and *progress*, sometimes called *session fidelity* and *deadlock freedom*, respectively. Compared to other recent treatments of linear session types [Balzer and Pfenning 2017; Pfenning and Griffith 2015], new challenges are presented by abstract time and the temporal modalities.

5.1 Configuration Typing

A key question is how we type configurations C . Configurations consist of multiple processes and messages, so they both *use* and *provide* a collection of channels. And even though we treat a configuration as a multiset, typing imposes a partial order on the processes and messages where a provider of a channel appears to the left of its client.

$$\text{Configuration } C ::= \cdot \mid C \ C' \mid \text{proc}(c, t, P) \mid \text{msg}(c, t, M)$$

We say $\text{proc}(c, t, P)$ and $\text{msg}(c, t, M)$ *provide* c . We stipulate that no two distinct processes or messages in a configuration provide the same channel c . Also recall that messages M are simply processes of a particular form and are typed as such. We can read off the possible messages (of which there is one for each type constructor) from the operational semantics. They are summarized here for completeness.

$$M ::= (c.k ; c \leftarrow c') \mid (c.k ; c' \leftarrow c) \mid \text{close } c \mid (\text{send } c \ d ; c' \leftarrow c) \mid (\text{send } c \ d ; c \leftarrow c')$$

The typing judgment has the form $\Omega' \vDash C :: \Omega$ meaning that if composed with a configuration that provides Ω' , the result will provide Ω .

$$\frac{}{\Omega \vDash (\cdot) :: \Omega} \text{empty} \quad \frac{\Omega_0 \vDash C_1 :: \Omega_1 \quad \Omega_1 \vDash C_2 :: \Omega_2}{\Omega_0 \vDash (C_1 \ C_2) :: \Omega_2} \text{compose}$$

To type processes and messages, we begin by considering *preservation*: we would like to achieve that if $\Omega' \vDash C :: \Omega$ and $C \mapsto C'$ then still $\Omega' \vDash C' :: \Omega$. Without the temporal modalities, this is guaranteed by the design of the sequent calculus: the right and left rules match just so that cut reduction (which is the basis for reduction in the operational semantics) leads to a well-typed deduction. The key here is what happens with time. Consider the special case

$$\frac{\Omega \vdash P :: A}{\Omega \vdash (\mathbf{delay} ; P) :: (x : \circ A)} \circ LR \quad \text{proc}(c, t, \mathbf{delay} ; P) \mapsto \text{proc}(c, t + 1, P)$$

Note that, inevitably, the type of the channel c changes in the transition, from $c : \circ A$ to $c : A$ and similarly for all channels used by P . So if in $\text{proc}(c, t, Q)$ we were to use the type of Q as the type of the semantic process object, preservation would fail. But while the type changes from $\circ A$ to A , *time* also advances from t to $t + 1$. This suggests the following rule should keep the configuration type invariant:

$$\frac{\Omega \vdash P :: (c : A)}{\circ^t \Omega \vDash \text{proc}(c, t, P) :: (c : \circ^t A)} \text{proc}^\circ$$

When we transition from $\mathbf{delay} ; P$ to P we strip one \circ modality from Ω and A , but because we also advance time from t to $t + 1$, the \circ modality is restored, keeping the interface type invariant.

When we also consider types $\square A$ and $\diamond A$ the situation is a little less straightforward because of their interaction with \circ , as we have already encountered in Section 4.4. We reuse the idea of the solution, allowing the subtraction of time from a type, possibly stopping when we meet a \square or \diamond .

$$\begin{aligned} [A]_L^{-0} &= A & [A]_R^{-0} &= A \\ [A]_L^{-(t+1)} &= [[A]_L^{-t}]_L^{-1} & [A]_R^{-(t+1)} &= [[A]_R^{-t}]_R^{-1} \end{aligned}$$

This is extended to channel declarations in the obvious way. Additionally, the imprecision of $\square A$ and $\diamond A$ may create temporal gaps in the configuration that need to be bridged by a weak form of subtyping $A <: B$ (not to be confused with the much stronger form $A \leq B$ in Section 6.1),

$$\frac{m \leq n}{\circ^m \square A <: \circ^n \square A} \square_{\text{weak}} \quad \frac{m \geq n}{\circ^m \diamond A <: \circ^n \diamond A} \diamond_{\text{weak}} \quad \frac{}{A <: A} \text{refl}$$

This relation is specified to be reflexive and clearly transitive. We extend it to contexts Ω in the obvious manner. In our final rules we also account for some channels that are not used by P or M but just passed through.

$$\frac{\Omega' <: \Omega \quad [\Omega]_L^{-t} \vdash P :: [c : A]_R^{-t} \quad A <: A'}{\Omega_0, \Omega' \vDash \text{proc}(c, t, P) :: (\Omega_0, c : A')} \text{proc} \quad \frac{\Omega' <: \Omega \quad [\Omega]_L^{-t} \vdash M :: [c : A]_R^{-t} \quad A <: A'}{\Omega_0, \Omega' \vDash \text{msg}(c, t, M) :: (\Omega_0, c : A')} \text{msg}$$

5.2 Type Preservation

With the four rules for typing configurations (empty, compose, proc and msg), type preservation is relatively straightforward. We need some standard lemmas about being able to split a configuration and be able to move a provider (whether process or message) to the right in a typing derivation until it rests right next to its client. Regarding time shifts, we need the following properties.

LEMMA 5.1 (TIME SHIFT).

- (i) If $[A]_L^{-t} = [B]_R^{-t}$ and both are defined then $A = B$.
- (ii) $[[A]_L^{-t}]_L^{-s} = [A]_L^{-(t+s)}$ and if either side is defined, the other is as well.
- (iii) $[[A]_R^{-t}]_R^{-s} = [A]_R^{-(t+s)}$ and if either side is defined, the other is as well.

THEOREM 5.2 (TYPE PRESERVATION). If $\Omega' \vDash C :: \Omega$ and $C \mapsto D$ then $\Omega' \vDash D :: \Omega$.

PROOF. By case analysis on the transition rule, applying inversion to the given typing derivation, and then assembling a new derivation of \mathcal{D} . \square

Type preservation on basic session types is a simple special case of this theorem.

5.3 Global Progress

We say a process or message is *poised* if it is trying to communicate along the channel that it provides. A poised process is comparable to a value in a sequential language. A configuration is poised if every process or message in the configuration is poised. Conceptually, this implies that the configuration is trying to communicate externally, i.e. along one of the channel it provides. The progress theorem then shows that either a configuration can take a step or it is poised. To prove this we show first that the typing derivation can be rearranged to go strictly from right to left and then proceed by induction over this particular derivation. This much is standard, even for significantly more complicated session-typed languages [Balzer and Pfenning 2017].

The question is how can we prove that processes are either at the same time (for most interactions) or that the message recipient is ready before the message arrives (for **when?**, **now!**, and some forwards)? The key insight here is in the following lemma.

LEMMA 5.3 (TIME INVERSION).

- (i) If $[A]_R^{-s} = [A]_L^{-t}$ and either side starts with a basic session type constructor then $s = t$.
- (ii) If $[A]_L^{-t} = \square B$ and $[A]_R^{-s} \neq \circ(-)$ then $s \leq t$ and $[A]_R^{-s} = \square B$.
- (iii) If $[A]_R^{-t} = \diamond B$ and $[A]_L^{-s} \neq \circ(-)$ then $s \leq t$ and $[A]_L^{-s} = \diamond B$.

THEOREM 5.4 (GLOBAL PROGRESS). If $\cdot \vdash C :: \Omega$ then either

- (i) $C \mapsto C'$ for some C' , or
- (ii) C is poised.

PROOF. By induction on the right-to-left typing of C so that either C is empty (and therefore poised) or $C = (\mathcal{D} \text{ proc}(c, t, P))$ or $C = (\mathcal{D} \text{ msg}(c, t, M))$. By induction hypothesis, \mathcal{D} can either take a step (and then so can C), or \mathcal{D} is poised. In the latter case, we analyze the cases for P and M , applying multiple steps of inversion and Lemma 5.3 to show that in each case either C can take a step or is poised. \square

6 TIME RECONSTRUCTION

The process expressions introduced so far have straightforward syntax-directed typing rules. This requires the programmer to write a significant number of explicit **delay**, **when?**, and **now!** constructs in their code. This in turn hampers reuse: we would like to be able to provide multiple types for the same process definition so it can be used in different contexts, with different types, even under a single, fixed cost model.

In this section we introduce an implicit system which may be thought of as a *temporal refinement* of the basic session type system in Section 2. The **delay**, **when?**, and **now!** constructs never appear in the source, and, as before, **tick** is added before type-checking and never by the programmer. The rules for the new judgment $\Omega \vdash P :: (x : A)$ are shown in Figure 6; the other rules remain the same (except for **def**, see below). We still need an explicit rule for the **tick** synonym of **delay** which captures the cost model.

These rules are trivially sound and complete with respect to the explicit system in Section 4 because from an implicit type derivation we can read off the explicit process expression and vice versa. They are also manifestly decidable because the types in the premises are smaller than those in the conclusion, with one possible exception: In the $\circ LR$ rule the premise may be equal to the

$$\begin{array}{c}
\frac{[\Omega]_L^{-1} \sharp P :: [x : A]_R^{-1}}{\Omega \sharp P :: (x : A)} \circ LR \qquad \frac{[\Omega]_L^{-1} \sharp P :: [x : A]_R^{-1}}{\Omega \sharp (\text{tick} ; P) :: (x : A)} \circ LR' \\
\frac{\Omega \sharp P :: (x : A)}{\Omega \sharp P :: (x : \diamond A)} \diamond R \qquad \frac{\Omega \text{ delayed}^\square \quad \Omega, x:A \sharp Q :: (z : C) \quad C \text{ delayed}^\diamond}{\Omega, x:\diamond A \sharp Q :: (z : C)} \diamond L \\
\frac{\Omega \text{ delayed}^\square \quad \Omega \sharp P :: (x : A)}{\Omega \sharp P :: (x : \square A)} \square R \qquad \frac{\Omega, x:A \sharp Q :: (z : C)}{\Omega, x:\square A \sharp Q :: (z : C)} \square L
\end{array}$$

Fig. 6. Implicit Temporal Rules

$$\begin{array}{c}
\frac{}{A \leq A} \text{ refl} \qquad \frac{A \leq B}{\square A \leq \square B} \square \square \qquad \frac{\square A \leq B}{\square A \leq \square B} \square \square \qquad \frac{A \leq \diamond B}{\square A \leq \diamond B} \square \diamond \\
\frac{\square^n \square A \leq B}{\square^n \square A \leq \square B} \square R \qquad \frac{A \leq B}{\square A \leq B} \square L \qquad \frac{A \leq B}{A \leq \diamond B} \diamond R \qquad \frac{A \leq \square^n \diamond B}{\diamond A \leq \square^n \diamond B} \diamond L
\end{array}$$

Fig. 7. Subtyping Rules

conclusion if neither Ω nor A contain a type of the form $\square(-)$. In this case, $B = \square B'$ for every $y : B$ in Ω and $A = \diamond A'$ and there P can delay by any finite number of time steps. Time reconstruction avoids such an arbitrary delay.

Our examples revealed a significant shortcoming in these rules: when calling upon a process definition, the types in the antecedent and succedent often do not match the types of the process to be spawned. For example, the process *skip1s* in Section 4.1 we have

```

bits =  $\oplus\{b0 : \square\text{bits}, b1 : \square\text{bits}, \$ : \square 1\}$ 
sbits =  $\oplus\{b0 : \square\text{sbits}, b1 : \square\diamond\text{sbits}, \$ : \square 1\}$ 

y : bits  $\vdash$  compress :: (x :  $\square\text{sbits}$ )
y : bits  $\vdash$  skip1s :: (x :  $\square\diamond\text{sbits}$ )

x  $\leftarrow$  skip1s  $\leftarrow$  y =
  case y (b1  $\Rightarrow$  tick ;
    x  $\leftarrow$  skip1s  $\leftarrow$  y
    | ...)
    % y : bits  $\vdash$  x :  $\diamond\text{sbits}$ 
    % does not type-check!

```

The indicated line does not type-check (neither in the explicit nor the implicit system presented so far) since the type $\square\diamond\text{sbits}$ offered by *skip1s* does not match $\diamond\text{sbits}$. We had to write a process *idle* to account for this mismatch:

```

x' :  $\square\diamond A \vdash$  idle :: (x :  $\diamond A$ )
x  $\leftarrow$  idle  $\leftarrow$  x' = delay ; x  $\leftarrow$  x'

```

In the implicit system the version with an explicit identity *can* in fact be reconstructed:

```

x  $\leftarrow$  skip1s  $\leftarrow$  y =
  case y (b1  $\Rightarrow$  tick ;
    x'  $\leftarrow$  skip1s  $\leftarrow$  y
    | ...)
    % y : bits  $\sharp$  x :  $\diamond\text{sbits}$ 
    % x' :  $\square\diamond\text{sbits} \sharp$  x :  $\diamond\text{sbits}$ 
    % x' :  $\diamond\text{sbits} \sharp$  x :  $\diamond\text{sbits}$  using rule  $\circ LR$ 
    x  $\leftarrow$  x'

```

6.1 Subtyping

Extrapolating from the example of *skip1s* above, we can generalize process invocations by allowing *subtyping* on all used channels. The implicit rule for process invocation then reads

$$\frac{\Omega' \leq \Omega_f \quad (\Omega_f \vdash f = P_f :: (x : A)) \in \Sigma \quad \Omega, x:A \vdash Q :: (z : C)}{\Omega, \Omega' \vdash (x \leftarrow f \leftarrow \Omega' ; Q) :: (z : C)} \text{ def}$$

But how do we define subtyping $A \leq B$? We would like the coercion to be an identity on basic session types and just deal with temporal mismatches through appropriate **delay**, **when?**, and **now!** actions. In other words, A should be a subtype of B if and only if $y : A \vdash x \leftarrow y :: (x : B)$. Given this desired theorem, we can just read off the subtyping rules from the implicit typing rules in Figure 6 by using the forwarding process $x \leftarrow y$ as the subject in each rule! This form of subtyping is independent from subtyping between basic session types [Gay and Hole 2005], which we believe can be added to our system in a sound way, even if it would not be complete for asynchronous communication [Lange and Yoshida 2017].

This approach yields the rules in Figure 7, where we have split the $\circ LR$ rule into three different cases. We have expanded the definitions of *patient* types to make it syntactically more self-contained.

THEOREM 6.1 (SUBTYPING IDENTITY). $A \leq B$ iff $y : A \vdash x \leftarrow y :: (x : B)$

PROOF. In each direction by induction over the structure of the given deduction. \square

The subtyping rules are manifestly decidable. In the bottom-up search for a subtyping derivation, the rules $\circ\circ$, $\square R$, and $\diamond L$ can be applied eagerly without losing completeness. There is a nontrivial decision point between the $\square\circ$ and $\square L$ rules. The examples $\square S \leq \circ\square S$ and $\square\circ S \leq \circ S$ for a basic session type S show that sometimes $\square\circ$ must be chosen and sometimes $\square L$ when both rules apply. A dual non-deterministic choice exists between $\circ\diamond$ and $\diamond R$. The cost of backtracking is minimal in all examples we have considered.

We already know that subtype coercions are identities. To verify that we have a sensible subtype relation it remains to prove that transitivity is admissible. For this purpose we need two lemmas regarding patient types, as they appear in the $\square R$ and $\diamond L$ rules.

LEMMA 6.2 (PATIENCE).

- (i) If $A \leq \circ^n \square B$ then $A = \circ^k \square A'$ for some k and A' .
- (ii) If $\circ^n \diamond A \leq B$ then $B = \circ^k \diamond B'$ for some k and B' .

PROOF. By separate inductions over the structure of the given deductions. \square

LEMMA 6.3 (IMPATIENCE).

- (i) If $\circ\circ^n \square A \leq B$ then $\circ^n \square A \leq B$.
- (ii) If $A \leq \circ\circ^n \diamond B$ then $A \leq \circ^n \diamond B$.

PROOF. By separate inductions over the structure of the given deductions. \square

THEOREM 6.4 (TRANSITIVITY OF SUBTYPING).

If $A \leq B$ and $B \leq C$ then $A \leq C$.

PROOF. By simultaneous induction on the structure of the deductions \mathcal{D} of $A \leq B$ and \mathcal{E} of $B \leq C$ with appeals to the preceding lemmas in four cases. \square

7 FURTHER EXAMPLES

In this section we present example analyses of some of the properties that we can express in the type system, such as the message rates of streams, the response time of concurrent data structures, and the span of a fork/join parallel program.

In some examples we use parametric definitions, both at the level of types and processes. For example, stack_A describes stacks parameterized over a type A , $\text{list}_A[n]$ describes lists of n elements, and $\text{tree}[h]$ describes binary trees of height h . Process definitions are similarly parameterized. We think of these as families of ordinary definitions and calculate with them accordingly, at the metalevel, which is justified since they are only implicitly quantified across whole definitions. This common practice (for example, in work on interaction nets [Gimenez and Moser \[2016\]](#)) avoids significant syntactic overhead, highlighting conceptual insight. It is of course possible to internalize such parameters (see, for example, work on refinement of session types [\[Griffith and Gunter 2013\]](#) or explicitly polymorphic session types [\[Caires et al. 2013; Griffith 2016\]](#)).

7.1 Response Times: Stacks and Queues

To analyze response times, we study concurrent stacks and queues. A stack data structure provides a client with a choice between a push and a pop. After a push, the client has to send an element, and the provider will again behave like a stack. After a pop, the provider will reply either with the label none and terminate (if there are no elements in the stack), or send an element and behave again like a stack. In the cost-free model, this is expressed in the following session type.

$$\begin{aligned} \text{stack}_A &= \&\{ \text{push} : A \multimap \text{stack}_A, \\ &\quad \text{pop} : \oplus\{ \text{none} : \mathbf{1}, \text{some} : A \otimes \text{stack}_A \} \} \end{aligned}$$

We implement a stack as a chain of processes. The bottom to the stack is defined by the process *empty*, while a process *elem* holds a top element of the stack as well as a channel with access to the top of the remainder of the stack.

$$\begin{aligned} x : A, t : \text{stack}_A \vdash \text{elem} &:: (s : \text{stack}_A) \\ \cdot \vdash \text{empty} &:: (s : \text{stack}_A) \end{aligned}$$

The cost model we would like to consider here is \mathcal{RS} where both receives and sends cost one unit of time. Because a receive costs one unit, every continuation type must be delayed by one tick of the clock, which we have denoted by prefixing continuations by the \circ modality. This delay is not an artifact of the implementation, but an inevitable part of the cost model—one reason we have distinguished the synonyms **tick** (delay of one, due to the cost model) and **delay** (delay of one, to correctly time the interactions). In this section of examples we will make the same distinction for the next-time modality: we write $\prime A$ for a step in time mandated by the cost model, and $\circ A$ for a delay necessitated by a particular set of process definitions.

As a first approximation, we would have

$$\begin{aligned} \text{stack}_A &= \&\{ \text{push} : \prime(A \multimap \prime \text{stack}_A), \\ &\quad \text{pop} : \oplus\{ \text{none} : \prime \mathbf{1}, \text{some} : \prime(A \otimes \prime \text{stack}_A) \} \} \end{aligned}$$

There are several problems with this type. The stack is a data structure and has little or no control over *when* elements will be pushed onto or popped from the stack. Therefore we should use a type $\square \text{stack}_A$ to indicate that the client can choose the times of interaction with the stack. While the elements are held by the stack time advances in an indeterminate manner. Therefore, the elements stored in the stack must also have type $\square A$, not A (so that they are always available).

$$\begin{aligned} \text{stack}_A &= \&\{ \text{push} : \prime(\square A \multimap \prime \square \text{stack}_A), \\ &\quad \text{pop} : \oplus\{ \text{none} : \prime \mathbf{1}, \text{some} : \prime(\square A \otimes \prime \square \text{stack}_A) \} \} \\ x : \square A, t : \square \text{stack}_A \vdash \text{elem} &:: (s : \square \text{stack}_A) \end{aligned}$$

close s)

The difference between the *response times* of stacks and queues in the cost model is minimal: both are constant, with the queue being two units slower. This is in contrast to the total work [Das et al. 2017] which is constant for the stack but linear in the number of elements for the queue.

This difference in response times can be realized by typing clients of both stacks and queues. We compare clients S_n and Q_n that insert n elements into a stack and queue, respectively, send the result along channel d , and then terminate. We show only their type below, omitting the implementations.

$$\begin{aligned} x_1 : \Box A, \dots, x_n : \Box A, s : \Box \text{stack}_A \vdash S_n &:: (d : \circ^{2n} (\Box \text{stack}_A \otimes '1)) \\ x_1 : \Box A, \dots, x_n : \Box A, s : \Box \text{queue}_A \vdash Q_n &:: (d : \circ^{4n} (\Box \text{queue}_A \otimes '1)) \end{aligned}$$

The types demonstrate that the total execution time of S_n is only $2n + 1$, while it is $4n + 1$ for Q_n . The difference comes from the difference in response times. Note that we can infer precise execution times, even in the presence of the \Box modality in the stack and queue types.

7.2 Parametric Rates: Lists and Streams

Lists describe an interface that sends either `nil` and ends the session, or sends `cons` followed by a channel of some type A and then behaves again like a list. In the cost-free setting:

$$\text{list}_A = \oplus\{ \text{cons} : A \otimes \text{list}_A, \text{nil} : 1 \}$$

Here is the straightforward definition of *append*.

$$l_1 : \text{list}_A, l_2 : \text{list}_A \vdash \text{append} : (l : \text{list}_A)$$

$$\begin{aligned} l \leftarrow \text{append} \leftarrow l_1 \ l_2 = & \\ \text{case } l_1 \ (\text{cons} \Rightarrow & x \leftarrow \text{recv } l_1 ; & \% \text{ receive element } x \text{ from } l_1 \\ & l.\text{cons} ; \text{send } l \ x ; & \% \text{ send } x \text{ along } l \\ & l \leftarrow \text{append} \leftarrow l_1 \ l_2 & \% \text{ recurse} \\ | \text{nil} \Rightarrow & \text{wait } l_1 ; & \% \text{ wait for } l_1 \text{ to close} \\ & l \leftarrow l_2 & \% \text{ identify } l \text{ and } l_2 \end{aligned}$$

In this example we are interested in analyzing the timing of several processes precisely, but parametrically over an arrival rate. Because it takes two units of time to copy the inputs to the outputs, the arrival rate needs to be at least 2, which we represent by writing it as $r + 2$. Since we append the two lists, the second list will be idle while we copy the elements from the first list to the output. We could give this list type $\Box(-)$, but we can also precisely determine the delay if we index lists by the number of elements. We write $\text{list}_A[n]$ for a list sending exactly n elements. We have the following types in the \mathcal{RS} cost model:

$$\begin{aligned} \text{list}_A[0] &= \oplus\{ \text{nil} : '1 \} \\ \text{list}_A[n+1] &= \oplus\{ \text{cons} : '(\Box A \otimes '\circ^{r+2} \text{list}_A[n]) \} \end{aligned}$$

As before, the tick marks account for the delay mandated by the cost model. The \circ^{r+2} accounts for the arrival rate of $r + 2$. We use type $\Box A$ for the elements since they will be in the lists for an indeterminate amount of time. The precise type of *append* then becomes

$$l_1 : \text{list}_A[n], l_2 : \circ^{(r+4)n+2} \text{list}_A[k] \vdash \text{append} :: (l : \circ\circ \text{list}_A[n+k])$$

It expresses that the output list has the same rate as the input lists, but with a delay of 2 cycles relative to l_1 . The channel l_2 has to sit idle for $r + 4$ cycles for each element of l_1 , accounting for the two inputs along l_1 and two outputs along l_2 . It takes 2 further cycles to input the `nil` and the end token for the list.

With our type system and just a little bit of arithmetic we can verify this type, checking the definition twice: once for a list of length 0 and once for $n + 1$. We show here the latter, where $l_1 : \text{list}_A[n + 1]$.

$$\begin{aligned}
& l \leftarrow \text{append} \leftarrow l_1 \ l_2 = \\
& \text{case } l_1 \ (\text{cons} \Rightarrow \quad \% l_1 : \square A \otimes ' \circ^{r+2} \text{list}_A[n], l_2 : [\circ^{(r+4)(n+1)+2} \text{list}_A[k]]_L^{-1} \vdash l : \circ \text{list}_A[(n+1)+k] \\
& \quad x \leftarrow \text{recv } l_1 ; \quad \% x : \square A, l_1 : \circ^{r+2} \text{list}_A[n], l_2 : [\circ^{(r+4)(n+1)+2} \text{list}_A[k]]_L^{-2} \vdash l : \text{list}_A[(n+1)+k] \\
& \quad l.\text{cons} ; \quad \% x : \square A, l_1 : \circ^{r+1} \text{list}_A[n], l_2 : [\circ^{(r+4)(n+1)+2} \text{list}_A[k]]_L^{-3} \vdash l : \square A \otimes ' \circ^{r+2} \text{list}_A[n+k] \\
& \quad \text{send } l \ x ; \quad \% l_1 : \circ^r \text{list}_A[n], l_2 : [\circ^{(r+4)(n+1)+2} \text{list}_A[k]]_L^{-4} \vdash l : \circ^{r+2} \text{list}_A[n+k] \\
& \quad \% \text{delay}' \quad \% l_1 : \text{list}_A[n], l_2 : [\circ^{(r+4)(n+1)+2} \text{list}_A[k]]_L^{-4-r} \vdash l : \circ^2 \text{list}_A[n+k] \\
& \quad \quad \% l_1 : \text{list}_A[n], l_2 : \circ^{(r+4)n+2} \text{list}_A[k] \vdash l : \circ \circ \text{list}_A[n+k] \\
& \quad l \leftarrow \text{append} \leftarrow l_1 \ l_2 \\
& \quad | \text{nil} \Rightarrow \dots)
\end{aligned}$$

We showed only the one delay by r units inserted by time reconstruction since it is the critical step. The case for nil does not apply for $l_1 : \text{list}_A[n+1]$. Here is the typing derivation when $l_1 : \text{list}_A[0]$ where the cons branch does not apply.

$$\begin{aligned}
& l \leftarrow \text{append} \leftarrow l_1 \ l_2 = \\
& \quad \text{case } l_1 \ (\text{cons} \Rightarrow \dots \\
& \quad \quad | \text{nil} \Rightarrow \quad \% l_1 : \mathbf{1}, l_2 : \circ \text{list}_A[k] \vdash l : \circ \text{list}_A[k] \\
& \quad \quad \quad \text{wait } l_1 ; \quad \% l_2 : \text{list}_A[k] \vdash l : \text{list}_A[k] \\
& \quad \quad \quad l \leftarrow l_2)
\end{aligned}$$

As a related example we consider a process that alternates the elements between two infinite input streams. At first we might expect if the two input streams come in with a rate of 2 then the output stream will have a rate of 1. However, in the \mathcal{RS} cost model one additional tick is required for sending on the messages which means that the input streams need to have rate 3 and be offset by 2 cycles. We parameterize the type of stream by its rate k

$$\begin{aligned}
& \text{stream}_A^k = \square A \otimes ' \circ^k \text{stream}_A^k \\
& l_1 : \text{stream}_A^3, l_2 : \circ^2 \text{stream}_A^3 \vdash \text{alternate} :: (l : \circ^1 \text{stream}_A^1) \\
& l \leftarrow \text{alternate} \leftarrow l_1 \ l_2 = \\
& \quad x \leftarrow \text{recv } l_1 ; \quad \% x : \square A, l_1 : \circ^3 \text{stream}_A^3, l_2 : \circ^1 \text{stream}_A^3 \vdash l : \text{stream}_A^1 \\
& \quad \text{send } l \ x ; \quad \% \quad \quad \quad l_1 : \circ^2 \text{stream}_A^3, l_2 : \text{stream}_A^3 \vdash l : \circ^1 \text{stream}_A^1 \\
& \quad l \leftarrow \text{alternate} \leftarrow l_2 \ l_1)
\end{aligned}$$

A more general parametric type for the same code would be

$$l_1 : \text{stream}_A^{2k+3}, l_2 : \circ^{k+2} \text{stream}_A^{2k+3} \vdash \text{alternate} :: (l : \circ^1 \text{stream}_A^{k+1})$$

from which we can recover the more specialized one with $k = 0$.

7.3 Span Analysis: Trees

We use trees to illustrate an example that is typical for fork/join parallelism and computation of *span*. In order to avoid integers, we just compute the parity of a binary tree of height h with boolean values at the leaves. We do not show the obvious definition of *xor*, which in the \mathcal{RS} cost model requires a delay of four from the first input.

$$\begin{aligned}
& \text{bool} = \oplus \{ b0 : '1, b1 : '1 \} \\
& a : \text{bool}, b : \circ^2 \text{bool} \vdash \text{xor} :: (c : \circ^4 \text{bool})
\end{aligned}$$

In the definition of *leaf* and *node* we have explicated the delays inferred by time reconstruction, but not the *tick* delays. The type of *tree*[h] gives the *span* of this particular parallel computation as $5h + 2$. This is the time it takes to compute the parity under maximal parallelism, assuming that *xor* takes 4 cycles as shown in the type above.

$$\text{tree}[h] = \&\{ \text{parity} : ' \circ^{5h+2} \text{bool} \}$$

```

· ⊢ leaf :: (t : tree[h])
t ← leaf =
  case t (parity ⇒
    % delay5h+2 % · ⊢ t : ○5h+2 bool
    % · ⊢ t : bool
    t.b0 ; % · ⊢ t : 1
  close t)

l : ○1tree[h], r : ○3tree[h] ⊢ node :: (t : tree[h + 1])
t ← node ← l r =
  case t (parity ⇒
    % l : tree[h], r : ○2tree[h] ⊢ t : ○5(h+1)+2 bool
    l.parity ; % l : ○5h+2 bool, r : ○1tree[h] ⊢ t : ○5(h+1)+1 bool
    % delay % l : ○5h+1 bool, r : tree[h] ⊢ t : ○5h+5 bool
    r.parity ; % l : ○5h bool, r : ○5h+2 bool ⊢ t : ○5h+4 bool
    % delay5h % l : bool, r : ○2bool ⊢ t : ○4 bool
  t ← xor ← l r)

```

The type $l : \circ^1 \text{tree}[h]$ comes from the fact that, after receiving a parity request, we first send out the parity request to the left subtree l . The type $r : \circ^3 \text{tree}[h]$ is determined from the delay of 2 between the two inputs to xor . The magic number 5 in the type of $tree$ was derived in reverse from setting up the goal of type-checking the $node$ process under the constraints already mentioned. We can also think of it as $4+1$, where 4 is the time to compute the exclusive or at each level and 1 as the time to propagate the parity request down each level.

As is often done in abstract complexity analysis, we can also impose an alternative cost model. For example, we may count only the number of calls to xor while all other operations are cost free. Then we would have

```

a : bool, b : bool ⊢ xor :: (c : ○bool)      · ⊢ leaf :: (t : tree[h])
tree[h] = &{ parity : ○h bool }           l : tree[h], r : tree[h] ⊢ node :: (t : tree[h + 1])

```

with the same code but different times and delays from before. The reader is invited to reconstruct the details.

7.4 A Higher-Order Example

As an example of higher-order programming we show how to encode a process analogue of a fold function. Because our language is purely linear the process to fold over a list has to be recursively defined. In the cost-free setting we would write

```

folderAB = &{ next : A → (B → (B ⊗ folderAB)), done : 1 }
l : listA, f : folderAB, b : B ⊢ fold :: (r : B)
r ← fold ← l f b =
  case l (cons ⇒ x ← recv l ;
    f.next ; send f x ; send f b ; % send x and b to folder f
    y ← recv f ; r ← fold ← l f y % receive y from f and recurse
  | nil ⇒ wait l ; f.done ; wait f ; r ← b)

```

If we want to assign precise temporal types to the $fold$ process then the incoming list should have a delay of at least 4 between successive elements. Working backwards from the code we obtain the following types.

```

listA[0] = ⊕{ nil : '1 }
listA[n + 1] = ⊕{ cons : '(□A ⊗ '○k+4 listA[n]) }

```

$$\text{folder}_{AB} = \&\{ \text{next} : '(\Box A \multimap '(B \multimap ' \circ^k (\circ^5 B \otimes ' \circ^2 \text{folder}_{AB}))), \text{done} : '1 \}$$

$$l : \text{list}_A[n], f : \circ^2 \text{folder}_{AB}, b : \circ^4 B \vdash \text{fold} :: (r : \circ^{(k+5)n+4} B)$$

The type of *fold* indicates that if the combine function of folder_{AB} takes k time units to compute, the result r is produced after $(k + 5)n + 4$ time units in the \mathcal{RS} cost model.

8 RELATION TO THE STANDARD SEMANTICS

While our temporal semantics stands on its own, one may ask precisely in which way it captures properties of the standard semantics. We analyze this here for the fragment with only the next-time modality $\circ A$; the general case including $\Box A$ and $\Diamond A$ is just slightly more complicated. We proceed in several steps.

Step 1: Standard Semantics. The standard operational semantics for basic session types (without temporal modalities) is precisely that in Figure 4 where the time t always remains at 0.

Step 2: Measuring Span. We are interested in capturing the *span* of a standard computation, which is the number of steps required to completion under the assumption that any action takes place as soon as possible, subject only to its dependencies. We can measure this by *instrumenting* the standard semantics, as was done by Silva et al. [2016], except that we increment time on every step instead of just when messages are received. We write $\text{proc}^*(c, t, P)$ and $\text{msg}^*(c, t, M)$ to distinguish them because the time t has a different interpretation, namely the earliest time the process could have arrived at this point in the computation under an asynchronous model of communication. We show the two rules for internal choice to illustrate this semantics.

$$\text{proc}^*(c, t, c.k ; P) \mapsto \text{proc}^*(c', t + 1, [c'/c]P), \text{msg}^*(c, t + 1, c.k ; c \leftarrow c') \quad (c' \text{ fresh})$$

$$\text{msg}^*(c, t, c.k ; c \leftarrow c'), \text{proc}^*(d, t', \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}) \mapsto \text{proc}^*(d, \max(t, t' + 1), [c'/c]Q_k)$$

Step 3: Relating Computations. We define $|A|$ and $|P|$ as the result of erasing all next-time modalities from A and **delay** actions from $|P|$, respectively. We would like to relate computations of $\Omega \vdash P : A$ to those of $|\Omega| \vdash |P| :: (c : |A|)$. However, we have to ensure that there are enough **delay** operators in P to account for the fact that the cost model for standard computations counts every step. To this end, we define two mutually recursive relations $P \geq Q$ and $P > Q$ where $Q = |P|$ but there are further constraints on P . $P \geq Q$ expresses that P and Q start with the same action (which cannot be a **delay**) and the continuations P' and Q' are related with $P' > Q'$. This in turn requires one or more initial delays in P' with its continuation $P'' \geq Q'$. These relations embody the idea that a delay of the cost model *precedes* each action, which is necessary since close and forwarding actions have no continuation. We only show the rules for sending and receiving labels.

$$\frac{P > Q}{c.k ; P \geq c.k ; Q} \quad \frac{(\forall \ell \in L) P_\ell > Q_\ell}{\text{case } c (\ell \Rightarrow P_\ell)_{\ell \in L} \geq \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}} \quad \frac{P > Q}{\text{delay} ; P > Q} \quad \frac{P \geq Q}{\text{delay} ; P > Q}$$

We then extend this relation to all semantic objects, expressing that the temporal semantics may be an over-approximation of the standard semantics because it may contain arbitrarily many additional delay actions. This is expressed formally in the conditions on the time stamps of corresponding processes and messages.

$$\text{proc}(c, s, P) \geq \text{proc}^*(c, t, Q) \quad \text{if } (P > Q \text{ and } s \geq t) \text{ or } (P \geq Q \text{ and } s > t)$$

$$\text{msg}(c, s, M) \geq \text{msg}^*(c, t, M) \quad \text{if } s \geq t$$

We then extend this compositionally to full configurations, $C \geq \mathcal{D}$.

Bisimulation. The \geq relation is a weak bisimulation in the following sense, where $\mapsto^{\leq 1}$ means at most one transition and $\mapsto^{\geq 1}$ means at least one transition. Note that the approximation property comes from the relation between time stamps in the two configurations, not the number of transitions.

THEOREM 8.1. *Assume $\cdot \vDash C :: \Omega$ and $\cdot \vDash \mathcal{D} :: |\Omega|$.*

- (i) *If $C \geq \mathcal{D}$ and $C \mapsto C'$ then $\mathcal{D} \mapsto^{\leq 1} \mathcal{D}'$ for some \mathcal{D}' with $C' \geq \mathcal{D}'$.*
- (ii) *If $C \geq \mathcal{D}$ and $\mathcal{D} \mapsto \mathcal{D}'$ then $C \mapsto^{\geq 1} C'$ for some C' with $C' \geq \mathcal{D}'$.*

PROOF. In each direction, by analyzing each case of the given reduction, with a case analysis or induction over the definition of \geq . □

The treatment of $\square A$ and $\diamond A$ is slightly more complicated. We obtain the simplest generalization by defining $|\square A| = \&\{\text{now} : |A|\}$ and $|\diamond A| = \oplus\{\text{now} : |A|\}$ with the corresponding erasure in the process expressions. If we want to avoid such “administrative messages” we can instead fully erase all temporal constructs but enforce a normal form on process expressions where every **when?** x action is always immediately followed by another receive action along x .

9 FURTHER RELATED WORK

In addition to the related work already mentioned, we highlight a few related threads of research.

Session types and process calculi. In addition to the work on timed multiparty session types [Bocchi et al. 2014; Neykova et al. 2014], time has been introduced into the π -calculus (see, for example, Saeedloei and Gupta [2014]) or session-based communication primitives (see, for example, López et al. [2009]) but generally these works do not develop a type system. Kobayashi [2002] extends a (synchronous) π -calculus with means to count parallel reduction steps. He then provides a type system to verify time-boundedness. This is more general in some dimension than our work because of a more permissive underlying type and usage system, but it lacks internal and external choice, genericity in the cost model, and provides bounds rather than a fine gradation between exact and indefinite times. Session types can also be derived by a Curry-Howard interpretation of *classical linear logic* [Wadler 2012] but we are not aware of temporal extensions. We conjecture that there is a classical version of our system where \square and \diamond are dual and \circ is self-dual.

Reactive programming. Synchronous data flow languages such as Lustre [Halbwachs et al. 1991], Esterel [Berry and Gonthier 1992], or Lucid Synchrone [Pouzet 2006] are time-synchronous with uni-directional flow and thus may be compared to the fragment of our language with internal choice (\oplus) and the next-time modality ($\circ A$), augmented with existential quantification over basic data values like booleans and integers (which we have omitted here only for the sake of brevity). The global clock would map to our underlying notion of time, but data-dependent local clocks would have to be encoded at a relatively low level using streams of option type, compromising the brevity and elegance of these languages. Furthermore, synchronous data flow languages generally permit sharing of channels, which, although part of many session-typed languages [Balzer and Pfenning 2017; Caires and Pfenning 2010], require further investigation in our setting. On the other hand, we support a number of additional types such as external choice ($\&$) for bidirectional communication and higher-order channel-passing ($A \multimap B$, $A \otimes B$). In the context of functional reactive programming, a Nakano-style [Nakano 2000] temporal modality has been used to ensure productivity [Krishnaswami and Benton 2011]. A difference in our work is that we consider concurrent processes and that our types prescribe the timing of messages.

Computational interpretations of $\circ A$. A first computational interpretation of the next-time modality under a proofs-as-programs paradigm was given by Davies [1996]. The basis is natural deduction for a (non-linear!) intuitionistic linear-time temporal logic with only the next-time modality. Rather than capturing cost, the programmer could indicate *staging* by stipulating that some subexpressions should be evaluated “at the next time”. The natural operational semantics then is a logically-motivated form of *partial evaluation* which yields a residual program of type $\circ A$. This idea was

picked up by [Feltman et al. \[2016\]](#) to instead *split* the program statically into two stages where results from the first stage are communicated to the second. Again, neither linearity (in the sense of linear logic), nor any specific cost semantics appears in this work.

Other techniques. Inferring the cost of concurrent programs is a fundamental problem in resource analysis. [Hoffmann and Shao \[2015\]](#) introduce the first automatic analysis for deriving bounds on the worst-case evaluation cost of parallel first-order functional programs. Their main limitation is that they can only handle parallel computation; they don't support message-passing or shared memory based concurrency. [Blelloch and Reid-Miller \[1997\]](#) use pipelining [[Paul et al. 1983](#)] to improve the complexity of parallel algorithms. However, they use futures [[Halstead 1985](#)], a parallel language construct to implement pipelining without the programmer having to specify them explicitly. The runtime of algorithms is determined by analyzing the work and depth in a language-based cost model. The work relates to ours in the sense that pipelines can have delays, which can be data dependent. However, the algorithms they analyze have no message-passing concurrency or other synchronization constructs. [Albert et al. \[2015\]](#) devised a static analysis for inferring the parallel cost of distributed systems. They first perform a block-level analysis to estimate the serial cost, then construct a distributed flow graph (DFG) to capture the parallelism and then obtain the parallel cost by computing the maximal cost path in the DFG. However, the bounds they produce are modulo a points-to and serial cost analysis. Hence, an imprecise points-to analysis will result in imprecise parallel cost bounds. Moreover, since their technique is based on static analysis, it is not compositional and a whole program analysis is needed to infer bounds on each module. Recently, a bounded linear typing discipline [[Ghica and Smith 2014](#)] modeled in a semiring was proposed for resource-sensitive compilation. It was then used to calculate and control execution time in a higher-order functional programming language. However, this language did not support recursion.

10 CONCLUSION

We have developed a system of temporal session types that can accommodate and analyze concurrent programs with respect to a variety of different cost models. Types can vary in precision, based on desired and available information, and includes latency, rate, response time, and span of computations. It is constructed in a modular way, on top of a system of basic session types, and therefore lends itself to easy generalization. We have illustrated the type system through a number of simple programs on streams of bits, binary counters, lists, stacks, queues, and trees. Time reconstruction and subtyping go some way towards alleviating demands on the programmer and supporting program reuse. In ongoing work we are exploring an implementation with an eye toward practical aspects of time reconstruction and, beyond that, automatic resource analysis based on internal measures of processes such as the length of a list or the height of a tree—so far, we have carried out these analyses by hand.

ACKNOWLEDGMENTS

This article is based on research that has been supported, in part, by AFRL under DARPA STAC award FA8750-15-C-0082 and in part by the NSF Foundation under Grants No. 1718267 and 1812876. Any opinions, findings, and conclusions contained in this document are those of the authors and do not necessarily reflect the views of the sponsoring organizations.

REFERENCES

Elvira Albert, Jesús Correias, Einar Broch Johnsen, and Guillermo Román-Díez. 2015. Parallel Cost Analysis of Distributed Systems. In *Static Analysis*, Sandrine Blazy and Thomas Jensen (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 275–292.

- Martin Avanzini, Ugo Dal Lago, and Georg Moser. 2015. Analysing the Complexity of Functional Programs: Higher-Order Meets First-Order. In *29th Int. Conf. on Functional Programming (ICFP'15)*.
- Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. In *International Conference on Functional Programming (ICFP)*. ACM, 37:1–37:29.
- G erard Berry and Georges Gonthier. 1992. The ESTEREL Synchronous Programming Language: Design, Semantics, Implementation. *Sci. Comput. Program.* 19, 2 (Nov. 1992), 87–152. [https://doi.org/10.1016/0167-6423\(92\)90005-V](https://doi.org/10.1016/0167-6423(92)90005-V)
- Guy E. Blelloch and Margaret Reid-Miller. 1997. Pipelining with Futures. In *Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA '97)*. ACM, New York, NY, USA, 249–259. <https://doi.org/10.1145/258492.258517>
- Laura Bocchi, Weizhen Yang, and Nobuko Yoshida. 2014. Timed Multiparty Session Types. In *CONCUR 2014 – Concurrency Theory*, Paolo Baldan and Daniele Gorla (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 419–434.
- Luis Caires, Jorge A. P erez, Frank Pfenning, and Bernardo Toninho. 2013. Behavioral Polymorphism and Parametricity in Session-Based Communication. In *Proceedings of the European Symposium on Programming (ESOP'13)*, M.Felleisen and P.Gardner (Eds.). Springer LNCS 7792, Rome, Italy, 330–349.
- Luis Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *Proceedings of the 21st International Conference on Concurrency Theory (CONCUR 2010)*. Springer LNCS 6269, Paris, France, 222–236.
- Luis Caires, Frank Pfenning, and Bernardo Toninho. 2016. Linear Logic Propositions as Session Types. *Mathematical Structures in Computer Science* 26, 3 (2016), 367–423.
- Iliano Cervesato and Andre Scedrov. 2009. Relating State-Based and Process-Based Concurrency through Linear Logic. *Information and Computation* 207, 10 (Oct. 2009), 1044–1077.
- Norman Danner, Daniel R. Licata, and Ramyaa Ramyaa. 2015. Denotational Cost Semantics for Functional Languages with Inductive Types. In *29th Int. Conf. on Functional Programming (ICFP'15)*.
- Ankush Das, Jan Hoffmann, and Frank Pfenning. 2017. Work Analysis with Resource-Aware Session Types. *CoRR abs/1712.08310* (2017). arXiv:1712.08310 <http://arxiv.org/abs/1712.08310>
- Rowan Davies. 1996. A Temporal Logic Approach to Binding-Time Analysis. In *Proceedings of the Eleventh Annual Symposium on Logic in Computer Science*, E. Clarke (Ed.). IEEE Computer Society Press, New Brunswick, New Jersey, 184–195. <http://www.cs.cmu.edu/afs/cs/user/rowan/www/papers/multbta.ps.Z>
- Nicolas Feltman, Carlo Angiuli, Umut Acar, and Kayvon Fatahalian. 2016. Automatically Splitting a Two-Stage Lambda Calculus. In *Proceedings of the 25th European Symposium on Programming (ESOP)*, P. Thiemann (Ed.). Springer LNCS 9632, Eindhoven, The Netherlands, 255–281.
- J er me Fortier and Luigi Santocanale. 2013. Cuts for Circular Proofs: Semantics and Cut Elimination. In *22nd Conference on Computer Science Logic (LIPICs)*, Vol. 23. 248–262.
- Simon J. Gay and Malcolm Hole. 2005. Subtyping for Session Types in the π -Calculus. *Acta Informatica* 42, 2–3 (2005), 191–225.
- Dan R. Ghica and Alex I. Smith. 2014. Bounded Linear Types in a Resource Semiring. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410*. Springer-Verlag New York, Inc., New York, NY, USA, 331–350. https://doi.org/10.1007/978-3-642-54833-8_18
- St ephane Gimenez and Georg Moser. 2016. The Complexity of Interaction. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '16)*. ACM, New York, NY, USA, 243–255. <https://doi.org/10.1145/2837614.2837646>
- Dennis Griffith. 2016. *Polarized Substructural Session Types*. Ph.D. Dissertation. University of Illinois at Urbana-Champaign.
- Dennis Griffith and Elsa L. Gunter. 2013. Liquid Pi: Inferrable Dependent Session Types. In *Proceedings of the NASA Formal Methods Symposium*. Springer LNCS 7871, 186–197.
- Sumit Gulwani, Krishna K. Mehra, and Trishul M. Chilimbi. 2009. SPEED: Precise and Efficient Static Estimation of Program Computational Complexity. In *36th ACM Symp. on Principles of Prog. Langs. (POPL'09)*. 127–139.
- N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud. 1991. The synchronous data flow programming language LUSTRE. *Proc. IEEE* 79, 9 (Sep 1991), 1305–1320. <https://doi.org/10.1109/5.97300>
- Robert H. Halstead, Jr. 1985. MULTILISP: A Language for Concurrent Symbolic Computation. *ACM Trans. Program. Lang. Syst.* 7, 4 (Oct. 1985), 501–538. <https://doi.org/10.1145/4472.4478>
- Jan Hoffmann, Ankush Das, and Shu-Chun Weng. 2017. Towards Automatic Resource Bound Analysis for OCaml. In *44th Symposium on Principles of Programming Languages (POPL'17)*.
- Jan Hoffmann and Zhong Shao. 2015. Automatic Static Cost Analysis for Parallel Programs. In *Proceedings of the 24th European Symposium on Programming on Programming Languages and Systems - Volume 9032*. Springer-Verlag New York, Inc., New York, NY, USA, 132–157. https://doi.org/10.1007/978-3-662-46669-8_6
- Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *7th European Symposium on Programming Languages and Systems (ESOP'98)*. Springer LNCS 1381, 122–138.

- Naoki Kobayashi. 2002. A Type System for Lock-Free Processes. *Information and Computation* 177 (2002), 122–159.
- Neelakantan R. Krishnaswami and Nick Benton. 2011. Ultrametric Semantics of Reactive Programs. In *26th IEEE Symposium on Logic in Computer Science, (LICS'11)*. 257–266.
- Ugo Dal Lago and Marco Gaboardi. 2011. Linear Dependent Types and Relative Completeness. In *26th IEEE Symp. on Logic in Computer Science (LICS'11)*. 133–142.
- Julien Lange and Nobuko Yoshida. 2017. On the Undecidability of Asynchronous Session Subtyping. In *Proceedings of the 20th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)*. Springer LNCS 10203, 441–457.
- Hugo A. López, Carlos Olarte, and Jorge A. Pérez. 2009. Towards a Unified Framework for Declarative Structure Communications. In *Proceedings of the Workshop on Programming Language Approaches to Concurrency and Communication-Centric Software (PLACES)*, A. Beresford and S. Gay (Eds.). EPTCS 17, 1–15.
- Hiroshi Nakano. 2000. A Modality for Recursion. In *15th IEEE Symposium on Logic in Computer Science (LICS'00)*. 255–266.
- Rumyana Neykova, Laura Bocchi, and Nobuko Yoshida. 2014. Timed Runtime Monitoring for Multiparty Conversations. In *3rd International Workshop on Behavioural Types (BEAT 2014)*.
- W. Paul, U. Vishkin, and H. Wagener. 1983. Parallel dictionaries on 2–3 trees. In *Automata, Languages and Programming*, Josep Diaz (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 597–609.
- Frank Pfenning and Dennis Griffith. 2015. Polarized Substructural Session Types. In *Proceedings of the 18th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2015)*, A. Pitts (Ed.). Springer LNCS 9034, London, England, 3–22. Invited talk.
- Amir Pnueli. 1977. The Temporal Logic of Programs. In *Proceedings of the 18th Symposium on Foundations of Computer Science (FOCS'77)*. IEEE Computer Society, 46–57.
- Marc Pouzet. 2006. Lucid Synchrone Release, version 3.0 Tutorial and Reference Manual. (2006).
- Neda Saeedloei and Gopal Gupta. 2014. Timed π -Calculus. In *8th International Symposium on Trustworthy Global Computing - Volume 8358 (TGC 2013)*. Springer-Verlag New York, Inc., New York, NY, USA, 119–135. https://doi.org/10.1007/978-3-319-05119-2_8
- Miguel Silva, Mário Florido, and Frank Pfenning. 2016. Non-Blocking Concurrent Imperative Programming with Session Types. In *Fourth International Workshop on Linearity*.
- Bernardo Toninho, Luís Caires, and Frank Pfenning. 2013. Higher-Order Processes, Functions, and Sessions: A Monadic Integration. In *Proceedings of the European Symposium on Programming (ESOP'13)*, M.Felleisen and P.Gardner (Eds.). Springer LNCS 7792, Rome, Italy, 350–369.
- Bernardo Toninho, Luís Caires, and Frank Pfenning. 2014. Corecursion and Non-Divergence in Session-Typed Processes. In *Proceedings of the 9th International Symposium on Trustworthy Global Computing (TGC 2014)*, M. Maffei and E. Tuosto (Eds.). Springer LNCS 8902, Rome, Italy, 159–175.
- Philip Wadler. 2012. Propositions as Sessions. In *Proceedings of the 17th International Conference on Functional Programming (ICFP 2012)*. ACM Press, Copenhagen, Denmark, 273–286.
- Ezgi Çiçek, Gilles Barthe, Marco Gaboardi, Deepak Garg, and Jan Hoffmann. 2017. Relational Cost Analysis. In *44th Symposium on Principles of Programming Languages (POPL'17)*.