

Work Analysis with Resource-Aware Session Types

Ankush Das
Carnegie Mellon University

Jan Hoffmann
Carnegie Mellon University

Frank Pfenning
Carnegie Mellon University

Abstract

While there exist several successful techniques for supporting programmers in deriving static resource bounds for sequential code, analyzing the resource usage of message-passing concurrent processes poses additional challenges. To meet these challenges, this article presents an analysis for statically deriving worst-case bounds on the total work performed by message-passing processes. To decompose interacting processes into components that can be analyzed in isolation, the analysis is based on novel resource-aware session types, which describe protocols and resource contracts for inter-process communication. A key innovation is that both messages and processes carry potential to share and amortize cost while communicating. To symbolically express resource usage in a setting without static data structures and intrinsic sizes, resource contracts describe bounds that are functions of interactions between processes. Resource-aware session types combine standard binary session types and type-based amortized resource analysis in a linear type system. This type system is formulated for a core session-type calculus of the language SILL and proved sound with respect to a multiset-based operational cost semantics that tracks the total number of messages that are exchanged in a system. The effectiveness of the analysis is demonstrated by analyzing standard examples from amortized analysis and the literature on session types and by a comparative performance analysis of different concurrent programs implementing the same interface.

1 Introduction

In the past years, there has been increasing interest in supporting developers to statically reason about the resource usage of their code. This research has numerous applications such as prevention of side channels leaking secret information [4, 23, 26], identification of complexity bugs [27], support of scheduling decisions [1], and help in profiling [16]. While there has been great progress in analyzing sequential code, relatively little research has been done on analyzing the resource consumption of concurrent and distributed programs [2, 3, 13]. The lack of analysis tools is in sharp contrast to the need of programming language support in this area: concurrent and distributed programs are increasingly pervasive and particularly difficult to analyze. For shared memory concurrency, we need to precisely predict scheduling to account for synchronization cost. Similarly, the interactive nature of message-passing systems makes it difficult to decompose the system into components that can be analyzed in isolation. After all, the resource usage of each component crucially depends on its interactions.

In this paper, we study the foundations of worst-case resource analysis for message-passing processes. A key idea of our approach is to rely on *resource-aware session types* to describe the structure, protocols and resource bounds for inter-process communication and to perform a compositional and precise amortized analysis. *Session types* [22] prescribe bidirectional communication protocols for message-passing processes. *Binary session types* govern the

interaction of two processes along a single channel, prescribing complementary send and receive actions for the processes at the two endpoints of a channel. Recently, message-passing concurrency has been put onto a firm logical foundation by exhibiting a Curry-Howard isomorphism between intuitionistic linear logic and session-typed communication [8]. We use such protocols as the basis of resource usage contracts that not only specify the type, but also the potential of a message that is sent along a channel. The potential (in the sense of classical amortized analysis [31]) may be spent by sending other messages as part of the network of interacting processes, or maintained locally for future interactions. Resource analysis is static, using the type system, and the runtime behavior of programs is not affected.

We focus on bounds on the total work performed by a system, counting the number of messages that are exchanged. While this alone does not yet account for the concurrent nature of message-passing programs, it constitutes a significant and necessary first step. The bounds we derive are also useful in their own right. For example, the information can be used in scheduling decisions, to derive the number of messages that are sent along a specific channel, or to statically decide whether we should spawn a new thread of control or execute sequentially when possible. Additionally, bounds on the work of a process also serve as input to a Brent-style theorem [7] that relates the complexity of the execution of a program on a k -processor machine to the program's work (the focus of this paper) and span (the resource usage if we assume an unlimited number of processors). We are working on a companion paper for deriving bounds on the span, which is both conceptually and technically quite different.

Our analysis is based on a linear type system that combines standard binary session types as available in the SILL language [28, 33], and type-based amortized resource analysis [17, 19]. Both techniques are based on linear or affine type systems, making their combination natural. Each session type constructor is decorated with a natural number that declares a potential that must be transferred (conceptually!) along with the corresponding message. Since the interface to a process is characterized entirely by the resource-aware session types of the channels it interacts with, this design provides a compositional resource specification. For closed programs (which evolve into a closed network of interacting processes) the bound becomes a single constant. In addition to the natural compositionality of type systems we also preserve the necessary support for deriving resource annotations via LP solving, a key feature of type-based amortized analysis. While we have not yet implemented a type inference algorithm, we designed the system to support inference. Moreover, resource-aware session types integrate well with type-based amortized analysis for functional programs because they are based on compatible logical foundations (intuitionistic linear logic and intuitionistic logic, respectively), as exemplified in the design of SILL [28, 33] that combines them monadically.

A conceptual challenge is to express symbolic bounds in a setting without static data structures and intrinsic sizes. Our innovation is that resource-aware session types describe bounds as functions

of interactions (messages sent) on a channel. A major technical challenge is to account for the global number of messages sent with local derivation rules: operationally, local message counts are forwarded to a parent process when a sub-process terminates. As a result, local message counts are incremented by sub-processes in a non-local fashion. Our solution is that messages and processes carry potential to amortize the cost of a terminating sub-process proactively as a side-effect of the communication.

Our main contributions are as follows. We present the first session type system for deriving parametric bounds on the resource usage of message-passing processes. We prove the nontrivial soundness of the type system with respect to an operational cost semantics that tracks the total number of messages exchanged in a network of communicating processes. We demonstrate the effectiveness of the technique by deriving tight bounds for some standard examples of amortized analysis from the literature on session types. We also show how resource-aware session types can be used to specify and compare the performance characteristics of different implementations of the same protocol. The analysis is currently manual, with automation left for future work, as is a companion type system for deriving the *span* of concurrent computations.

2 Overview

In this section, we motivate and informally introduce our resource-aware session types and show how they can be used to analyze the resource usage of message-passing processes. We start with building some intuition about session types.

Session Types. Session types were introduced by Honda [22] to describe the structure of communication just like standard data types describe the structure of data. We follow the approach and syntax of SILL [28, 33] which is based on a Curry-Howard isomorphism between intuitionistic linear logic and session types, extended by recursively defined types and processes. In the intuitionistic approach, every channel has a *provider* and a *client*. We view the session type as describing the communication from the provider’s point of view, with the client having to perform matching actions.

As a first simple example, we consider natural numbers in binary form. A process *providing* a natural number sends a stream of bits starting with the least significant bit. These bits are represented by messages zero and one, eventually terminated by \$. Because the provider chooses which messages to send, we call this an *internal choice*, which is written as

$$\text{bits} = \oplus\{\text{zero} : \text{bits}, \text{one} : \text{bits}, \$: \mathbf{1}\}$$

Here, $\oplus\{l_1 : A_1, \dots, l_n : A_n\}$ is an n-ary, labelled generalization of $A \oplus B$ of linear logic, and $\mathbf{1}$ is the multiplicative unit of linear logic. Operationally, $\mathbf{1}$ means the provider has to send an *end* message, closing the channel and terminating the communication session. For example, the number $6 = (110)_2$ would be represented by the sequence of messages zero, one, one, \$, *end*. A client of a channel $c : \text{bits}$ has to branch on whether it receives zero, one, or \$. Note that as we proceed in a session, the type of a channel must change according to the protocol. For example, if a client receives the message \$ along $c : \text{bits}$ then c must afterwards have type $\mathbf{1}$. The next message along c must be *end* and its client has to wait for that after receiving \$ so the session is properly closed.

As a second example, we describe the interface to a counter. As a client, we can repeatedly send *inc* messages to a counter, until we want to read its value and send *val*. At that point the counter




Figure 1. Binary counter system representing $5 = (101)_2$ with messages triggered when *inc* message is received on s_4 .

will send a stream of bits representing its value as prescribed by the type bits. From the provider’s point of view, a counter presents an *external choice*, since the client chooses between *inc* or *val*.

$$\text{ctr} = \&\{\text{inc} : \text{ctr}, \text{val} : \text{bits}\}$$

The type former $\&\{l_1 : A_1, \dots, l_n : A_n\}$ is an n-ary labelled generalization of $A \& B$ of linear logic. Operationally, the provider must branch based on which of the labels l_i it receives. After receiving l_k along a channel $c : \&\{l_1 : A_1, \dots, l_n : A_n\}$, communication along c proceeds at type A_k . Such type formers can be arbitrarily nested to allow more complex bidirectional protocols.

Modeling a binary counter. We describe an implementation of a counter and use our resource-aware session types to analyze its resource usage. Like in the rest of the paper, the resource we are interested in is the total number of messages sent along all channels in the system.

A well-known example of amortized analysis counts the number of bits that must be flipped to increment a counter. It turns out the amortized cost per increment is 2, so that n increments require at most $2n$ bits to be flipped. We observe this by introducing a potential of 1 for every bit that is 1 and using this potential to *pay* for the expensive case in which an increment triggers many flips. When the lowest bit is zero, we flip it to one (costing 1) and also store a remaining potential of 1 with this bit. When the lowest bit is one we use the stored potential to flip the bit back to zero (with no stored potential) and the remaining potential of 2 is passed along for incrementing the higher bits.

We model a binary counter as a chain of processes where each process represents a single bit (process b_0 or b_1) with a final process e at the end. Each of the processes in the chain *provides* a channel of the *ctr* type, and each (except the last) also *uses* a channel of this type representing the higher bits. For example, in the first chain in Figure 1, the process b_0 offers along channel s_3 (indicated by • between b_0 and s_3) and uses channel s_2 . In our notation, we would write this as

$$\begin{array}{l} \cdot \vdash e :: (s_1 : \text{ctr}) \quad s_1 : \text{ctr} \vdash b_1 :: (s_2 : \text{ctr}) \\ s_2 : \text{ctr} \vdash b_0 :: (s_3 : \text{ctr}) \quad s_3 : \text{ctr} \vdash b_1 :: (s_4 : \text{ctr}) \end{array}$$

We see that, logically, parallel composition with a private shared channel corresponds to an application of the cut rule. We do not show here the *definitions* of e , b_0 , and b_1 , which can be found in Figure 3. The only channel visible to an outside client (not shown) is s_4 . Figure 1 shows the messages triggered if an increment message is received along s_4 .

Expressing resource bounds. Our basic approach is that *messages carry potential* and *processes store potential*. This means the sender has to pay not just 1 unit for sending the message, but whatever additional units to amortize future costs. In the amortized analysis

of the counter, each bit flip corresponds exactly to an `inc` message, because that is what triggers a bit to be flipped. Our cost model focuses on messages as prescribed by the session type and does not count other operations, such as spawning a new process or terminating a process. This choice is not essential to our approach, but convenient here.

To capture the informal analysis we need to express *in the type* that we have to send 1 unit of potential with the label `inc`. We do this using a superscript indicating the required potential with the label, postponing the discussion of `val`.

$$\text{ctr} = \&\{\text{inc}^1 : \text{ctr}, \text{val}^? : \text{bits}\}$$

When we assign types to the processes, we now use these more expressive types. We also indicate the potential stored in a particular process as a superscript on theturnstile.

$$t : \text{ctr} \stackrel{0}{\vdash} b0 :: (s : \text{ctr}) \quad (1)$$

$$t : \text{ctr} \stackrel{1}{\vdash} b1 :: (s : \text{ctr}) \quad (2)$$

$$\cdot \stackrel{0}{\vdash} e :: (s : \text{ctr}) \quad (3)$$

With our formal typing rules (see Section 5) we can verify these typing constraints, using the definitions of $b0$, $b1$, and e . Informally, we can reason as follows:

- $b0$: When $b0$ receives `inc` it receives 1 unit of potential. It continues as $b1$ (which requires no communication) which stores this 1 unit (as prescribed from the type of $b1$ in Equation 2).
- $b1$: When $b1$ receives `inc` it receives 1 unit of potential which, when combined with the stored one, makes 2 units. It needs to send an `inc` message which consumes these 2 units (1 to send the message, and 1 to send along a potential of 1 as prescribed in the type). It has no remaining potential, which is sufficient because it transitions to $b0$ which stores no potential (inferred from the type of $b0$ in Equation 1).
- e : When e receives `inc` it receives 1 unit of potential. It spawns a new process e and continues as $b1$. Spawning a process is free, and e requires no potential, so it can store the potential it received with $b1$ as required.

How do we handle the type annotation $\text{val}^? : \text{bits}$ of the label `val`? Recall that $\text{bits} = \oplus\{\text{zero} : \text{bits}, \text{one} : \text{bits}, \$: 1\}$. In our implementation, upon receiving a `val` message, a $b0$ or $b1$ process will first respond with zero or one respectively. It then sends `val` along the channel it uses (representing the higher bits of the number) and terminates by *forwarding* further communication to the higher bits in the chain. Figure 2 demonstrates the messages triggered when `val` message is received along s_4 . The e process will just send `$` and `end`, indicating the empty stream of bits.

We know we will have enough potential to carry out the required send operations if each process ($b0$, $b1$, and e) carries an additional 2 units of potential. We could impart these with the `inc` and `val` messages by sending 2 more units with `inc` and 2 units with `val`. That is, the following type is correct:

$$\text{bits} = \oplus\{\text{zero}^0 : \text{bits}, \text{one}^0 : \text{bits}, \$^0 : 1^0\}$$

$$\text{ctr} = \&\{\text{inc}^3 : \text{ctr}, \text{val}^2 : \text{bits}\}$$

Here, the superscript 0 in the type of bits indicates that the corresponding messages carry no potential.

However, this type is a gross over-approximation! The processes of a counter of value n , would carry $2n$ additional potential while only $2 \lceil \log(n+1) \rceil + 2$ are needed. To obtain this more precise bound, we need *families of session types*.




Figure 2. Binary counter system representing $5 = (101)_2$ with messages triggered when `val` message is received on s_4 .

A more precise analysis. This requires that *in the type* we can refer either to the number of bits in the representation of a number or its value. This form of internal measure is needed only for type-checking purposes, not at runtime. It is also not intrinsically tied to a property of a representation, the way the length of a list in a functional language is tied to its memory requirements. We indicate these measures using square brackets, so that $\text{ctr}[n]$ is a family of types, and $\text{ctr}[0]$, for example, is a counter with value 0. Such type refinements have been considered in the literature on session types (see [15]) with respect to type-checking and inference. Here, we treat it as a meta-level notation to denote families of types. Following the reasoning above, we obtain the following type:

$$\text{bits} = \oplus\{\text{zero}^0 : \text{bits}, \text{one}^0 : \text{bits}, \$^0 : 1^0\}$$

$$\text{ctr}[n] = \&\{\text{inc}^1 : \text{ctr}[n+1], \text{val}^{2 \lceil \log(n+1) \rceil + 2} : \text{bits}\}$$

To check the types of our implementation, we need to revisit and refine the typing of the $b0$, $b1$ and e processes.

$$\begin{aligned} t : \text{ctr}[n] &\stackrel{0}{\vdash} b0 :: (s : \text{ctr}[2n]) \\ t : \text{ctr}[n] &\stackrel{1}{\vdash} b1 :: (s : \text{ctr}[2n+1]) \\ \cdot &\stackrel{0}{\vdash} e :: (s : \text{ctr}[0]) \end{aligned}$$

Our type system verifies these types against the implementation of $b0$, $b1$, and e (see Figure 3). The typing rules reduce the well-typedness of these processes to arithmetic inequalities which we can solve by hand, for example, using that $\log(2n) = \log(n) + 1$. The intrinsic measure n and the precise potential annotations are not automatically derived, but come from our insight about the nature of the algorithms and programs.

Before introducing the formalism in which the programs are expressed, together with the typing rules that let us perform rigorous amortized analysis of the code (as expressed in the soundness theorem in Section 6), we again emphasize the *compositional nature* of the way resource bounds are expressed in the types themselves and in the typing judgments for process expressions. Of course, they reveal some intensional property of the implementations, namely a bound on its cost, so different implementations of the same plain session type may have different resource annotations.

The typing derivation provides a proof certificate on the resource bound for a process. For closed processes typed as

$$\cdot \stackrel{p}{\vdash} Q :: (c : 1^0)$$

the number p provides a worst case bound on the number of messages sent during computation of Q , which always ends with the process sending `end` along c , indicating termination.

3 Resource-Aware SILL

We briefly introduce the linear, process-only fragment of SILL [28, 33], which integrates functional and concurrent computation. A program in SILL is a collection of processes exchanging messages

Type (current)	Continuation	Process Term (current)	Continuation	Description
$c : \oplus \{l_i^{q_i} : S_i\}$	$c : S_k$	$c.l_k ; P$ case $c (l_i \Rightarrow Q_i)_{i \in I}$	P Q_k	provider sends label l_k along c with potential q_k client receives label l_k along c with potential q_k
$c : \& \{l_i^{q_i} : S_i\}$	$c : S_k$	case $c (l_i \Rightarrow P_i)_{i \in I}$ $c.l_k ; Q$	P_k Q	provider receives label l_k along c with potential q_k client sends label l_k along c with potential q_k
$c : S \otimes T$	$c : T$	send $c w ; P$ $y \leftarrow \text{rcv } c ; Q_y$	P $[w/y]Q_y$	provider sends channel $w : S$ along c with potential q client receives channel $w : S$ along c with potential q
$c : S \multimap T$	$c : T$	$y \leftarrow \text{rcv } c ; P_y$ send $c w ; Q$	$[w/y]P_y$ Q	provider receives channel $w : S$ along c with potential q client sends channel $w : S$ along c with potential q
$c : 1^q$	–	close c wait $c ; Q$	– Q	provider sends <i>end</i> along c with potential q client receives <i>end</i> along c with potential q

Table 1. Linear resource-aware session types

through channels. A new process is *spawned* by invoking a process definition, which also creates a fresh channel. The newly spawned process acts as a *provider* of the fresh channel while the parent process acts as its *client*. The exacting nature of linear typing provides strong guarantees, including session fidelity (a form of preservation) and absence of deadlocks (a form of progress).

We present an overview of the session types in SILL along with a brief description of their communication protocol. They follow the type grammar below.

$$S, T ::= V \mid \oplus \{l_i : S\} \mid \& \{l_i : S\} \mid S \multimap T \mid S \otimes T \mid 1$$

V denotes a type variable. Types may be defined mutually recursively in a global signature, where type definitions are constrained to be *contractive* [12]. This allows us to treat them equi-recursively, meaning we can silently replace a type variable by its definition for type-checking purposes.

Internal choice $S \oplus T$ and external choice $S \& T$ have been generalized to n -ary labeled sums $\oplus \{l_i : S_i\}_{i \in I}$ and $\& \{l_i : S_i\}_{i \in I}$ (for some index set I) respectively. As a provider of internal choice $\oplus \{l_i : S_i\}_{i \in I}$, a process can send one of the labels l_i to its client. As a dual, a provider of external choice $\& \{l_i : S_i\}_{i \in I}$ receives one of the labels l_i which is sent by its client. We require external and internal choice to comprise at least one label, otherwise there would exist a linear channel without observable communication along it, which is computationally uninteresting and would complicate our proofs. The connectives \otimes and \multimap are used to send channels via other channels. A provider of $S \otimes T$ sends a channel of type S to its client and then behaves as a provider of T . Dually, a provider of $S \multimap T$ receives a channel of type S from its client. The types of the provider and client change consistently, and the process terms of a provider and client come in matching pairs.

Formally, the syntax of process expressions of Resource-Aware SILL is same as in SILL.

$$P ::= x \leftarrow X \leftarrow \bar{y} ; P \mid x \leftarrow y \mid \text{close } x \mid \text{wait } x ; P \\ \mid x.l_k ; P \mid \text{case } x (l_i \Rightarrow P_i)_{i \in I} \\ \mid \text{send } x w ; P \mid y \leftarrow \text{rcv } x ; P$$

The term $x \leftarrow X \leftarrow \bar{y} ; P$ invokes a process definition X to spawn a new process Q , which uses the channels in \bar{y} as a client and provides service along a fresh channel substituted for x in P . A forwarding process $x \leftarrow y$ (which provides channel x) identifies channels x and y and terminates. The effect is that clients of x will afterwards communicate along y (used in Figure 2). The rest of the program constructs concern communication between two processes and are guided by their corresponding session type. Table 1 provides

an overview of session types, associated process terms, and their operational description (ignore the annotations in red). For each connective in Table 1, the first row presents the perspective of the provider, while the second presents that of the client. The first two columns present the type of the channel before (**current**) and after (**continuation**) the interaction. Similarly, the next two columns present the process terms before and after the interaction. The last column provides the operational description.

We conclude by illustrating the syntax, types and semantics of SILL using a simple example. Recall the counter protocol (ignoring the resource annotations in red):

$$\text{bits} = \oplus \{\text{zero}^0 : \text{bits}, \text{one}^0 : \text{bits}, \$^0 : 1^0\} \\ \text{ctr}[n] = \& \{\text{inc}^1 : \text{ctr}[n+1], \text{val}^{2^{\lceil \log(n+1) \rceil + 2}} : \text{bits}\}$$

Figure 3 presents implementations of the b_0 , b_1 and e processes respectively that were analyzed in Section 2. The comments (starting with %) show the channel types after the interaction on each line (again ignoring the annotations in red). Since the b_0 process provides an external choice along s , b_0 needs to branch based on the label received (line 4). If it receives the label inc , the type of the channel s updates to ctr , as indicated on the typing in the comment. At this point, we spawn the b_1 process whose type (line 8) matches with the type on line 4. If instead b_0 receives the val label along s , it continues at type bits . It sends zero (since the lowest bit is indeed zero) and requests the value of the higher bits by sending val along channel t . Now both s and t have type bits (indicated in the typing on line 7) and b_0 can terminate by forwarding further communication along t to s .

The b_1 process operates similarly, taking care to handle the carry upon increment by sending an inc label along t . The e process spawns a new e process and continues as b_1 upon receiving the label inc and closes the channel after sending $\$$ when receiving val .

4 Cost Semantics

We present an operational cost semantics for Resource-Aware SILL tracking the work performed by the system. Our semantics is a substructural operational semantics [29] based on multiset rewriting [10] and asynchronous communication [28]. It can be seen as a combination of an asynchronous version of a recent synchronous session-type semantics [6] with the cost tracking semantics of Concurrent C0 [30]. The technical advantage of our semantics is that it avoids the complex operational artifacts of Silva et al. [30] such as message buffers: processes and messages can be typed with exactly the same typing rules, changing only the cost metric.


```

1: (t : ctr[n]) ⊢ b0 :: (s : ctr[2n])
2:   s ← b0 ← t =
3:   case s
4:     (inc ⇒ s ← b1 ← t           % (t : ctr[n]) ⊢ s : ctr[2n + 1]
5:     | val ⇒ s.zero ;               % (t : ctr[n]) ⊢ s : bits
6:     t.val ;                         % (t : bits) ⊢ s : bits
7:     s ← t)                          % (t : bits) ⊢ s : bits
8: (t : ctr[n]) ⊢ b1 :: (s : ctr[2n + 1])
9:   s ← b1 ← t =
10:  case s
11:    (inc ⇒ t.inc ;                 % (t : ctr[n + 1]) ⊢ s : ctr[2n + 2]
12:    s ← b0 ← t
13:    | val ⇒ s.one ;                % (t : ctr[n]) ⊢ s : bits
14:    t.val ;                         % (t : bits) ⊢ s : bits
15:    s ← t)                          % (t : bits) ⊢ s : bits
16: · ⊢ e :: (s : ctr[0])
17:   s ← e =
18:   case s
19:     (inc ⇒ t ← e ;                 % (t : ctr[0]) ⊢ (s : ctr[1])
20:     s ← b1 ← t
21:     | val ⇒ s.e ;                  % · ⊢ s : 1
22:     close s)

```

Figure 3. Implementations for the b_0 , b_1 and e processes with their type derivations demonstrating the binary counter.

Our cost semantics is asynchronous, that is, processes can continue their evaluation without waiting after sending a message. In order to guarantee session fidelity, the semantics must ensure that messages are received in the order they are sent. Intuitively, we can think of channels as FIFO message buffers, although we will formally define them differently. Synchronous communication can be implemented in our language in a type-safe, logically motivated manner exploiting adjoint shift operators (see [28]).

A collection of communicating processes is called a *configuration*. A configuration is formally modelled as a multiset of propositions $\text{proc}(c, w, P)$ and $\text{msg}(c, w, M)$. The predicate $\text{proc}(c, w, P)$ describes a process executing process expression P and providing channel c . The predicate $\text{msg}(c, w, M)$ describes the message M on channel c . In order to guarantee that messages are received in they order they are sent, only a *single message* can be on a given channel c . In order for computation to remain truly asynchronous, every send operation (except for close) on a channel c creates not only a fresh message, but also a fresh continuation channel c' for the next message. This continuation channel is encoded within the message via a forwarding operation. Remarkably, this simple device allows us to assign session types to messages just as if they were processes! Since M need only encode a message, it has a restricted grammar.

$$M ::= c \leftarrow c' \mid c.l_k ; c \leftarrow c' \mid c.l_k ; c' \leftarrow c \mid \text{send } c \ e ; c \leftarrow c' \mid \text{send } c \ e ; c' \leftarrow c \mid \text{close } c$$

The work is tracked by the local counter w in $\text{proc}(c, w, P)$ and $\text{msg}(c, w, M)$ propositions. For process P , w maintains the total work performed by P so far. When a process sends a message (i.e. creates a new msg predicate), we increment its counter w by the cost for sending. When a process terminates we remove the respective predicate from the configuration, but preserve its work done. A process can terminate either by sending a close message, or by forwarding. In either case, we conveniently preserve the process' work in the msg predicate to pass it on to the client process.

We will only count communication costs, ignoring internal computation. To this end, we introduce 3 costs, M^{label} , M^{channel} and M^{close} , for labels, channels, and close messages, respectively. A concrete semantics can be obtained by setting appropriate values for each of those metrics. For instance, setting each cost to 1 will lead to counting the total number of messages exchanged.

The semantics is defined by a set of rules rewriting the configuration that *consume* the proposition in the premise of the rule and *produce* the propositions in the conclusion (rules should be read top-down!). A step consists of non-deterministic application of a rule whose premises match a part of the configuration. Consider for instance the rule $\&C_s$ that describes a client that sends label l_k along channel c .

$$\frac{\text{proc}(d, w, c.l_k ; P) \quad (c' \text{ fresh})}{\text{proc}(d, w + M^{\text{label}}, [c'/c]P) \quad \text{msg}(c', 0, c.l_k ; c' \leftarrow c)}$$

The above rule can be applied to every proposition of the form $\text{proc}(d, w, c.l_k ; P)$. When applying the rule, we generate a fresh continuation channel c' and replace the premise by $\text{proc}(d, w + M^{\text{label}}, [c'/c]P)$ and $\text{msg}(c', 0, c.l_k ; c' \leftarrow c)$ propositions. The message predicate contains the process $c.l_k ; c' \leftarrow c$ which will eventually deliver the message to the provider along c and will continue communication along c' (which is achieved by $c' \leftarrow c$). The work of the sender is incremented by M^{label} to account for the sent message, while the work of the message is 0.

Conversely, the rule $\&C_r$ defines how a provider receives a label l_k along c .

$$\frac{\text{msg}(c', w, c.l_k ; c' \leftarrow c) \quad \text{proc}(c, w', \text{case } c (l_i \Rightarrow Q_i)_{i \in I})}{\text{proc}(c, w + w', [c'/c]Q_k)}$$

The rule replaces the msg and proc propositions in the configuration that match the premises, with the single proc proposition in the conclusion. Since the provider receives the label l_k , it continues as Q_k . However, we replace c with c' in Q_k since the forwarding $c' \leftarrow c$ in the message process tells us that the next message will arrive on channel c' . Any work w encoded in the message is transferred to the recipient.

The rest of the rules of cost semantics are given in Figure 4. The rule spawn_c describes the creation of a new channel c along with spawning a new process X implemented by P_c . This implementation is looked up in a signature for the semantics Σ which maps process names to the implementation code. The new process is spawned with 0 work (as it has not sent any messages so far), while Q_c continues with the same amount of work. In the rule fwd_s , a forwarding process creates a *forwarding message* and terminates. The work carried by this special message is the same as the work done by the process, now defunct. A forwarding message form does not carry any real information (except for the work w !); it just serves to identify the two channels c and d . In an implementation this could be as simple as concatenating two message buffers. We therefore do not count forwarding messages when computing the work. Another reason forward messages are special is that unlike all other forms of messages, they are neither prescribed by nor manifest in a channel's type. In our formal rules, the forwarding message can be absorbed either into the client (fwd_r^+) or provider (fwd_r^-), in both cases preserving the total amount of work.

The rules of the cost semantics are successively applied to a configuration until the configuration becomes empty or the configuration is stuck and none of the rules can be applied. At any point

$$\begin{array}{c}
\frac{\Sigma(\mathcal{X}) = x \leftarrow \mathcal{X} \leftarrow \bar{y} = P_{x,\bar{y}} \quad \text{proc}(d, w, x \leftarrow \mathcal{X} \leftarrow \bar{e}; Q_x)}{\text{proc}(c, 0, [c/x, \bar{e}/y]P_{x,\bar{y}}) \quad \text{proc}(d, w, [c/x]Q_x)} \text{spawn}_c \quad \frac{\text{proc}(c, w, c \leftarrow d)}{\text{msg}(c, w, c \leftarrow d)} \text{fwd}_s \quad \frac{\text{proc}(c, w, \text{close } c)}{\text{msg}(c, w + M^{\text{close}}, \text{close } c)} 1C_s \\
\frac{\text{proc}(e, w, P_c) \quad \text{msg}(c, w', c \leftarrow d)}{\text{proc}(e, w + w', [d/c]P_c)} \text{fwd}_r^- \quad \frac{\text{proc}(d, w, P) \quad \text{msg}(c, w', c \leftarrow d)}{\text{proc}(c, w + w', [c/d]P)} \text{fwd}_r^+ \quad \frac{\text{msg}(c, w, \text{close } c) \quad \text{proc}(d, w', \text{wait } c; Q)}{\text{proc}(d, w + w', Q)} 1C_r \\
\frac{\text{proc}(c, w, c.l_k; P) \quad (c' \text{ fresh})}{\text{proc}(c', w + M^{\text{label}}, [c'/c]P) \quad \text{msg}(c, 0, c.l_k; c \leftarrow c')} \oplus_{C_s} \quad \frac{\text{msg}(c, w, c.l_k; c \leftarrow c') \quad \text{proc}(d, w', \text{case } c (l_i \Rightarrow Q_i)_{i \in I})}{\text{proc}(d, w + w', [c'/c]Q_k)} \oplus_{C_r} \\
\frac{\text{proc}(c, w, \text{send } c e; P) \quad (c' \text{ fresh})}{\text{proc}(c', w + M^{\text{channel}}, [c'/c]P) \quad \text{msg}(c, 0, \text{send } c e; c \leftarrow c')} \otimes_{C_s} \quad \frac{\text{msg}(c, w, \text{send } c e; c \leftarrow c') \quad \text{proc}(d, w', x \leftarrow \text{recv } c; Q_x)}{\text{proc}(d, w + w', [c'/c]Q_e)} \otimes_{C_r} \\
\frac{\text{proc}(c', w + M^{\text{channel}}, [c'/c]P) \quad \text{msg}(c, 0, \text{send } c e; c \leftarrow c')}{\text{proc}(d, w, \text{send } c e; P) \quad (c' \text{ fresh})} \multimap_{C_s} \quad \frac{\text{msg}(c', w, \text{send } c e; c' \leftarrow c) \quad \text{proc}(c, w', x \leftarrow \text{recv } c; Q_x)}{\text{proc}(c, w + w', [c'/c]Q_e)} \multimap_{C_r} \\
\text{proc}(d, w + M^{\text{channel}}, [c'/c]P) \quad \text{msg}(c', 0, \text{send } c e; c' \leftarrow c) \quad \text{proc}(c, w + w', [c'/c]Q_e) \quad \multimap_{C_s} \quad \multimap_{C_r}
\end{array}$$

Figure 4. Cost semantics tracking total work for programs in SILL

in this local stepping, the total work performed by the system can be obtained by summing the local counters w for each predicate in the configuration. We will prove in Section 6 that this total work can be upper bounded by the initial potential of the configuration that is typed in our resource-aware type system.

5 Type System

We present the resource-aware type system of our language which extends the linear-only fragment of SILL [28, 33] with resource annotations. It is in turn based on intuitionistic linear logic [14] with sequents of the form

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

where A_1, \dots, A_n are the linear antecedents and A is the succedent. Under the Curry-Howard isomorphism for intuitionistic linear logic, propositions are related to session types, proofs to processes, and cut reduction in proofs to communication. Appealing to this correspondence, we assign a process term P to the above judgment and label each hypothesis as well as the conclusion with a channel.

$$(x_1 : A_1), (x_2 : A_2), \dots, (x_n : A_n) \vdash P :: (x : A)$$

The resulting judgment states that process P provides a service of session type A along channel x , using the services of session types A_1, \dots, A_n provided along channels x_1, \dots, x_n respectively. The assignment of a channel to the conclusion is convenient because, unlike functions, processes do not evaluate to a value but continue to communicate along their providing channel once they have been created. For the judgment to be well-formed, all the channel names have to be distinct. Whether a session type is used or provided is determined by its positioning to the left or right, respectively, of the turnstile.

Resource-aware session types are obtained by annotating simple session types with potential. They are defined by the following grammar.

$$S, T ::= V \mid \oplus\{l_i^{q_i} : S\} \mid \&\{l_i^{q_i} : S\} \mid S \multimap^q T \mid S \otimes^q T \mid 1^q$$

V is a type variable. The meaning of the types and the process terms associated with it are defined in Table 1 (annotations and descriptions pertaining to potentials are marked in red).

The typing judgment of Resource-Aware SILL has the form

$$\Sigma; \Omega \stackrel{q}{\vdash} P :: (x : S)$$

Intuitively, the judgment describes a process in state P using the context Ω and signature Σ and providing service along channel x of type S . In other words, P is the provider for channel $x : S$, and a client for all the channels in Ω . The resource annotation q is a natural number and defines the potential stored in the process P .

When reasoning about the work performed by a system, we reason parametrically in certain quantities, such as the value of a counter, the number of elements in a queue, the potential carried by a message, or even the type of the elements in a queue. In an implementation, we would have to make type families, index domains, constraint solving, etc. explicit, but fortunately we can avoid the notational overhead that this entails. This is because the types and rules are always *schematic* in their parameters and quantification over these parameters can remain entirely at the metalevel. We model this by allowing (conceptually) infinite signatures with all instances of parametric definitions. In this way, when we reason parametrically we can be assured that any instance of what we derive is indeed a valid judgment. This allows us to focus on the key conceptual and technical contributions of our approach.

Σ defines this signature containing type and process definitions. It is defined as a possibly infinite set of type definitions $V = S_V$ and process definitions $x : S \leftarrow \mathcal{X} @ q \leftarrow \bar{y} : \bar{W} = P_{x,\bar{y}}$. The equation $V = S_V$ is used to define the type variable V as S_V . We treat such definitions *equirecursively*. For instance, $\text{ctr}[n] = \&\{\text{inc}^1 : \text{ctr}[n+1], \text{val}^{2^{\lceil \log(n+1) \rceil + 2}} : \text{bits}\}$ exists in the signature for all $n \in \mathbb{N}$ for the binary counter system. The process definition $x : S \leftarrow \mathcal{X} @ q \leftarrow \bar{y} : \bar{W} = P_{x,\bar{y}}$ defines a (possibly recursive) process named \mathcal{X} implemented by $P_{x,\bar{y}}$ providing along channel $x : S$ and using the channels $\bar{y} : \bar{W}$ as a client. The process also stores a potential q , shown as $\mathcal{X} @ q$ in the definition. For instance, $s : \text{ctr}[2n] \leftarrow b0 @ 0 \leftarrow t : \text{ctr}[n] = P_{s,t}$ ($P_{s,t}$ defines the implementation of $b0$) exists in the signature for all $n \in \mathbb{N}$.

Messages are typed differently from processes as their work counters w (introduced in the predicate $\text{msg}(c, w, M)$) are not incremented when they actually deliver the message to the receiver. Hence, to type the messages, we define an auxiliary cost-free typing judgment, $\Sigma; \Omega \stackrel{q}{\vdash} P :: (x : S)$, which follows the same typing rules

as Figure 5, but with $M^{\text{label}} = M^{\text{channel}} = M^{\text{close}} = 0$. This avoids paying the cost for sending a message twice. A fresh signature Σ is used in the derivation of the cost-free judgment.

The principle behind the type system is that each message carries potential and the sending process pays the potential along with the cost of sending a message from its local potential. The receiving process receives the potential when it receives the message and adds it to its local potential. For example, consider the rule $\&L_k$ for a client sending a label l_k along channel x .

$$\frac{q \geq p + r_k + M^{\text{label}} \quad \Sigma; \Omega (x : S_k) \stackrel{p}{\vdash} Q :: (z : U)}{\Sigma; \Omega (x : \&\{l_i^{r_i} : S_i\}) \stackrel{q}{\vdash} x.l_k; Q :: (z : U)} \&L_k$$

$$\begin{array}{c}
\frac{q \geq p + r_k + M^{\text{label}} \quad \Sigma; \Omega \not\vdash P :: (x : S_k) \quad (k \in I)}{\Sigma; \Omega \not\vdash (x.l_k; P) :: (x : \oplus\{l_i^{r_i} : S_i\}_{i \in I})} \oplus R_k \\
\frac{q + r_i \geq q_i \quad \Sigma; \Omega (x : S_i) \not\vdash Q_i :: (z : U) \quad (\forall i \in I)}{\Sigma; \Omega (x : \oplus\{l_i^{r_i} : S_i\}_{i \in I}) \not\vdash \text{case } x (l_i \Rightarrow Q_i)_{i \in I} :: (z : U)} \oplus L \\
\frac{q + r \geq p \quad \Sigma; \Omega (y : S) \not\vdash P_y :: (x : T)}{\Sigma; \Omega \not\vdash (y \leftarrow \text{recv } x; P_y) :: (x : S \xrightarrow{r} T)} \multimap R \\
\frac{q \geq p + r + M^{\text{channel}} \quad \Sigma; \Omega (x : T) \not\vdash Q :: (z : U)}{\Sigma; \Omega (w : S) (x : S \xrightarrow{r} T) \not\vdash (\text{send } x \ w; Q) :: (z : U)} \multimap L \\
\frac{q \geq p + r + M^{\text{channel}} \quad \Sigma; \Omega \not\vdash P :: (x : T)}{\Sigma; (w : S) \Omega \not\vdash \text{send } x \ w; P :: (x : S \otimes T)} \otimes R \\
\frac{q + r \geq p \quad \Sigma; \Omega (y : S) (x : T) \not\vdash Q_y :: (z : U)}{\Sigma; \Omega (x : S \otimes T) \not\vdash y \leftarrow \text{recv } x; Q_y :: (z : U)} \otimes L \\
\frac{q \geq r + M^{\text{close}}}{\Sigma; \cdot \not\vdash \text{close } x :: (x : 1^r)} 1R \\
\frac{q + r \geq p \quad \Sigma; \Omega \not\vdash Q :: (z : U)}{\Sigma; \Omega (x : 1^r) \not\vdash \text{wait } x; Q :: (z : U)} 1L
\end{array}$$

Figure 5. Typing rules for session-typed programs

Since the continuation Q is typed with potential p , and the potential sent with the label is r_k , the total original potential need be at least $p + r_k + M^{\text{label}}$. Thus, we get the constraint $q \geq p + r_k + M^{\text{label}}$.

The rule $\&R$ describes a provider that is awaiting a message on channel x and has local potential q available.

$$\frac{q + r_i \geq q_i \quad \Sigma; \Omega \not\vdash P_i :: (x : S_i) \quad (\forall i \in I)}{\Sigma; \Omega \not\vdash \text{case } x (l_i \Rightarrow P_i)_{i \in I} :: (x : \&\{l_i^{r_i} : S_i\})} \&R$$

The second premise prescribes that the branch P_i is typed with potential q_i . Moreover, branch P_i is reached after receiving the label l_i with potential r_i . Hence, the initial potential q must be able to cover the difference $q_i - r_i$. Since potential q can typecheck all the branches, we get the constraint $q \geq q_i - r_i$ for all i .

To spawn a new process defined by X , we split the context Ω into $\Omega_1 \ \Omega_2$, and we use Ω_1 to type the newly spawned process and Ω_2 for the continuation Q_x .

$$\frac{r \geq p + q \quad x' : S \leftarrow X @ p \leftarrow \overline{y'} : W = P_{x', \overline{y'}} \in \Sigma}{\Omega_1 = \overline{y} : W \quad \Sigma; \Omega_2 (x : S) \not\vdash Q_x :: (z : U)} \text{spawn} \\
\Sigma; \Omega_1 \ \Omega_2 \not\vdash (x \leftarrow X \leftarrow \overline{y}; Q_x) :: (z : U)$$

If the spawned process needs potential p (indicated by the signature) and the continuation needs potential q then the whole process needs potential $r \geq p + q$.

A forwarding process $x \leftarrow y$ terminates and its potential q is lost. Since we do not count forwarding messages in our cost semantics, we don't need any potential to type the forward.

$$\frac{q \geq 0}{\Sigma; y : S \not\vdash x \leftarrow y :: (x : S)} \text{id}$$

The rest of the rules are given in Figure 5. They are similar to the discussed rules and we omit their explanation.

As an illustration, the implementation and resource-aware type derivation (marked in red) of the binary counter is presented in Figure 3. The derivation provides a proof certificate that increment has an amortized cost of 1, while reading a value costs $2 \lceil \log(n+1) \rceil + 2$.

$$\begin{array}{c}
\frac{}{\Sigma; (\cdot) \not\vdash (\cdot) :: (\cdot)} \text{emp} \\
\frac{\Sigma; \Omega \not\vdash C :: \Omega' \quad \Sigma; \Omega' \not\vdash C' :: \Omega''}{\Sigma; \Omega \not\vdash (C \ C') :: \Omega''} \text{compose} \\
\frac{\Sigma; \Omega \not\vdash (C \ C') :: \Omega''}{\Sigma; \Omega_1 \not\vdash P :: (x : A)} C_{\text{proc}} \\
\frac{\Sigma; \Omega \ \Omega_1 \not\vdash (\text{proc}(x, w, P)) :: (\Omega (x : A))}{\Sigma; \Omega_1 \not\vdash P :: (x : A)} C_{\text{msg}} \\
\frac{\Sigma; \Omega \ \Omega_1 \not\vdash (\text{msg}(x, w, P)) :: (\Omega (x : A))}{\Sigma; \Omega \ \Omega_1 \not\vdash P :: (x : A)} C_{\text{msg}}
\end{array}$$

Figure 6. Typing rules for a configuration

6 Soundness

This section concludes the discussion of Resource-Aware SILL by demonstrating the soundness of the resource-aware type system with respect to the cost semantics. So far, we have analyzed and type-checked processes in isolation. However, as our cost semantics indicates, processes always exist in a configuration interacting with other processes. Thus, we need to extend the typing rules to arbitrary configurations.

Configuration Typing At runtime, a program in Resource-Aware SILL is a set of processes interacting via channels. Such a set is represented as a multi-set of proc and msg predicates as described in Section 4. To type the resulting configuration C , we first need to define a well-formed signature.

A signature Σ is *well formed* if (a) every type definition $V = S_V$ is contractive, and (b) every process definition $x : S \leftarrow X @ p \leftarrow \overline{y} : W = P_{x, \overline{y}}$ in Σ is well typed according to the process typing judgment, i.e. $\Sigma; \overline{y} : W \not\vdash P_{x, \overline{y}} :: (x : S)$. Note that the same process name X can have different resource-aware types in the signature Σ . We pick the appropriate type while applying the spawn rule.

We use the following judgment to type a configuration.

$$\Sigma; \Omega_1 \not\vdash C :: \Omega_2$$

It states that Σ is well-formed and that the configuration C uses the channels in the context Ω_1 and provides the channels in the context Ω_2 . The natural number E denotes the sum of the total potential and work done by the system. We call E the energy of the configuration. The configuration typing judgment is defined using the rules presented in Figure 6. The rule emp defines that an empty configuration is well-typed with energy 0. The rule compose composes two configurations C and C' ; C provides service on the channels in Ω' while C' uses the channels in Ω' . The energy of the composed configuration $C \ C'$ is obtained by summing up their individual energies. The rule C_{proc} creates a configuration out of a single process. The energy of this singleton configuration is obtained by adding the potential of the process and the work performed by it. Similarly, the rule C_{msg} creates a configuration out of a single message. Messages are typed in a cost-free judgment where $M^{\text{label}} = M^{\text{channel}} = M^{\text{close}} = 0$, introduced in Section 5.

Soundness Theorem 6.1 is the main theorem of the paper. It is a stronger version of a classical type preservation theorem and the usual type preservation is a direct consequence. Intuitively, it states that the energy of a configuration never increases during an

evaluation step, i.e. the energy remains conserved with possibly some loss (due to throwing away potential, or friction).

Theorem 6.1 (Soundness). *Consider a well-typed configuration C w.r.t. a well-formed signature Σ such that $\Sigma; \Omega_1 \stackrel{E}{\models} C :: \Omega_2$. If $C \mapsto C'$, then there exist Ω'_1, Ω'_2 and E' such that $\Sigma; \Omega'_1 \stackrel{E'}{\models} C' :: \Omega'_2$ and $E' \leq E$.*

The proof of the soundness theorem is achieved by a case analysis on the cost semantics, followed by an inversion on the typing of a configuration. The complete proof is presented in Appendix B. The preservation theorem is a corollary since soundness implies that the configuration C' is well-typed.

The soundness implies that the energy of an initial configuration is an upper bound on the energy of any configuration it will ever step to. In particular, if a configuration starts with 0 work, the initial energy (equal to initial potential) is an upper bound on the total work performed by an evaluation starting in that configuration.

Corollary 6.2 (Upper Bound). *If $\Sigma; \Omega_1 \stackrel{E}{\models} C :: \Omega_2$, and $C \mapsto^* C'$, then $E \geq W'$, where W' is the total work performed by the configuration C' , i.e. the sum of the work performed by each process and message in C' . In particular, if the work done by the initial configuration C is 0, then the potential P of the initial configuration satisfies $P \geq W'$.*

Proof. Applying the Soundness theorem successively, we get that if $C \mapsto^* C'$ and $\Sigma; \Omega_1 \stackrel{E}{\models} C :: \Omega_2$ and $\Sigma; \Omega'_1 \stackrel{E'}{\models} C' :: \Omega'_2$, then $E' \leq E$. Also, $E' = P' + W'$, where P' is the total potential of C' , while W' is the total work performed so far in C' . Since $P' \geq 0$, we get that $W' \leq P' + W' = E' \leq E$. In particular, if $W = 0$, we get that $P = P + W = E \geq W'$, where P and W are the potential and work of the initial configuration respectively. \square

The progress theorem of Resource-Aware SILL is a direct consequence of progress in SILL [33]. Our cost semantics are a cost observing semantics, i.e. it is just annotated with counters observing the work. Hence, any runtime step that can be taken by a program in SILL can be taken in Resource-Aware SILL.

7 Case Study: Stacks and Queues

As an illustration of our type system, we present a case study on stacks and queues. Stacks and queues have the same interaction protocol: they store elements of a variable type A and support inserting and deleting elements. They only differ in their implementation and resource usage. We express their common interface type as the simple session type store_A (parameterized by type variable A).

$$\begin{aligned} \text{store}_A &= \&\{ \text{ins} : A \multimap \text{store}_A, \\ &\quad \text{del} : \oplus\{\text{none} : 1, \text{some} : A \otimes \text{store}_A\} \} \end{aligned}$$

The session type dictates that a process providing a service of type store_A gives a client the choice to either insert (ins) or delete (del) an element of type A . Upon receipt of the label ins , the providing process expects to receive a channel of type A to be enqueued and recurses. Upon receipt of the label del , the providing process either indicates that the queue is empty (none), in which case it terminates, or that there is an element stored in the queue (some), in which case it deletes this element, sends it to the client, and recurses.

To account for the resource cost, we add potential annotations leading to two different resource-aware types for stacks and queues. Since we are interested in counting the total number of messages

exchanged, we again set $M^{\text{label}} = M^{\text{channel}} = M^{\text{close}} = 1$ to obtain a concrete bound.

Stacks The type for stacks is defined as follows.

$$\begin{aligned} \text{stack}_A &= \&\{ \text{ins}^0 : A \multimap^0 \text{stack}_A, \\ &\quad \text{del}^2 : \oplus\{\text{none}^0 : 1^0, \text{some}^0 : A \otimes^0 \text{stack}_A\} \} \end{aligned}$$

A stack is implemented using a sequence of *elem* processes terminated by an *empty* process. Each *elem* process stores an element of the stack, while *empty* denotes the end of stack. The implementations and type derivations of *elem* and *empty* are presented in Appendix C.2.

Inserting an element simply spawns a new *elem* process (which has no cost in our semantics), thus having no resource cost. Deleting an element terminates the *elem* process at the head. Before termination, it sends two messages back to the client, either none followed by close or some followed by *element*. Thus, deletion has a resource cost of 2. This is reflected in the stack_A type, where ins and del are annotated with 0 and 2 units of potential respectively.

Queues The queue interface is achieved by using the same store_A type and annotating it with a different potential. The tight potential bound depends on the number of elements stored in the queue. Hence, a precise resource-aware type needs access to this internal measure in the type. A type $\text{queue}_A[n]$ intuitively defines a queue of size n (for $n > 0$).

$$\begin{aligned} \text{queue}_A[n] &= \&\{ \text{ins}^{2n} : A \multimap^0 \text{queue}_A[n+1], \\ &\quad \text{del}^2 : \oplus\{\text{none}^0 : 1^0, \text{some}^0 : A \otimes^0 \text{queue}_A[n \dot{-} 1]\} \} \end{aligned}$$

The $\dot{-}$ operator denotes the monus operator defined as $a \dot{-} b = \max(0, a - b)$. This prevents the type $\text{queue}_A[0]$ from referring the undefined type $\text{queue}_A[-1]$ in the del label. Resource-aware session types also allow us to provide a more precise type for queue_A , i.e. different types for $\text{queue}_A[0]$ and $\text{queue}_A[n]$ for $n > 0$.

$$\begin{aligned} \text{queue}_A[0] &= \&\{ \text{ins}^0 : A \multimap^0 \text{queue}_A[1], \\ &\quad \text{del}^2 : \oplus\{\text{none}^0 : 1^0\} \} \\ \text{queue}_A[n] &= \&\{ \text{ins}^{2n} : A \multimap^0 \text{queue}_A[n+1], \\ &\quad \text{del}^2 : \oplus\{\text{some}^0 : A \otimes^0 \text{queue}_A[n-1]\} \} \end{aligned}$$

Similar to a stack, a queue is also implemented by a sequence of *elem* processes, connected via channels, and terminated by the *empty* process. Their implementations are presented in Appendix C.4.

For each insertion, the ins label along with the element travels to the end of the queue. There, it spawns a new *elem* process that holds the inserted element. Hence, the resource cost of each insertion is $2n$ where n is the size of the queue. On the other hand, deletion is similar to that of stack and has a resource cost of 2. Again, this is reflected in the queue_A type, where ins and del are annotated with $2n$ and 2 units of potential respectively.

The resource-aware types show that stacks are more efficient than queues. The label ins is annotated by 0 for stack_A and with $2n$ for queue_A . The label del has the same annotation in both types. Hence, an efficiency comparison can be performed by simply observing the resource-aware session types.

Queues as two stacks In a functional language, a queue is often implemented with two lists. The idea is to enqueue into the first list and to dequeue from the second list. If the second list is empty then we copy the first list over, thereby reversing its order. Since

the cost of the dequeue operation varies drastically between the dequeue operations, amortized analysis is again instrumental in the analysis of the worst-case behavior and shows that the worst-case amortized cost for deletion is actually a constant.

Appendix C.5 contains an implementation of a functional queue in Resource-Aware SILL. The type of the queue is

$$\text{queue}'_A = \&\{\text{ins}^6 : A \overset{0}{\dashv} \text{queue}'_A, \\ \text{del}^2 : \oplus\{\text{some}^0 : A \otimes \text{queue}'_A, \text{none}^0 : \mathbf{1}^0\}\}$$

Resource-aware session types enable us to translate the amortized analysis to the distributed setting. The type prescribes that an insertion has an amortized cost of 6 while the deletion has an amortized cost of 2. The main idea here is that the elements are inserted with a constant potential in the first list. While deleting, if the second list is empty, then this stored potential in the first list is used to pay for copying the elements over to the second list. The exact potential annotations for the two lists can be found in Appendix C.5. As demonstrated from the resource-aware type, this implementation is more efficient than the previous queue implementation, which has a linear resource cost for insertion.

Generic clients The notion of efficiency of a store can be generalized and quantified by considering clients for the stack and queue interface. A client interacts with a generic store via a sequence of insertions and deletions. A provider can then implement the store as a stack, queue, priority queue, etc. (same interface) and just expose the resource-aware type for store_A . Our type system uses just the interface type and the generic client implementation to derive resource bounds on the client. For simplicity, the clients are typed in an affine type system which allows us to throw away dummy channels.

We provide a general mechanism for implementing clients for a generic store. We define a generic store_A type at which the potential annotations are arbitrary natural numbers (or functions of internal measure n).

$$\text{store}_A[n] = \&\{\text{ins}^i : A \overset{a}{\dashv} \text{store}_A[n+1], \\ \text{del}^d : \oplus\{\text{none}^p : \mathbf{1}^e, \text{some}^s : A \overset{t}{\otimes} \text{store}_A[n-1]\}\}$$

A client of the store_A interface is defined by a list ℓ of ins and del messages that it sends to the store. We index the client $C_{\ell,n}$ by ℓ and the internal measure n of $\text{store}_A[n]$. The channel along which the client provides is irrelevant for our analysis and is represented using a dummy channel $d : D$. For ease of notation, we define the potential needed for a client $C_{\ell,n}$ as a function $\phi(\ell, n)$.

We implement the client $C_{\ell,n}$ as follows. First, consider the case when $\ell = []$, i.e. an empty list. The client for an empty list just closes the channel d . We assume that all clients are typed with the cost-free metric to only count for the messages sent within the stores. Hence, $C_{[],n}$ needs 0 potential. For the potential function, this means $\phi([], n) = 0$. Next, consider the client when the head of the list ℓ is ins. The client sends an ins label followed by the element x . If $C_{\text{ins}::\ell,n}$ needs a potential q , then the type derivation informs us that $C_{\ell,n+1}$ needs a potential $q - i - a$. Thus, $\phi(\text{ins}::\ell, n) = \phi(\ell, n+1) + i + a$. Finally, consider the client when the head of the list ℓ is del. The client sends the del label and then case analyzes on the label it receives. If it receives the same label, it receives the element and then continues with $C_{\ell,n-1}$, else it receives the none label and waits for the channel s to close. In terms of the potential

function, this means

$$\phi(\text{del}::\ell, n) = \begin{cases} \phi(\ell, n-1) + d - s - t & \text{if } n > 0 \\ \max(0, d - p - e) & \text{otherwise} \end{cases}$$

Walking through the list ℓ and chaining the potential equations together, $\phi(\ell, n)$ achieves a resource bound on the client $C_{\ell,n}$. Appendix A provides a detailed mechanism for implementing such clients and deriving their resource bounds.

The stack_A and queue_A interface types are specific instantiations of the store_A type. The derivation for a client of the stack_A interface defines the following potential equations.

$$\phi([], n) = 0 \quad \phi(\text{ins}::\ell, n) = \phi(\ell, n+1) \quad \phi(\text{del}::\ell, n) = \phi(\ell, n-1) + 2$$

Similarly, considering the queue type as another instantiation defines the following potential equations (again $\phi([], n) = 0$)

$$\phi(\text{ins}::\ell, n) = \phi(\ell, n+1) + 2n \quad \phi(\text{del}::\ell, n) = \phi(\ell, n-1) + 2$$

This allows us to compare arbitrary clients of two interfaces and compare their resource cost. The resource-aware types are expressive enough to obtain these resource bounds without referring the implementation of the store interface. For instance, an important property of queues is that every insertion is more costly than the previous one. The cost of insertion depends on the size of the queue, which, in turn, increases with every insertion. Hence, the complexity of the queue system depends on the sequence in which inserts and deletes are performed. In particular, we can consider the efficiency of two different clients for the queue system, by solving the above system of equations.

Consider two clients $Q_{\ell_1,n}$ and $Q_{\ell_2,n}$, with two different message lists (but same number of insertions and deletions); $\ell_1 = [\text{ins}, \dots, \text{ins}, \text{del}, \dots, \text{del}]$, i.e. m insertions followed by m deletions, and $\ell_2 = [\text{ins}, \text{del}, \text{ins}, \text{del}, \dots, \text{ins}, \text{del}]$, i.e. m instances of alternate insertions and deletions. Both clients send the same number of insertions and deletions. However, their resource cost are completely different. Typing the two clients, we obtain $\phi(\ell_1, n) = 2mn + m(m-1) + 2m$, while $\phi(\ell_2, n) = 2m(n+1)$, which shows that the second client is an order of magnitude more efficient than the first one.

Appendix C provides an exhaustive list of examples (in addition to stacks and queues) with their resource-aware types and derived bounds. This includes standard list processes (*nil*, *cons* and *append*) and higher order processes (*map* and *fold*).

8 Related Work

Session types were introduced by Honda [22]. The technical development in this work is based on previous work in [28, 33]. By removing the potential annotation from the type rules in Section 5 we arrive at the type system of loc. cit. The internal measures and type families we use are inspired by [15]. In contrast to our work, the aforementioned articles do not discuss resource analysis.

In the context of process calculi, capabilities [32] and static analyses [24] have been used to statically restrict communication for controlling buffer sizes in languages without session types. For session-typed communication, upper bounding the size of message queues is simpler and studied in the compiler for Concurrent C0 [34]. In contrast to capabilities, our potential annotations do not control buffer sizes but provide a symbolic description of the number of messages exchanged at runtime. It is not clear how capabilities could be used to perform such an analysis.

Type systems for static resource bound analysis for sequential programs have been extensively studied (e.g., [11, 25]). Our work is based on type-based amortized resource analysis. Automatic amortized resource analysis (AARA) has been introduced as a type system to automatically derive linear [19] and polynomial bounds [17] for sequential functional programs. It can also be integrated with program logics to derive bounds for imperative programs [5, 9]. Moreover, it has been used to derive bounds for term-rewrite systems [21] and object-oriented programs [20]. A recent work also considers bounds on the parallel evaluation cost (also called *span*) of functional programs [18]. The innovation of our work is the integration of AARA and session types and the analysis of message-passing programs that communicate with the outside world. Instead of function arguments, our bounds depend on the messages that are sent along channels. As a result, the formulation and proof of the soundness theorem is quite different from the soundness of sequential AARA.

We are only aware of a couple of other works that study resource bounds for concurrent programs. Gimenez et al. [13] introduced a technique for analyzing the parallel and sequential space and time cost of evaluating interaction nets. While it also based on linear logic and potential annotations, the flavor of the analysis is quite different. Interaction nets are mainly used to model parallel evaluation while session types focus on the interaction of processes. A main innovation of our work is that processes can exchange potential via messages. It is not clear how we can represent the examples we consider in this article as interaction nets. Albert et al. [2, 3] have studied techniques for deriving bounds on the cost of concurrent programs that are based on the actor model. While the goals of the work are similar to ours, the used technique and considered examples are dissimilar. A major difference is that our method is type-based and compositional. A unique feature of our work is that types describe bounds as functions of the messages that are sent along a channel.

9 Conclusion and Future Work

We have introduced resource-aware session types, a linear type system that combines session types [22, 28] and type-based amortized resource analysis [17, 19] to reason about the resource usage of message-passing processes. The soundness of the type system has been proved for a core session-typed language with respect to a cost semantics that tracks the total communication cost in a system of processes. We have demonstrated that our technique can be used to prove tight resource bounds and supports amortized reasoning by analyzing standard session-type data structures such as distributed binary counters, stacks, and queues.

Our approach addresses some of the main challenges of analyzing message-passing programs such as compositionality and description of symbolic bounds. However, there are several open problems that we plan to tackle as part of future work. The technique we have developed in this paper does not yet account for the concurrent execution cost of processes, or the *span*. We are working on a companion paper that describes a type-based analysis to derive bounds on the span; the earliest time a concurrent computation terminates assuming an infinite number of processors. As another direction, we plan to implement our analysis. An advantage of our method is that it is based upon type-based amortized resource analysis for sequential programs. We designed the type system with automation in mind and we are confident that we can support

automatic type inference using templates and LP solving similar to AARA [17, 19]. To this end, we are working on an algorithmic version of the declarative type system presented here.

References

- [1] Umut A. Acar, Arthur Charguéraud, and Mike Rainey. 2016. Oracle-Guided Scheduling for Controlling Granularity in Implicitly Parallel Languages. *J. Funct. Programming* (2016).
- [2] Elvira Albert, Puri Arenas, Jesús Correas, Samir Genaim, Miguel Gómez-Zamalloa, Enrique Martín-Martín, Germán Puebla, and Guillermo Román-Díez. 2015. Resource Analysis: From Sequential to Concurrent and Distributed Programs. In *FM'15*.
- [3] Elvira Albert, Antonio Flores-Montoya, Samir Genaim, and Enrique Martín-Martín. 2016. May-Happen-in-Parallel Analysis for Actor-Based Concurrency. *ACM Trans. Comput. Log.* (2016).
- [4] Timos Antonopoulos, Paul Gazzillo, Michael Hicks, Eric Koskinen, Tachio Teruchi, and Shiyi Wei. 2017. Decomposition Instead of Self-composition for Proving the Absence of Timing Channels. In *PLDI'17*.
- [5] Robert Atkey. 2010. Amortised Resource Analysis with Separation Logic. In *ESOP'10*.
- [6] Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. In *ICFP'16*.
- [7] Richard P. Brent. 2013. *Algorithms for minimization without derivatives*. Courier Corporation.
- [8] Luís Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *CONCUR'10*.
- [9] Quentin Carbonneaux, Jan Hoffmann, Thomas Reps, and Zhong Shao. 2017. Automated Resource Analysis with Coq Proof Objects. In *CAV'17*.
- [10] Ilano Cervesato and Andre Scedrov. 2006. Relating State-Based and Process-Based Concurrency Through Linear Logic. *Electron. Notes Theor. Comput. Sci.* (2006).
- [11] Ezgi Çiçek, Deepak Garg, and Umut A. Acar. 2015. Refinement Types for Incremental Computational Complexity. In *ESOP'15*.
- [12] Simon J. Gay and Malcolm Hole. 2005. Subtyping for Session Types in the π -Calculus. *Acta Informatica* (2005).
- [13] Stéphane Gimenez and Georg Moser. 2016. The Complexity of Interaction. In *POPL'16*.
- [14] Jean-Yves Girard. 1987. Linear logic. *Theor. Comp. Sc.* (1987).
- [15] Dennis Griffith and Elsa L. Gunter. 2013. Liquid Pi: Inferrable Dependent Session Types. In *NASA Formal Methods Symp.'13*.
- [16] Rémy Haemmerlé, Pedro López-García, Umer Liqat, Maximiliano Klemen, John P. Gallagher, and Manuel V. Hermenegildo. 2016. A Transformational Approach to Parametric Accumulated-Cost Static Profiling. In *FLOPS'16*.
- [17] Jan Hoffmann, Ankush Das, and Shu-Chun Weng. 2017. Towards Automatic Resource Bound Analysis for OCaml. In *POPL'17*.
- [18] Jan Hoffmann and Zhong Shao. 2015. Automatic Static Cost Analysis for Parallel Programs. In *ESOP'15*.
- [19] Martin Hofmann and Steffen Jost. 2003. Static Prediction of Heap Space Usage for First-Order Functional Programs. In *POPL'03*.
- [20] Martin Hofmann and Steffen Jost. 2006. Type-Based Amortised Heap-Space Analysis. In *ESOP'06*.
- [21] Martin Hofmann and Georg Moser. 2015. Multivariate Amortised Resource Analysis for Term Rewrite Systems. In *TLCA'15*.
- [22] Kohei Honda. 1993. Types for dyadic interaction. In *CONCUR'93*.
- [23] Yu Feng Jia Chen and Isil Dillig. 2017. Precise Detection of Side-Channel Vulnerabilities using Quantitative Cartesian Hoare Logic. In *CCS'17*.
- [24] Naoki Kobayashi, Motoki Nakade, and Akinori Yonezawa. 1995. Static analysis of communication for asynchronous concurrent programming languages. In *SAS'95*.
- [25] Ugo Dal Lago and Barbara Petit. 2013. The Geometry of Types. In *POPL'13*.
- [26] Van Chan Ngo, Mario Dehesa-Azuara, Matthew Fredrikson, and Jan Hoffmann. 2017. Verifying and Synthesizing Constant-Resource Implementations with Types. In *S&P'17*.
- [27] Oswaldo Olivo, Isil Dillig, and Calvin Lin. 2015. Static Detection of Asymptotic Performance Bugs in Collection Traversals. In *PLDI'15*.
- [28] Frank Pfenning and Dennis Griffith. 2015. Polarized Substructural Session Types. In *FOSSACS'15*.
- [29] Frank Pfenning and Robert J. Simmons. 2009. Substructural Operational Semantics As Ordered Logic Programming. In *LICS'09*.
- [30] Miguel Silva, Mário Florido, and Frank Pfenning. 2016. Non-Blocking Concurrent Imperative Programming with Session Types. In *LINEARITY'16*.
- [31] Robert Endre Tarjan. 1985. Amortized Computational Complexity. *SIAM J. Alg. Disc. Methods* (1985).
- [32] Tachio Teruchi and Adam Megacz. 2008. Inferring Channel Buffer Bounds Via Linear Programming. In *ESOP'08*.
- [33] Bernardo Toninho, Luís Caires, and Frank Pfenning. 2013. Higher-Order Processes, Functions, and Sessions: A Monadic Integration. In *ESOP'13*.
- [34] Max Willsey, Rokhini Prabhu, and Frank Pfenning. 2016. Design and Implementation of Concurrent Co. In *LINEARITY'16*.

A Generic Clients

The notion of efficiency of a store can be generalized and quantified by considering clients for the stack and queue interface. A client interacts with a generic store via a sequence of insertions and deletions. A provider can then implement the store as a stack, queue, priority queue, etc. (same interface) and just expose the resource-aware type for store_A . Our type system uses just the interface type and the generic client implementation to derive resource bounds on the client. For simplicity, the clients are typed in an affine type system which allows us to throw away dummy channels.

We provide a general mechanism for implementing clients for a generic store. We define a generic store_A type at which the potential annotations are arbitrary natural numbers.

$$\text{store}_A[n] = \&\{\text{ins}^i : A \overset{a}{\multimap} \text{store}_A[n+1], \\ \text{del}^d : \oplus\{\text{none}^p : 1^e, \text{some}^s : A \overset{t}{\otimes} \text{store}_A[n-1]\}\}$$

A client of the store_A interface is defined by a list ℓ of ins and del messages that it sends to the store. We index the client $C_{\ell,n}$ by ℓ and the internal measure n of $\text{store}_A[n]$. The channel along which the client provides is irrelevant for our analysis and is represented using a dummy channel $d : D$. For ease of notation, we define the potential needed for a client $C_{\ell,n}$ as a function $\phi(\ell, n)$.

We implement the client $C_{\ell,n}$ as follows. First, consider the case when $\ell = []$, i.e. an empty list.

$$\cdot \overset{\circ}{\leftarrow} C_{[],n} :: (d : D) \\ d \leftarrow C_{[],n} = \text{close } d$$

The client for an empty list just closes the channel d . We assume that all clients are typed with the cost-free metric to only count for the messages sent inside the stores. So $C_{[],0}$ needs 0 potential. For the potential function, this means $\phi([], n) = 0$.

Next, Figure 7 implements the client when the head of the list ℓ is ins. The client sends an ins label followed by the element x . If $C_{\text{ins}::\ell,n}$ needs a potential q , then the type derivation informs us that $C_{\ell,n+1}$ needs a potential $q - i - a$. Thus, $\phi(\text{ins} :: \ell, n) = \phi(\ell, n+1) + i + a$.

Finally, Figure 8 shows the client implementation if the head of the list ℓ is del. The client sends the del label and then case analyzes on the label it receives. If it receives the some label, it receives the element and then continues with $C_{\ell,n-1}$, else it receives the none label and waits for the channel s to close. In terms of the potential function, this means

$$\phi(\text{del} :: \ell, n) = \begin{cases} \phi(\ell, n-1) + d - s - t & \text{if } n > 0 \\ \max(0, d - p - e) & \text{otherwise} \end{cases}$$

Walking through the list ℓ and chaining the potential equations together, we can achieve a resource bound on the client $C_{\ell,n}$ by computing $\phi(\ell, n)$.

The stack_A and queue_A interface types are specific instantiations of the store_A type. For the stack interface, plugging in appropriate potential annotations $i = a = p = e = s = t = 0$, and $d = 2$, we get (ignoring the case where the stack becomes empty)

$$\begin{aligned} \phi([], n) &= 0 \\ \phi(\text{ins} :: \ell, n) &= \phi(\ell, n+1) \\ \phi(\text{del} :: \ell, n) &= \phi(\ell, n-1) + 2 \end{aligned}$$

Similarly, considering the queue type as another instantiation of the store_A type, and plugging $a = p = e = s = t = 0$, $i = 2n$ and

$d = 2$, we get

$$\begin{aligned} \phi([], n) &= 0 \\ \phi(\text{ins} :: \ell, n) &= \phi(\ell, n+1) + 2n \\ \phi(\text{del} :: \ell, n) &= \phi(\ell, n-1) + 2 \end{aligned}$$

Finally, looking at queue'_A as another instantiation of the store_A type and plugging $a = p = e = s = t = 0$, $i = 6$ and $d = 2$, we get

$$\begin{aligned} \phi([], n) &= 0 \\ \phi(\text{ins} :: \ell, n) &= \phi(\ell, n+1) + 6 \\ \phi(\text{del} :: \ell, n) &= \phi(\ell, n-1) + 2 \end{aligned}$$

This allows us to compare arbitrary clients of two (same or different) interfaces and compare their resource cost. The resource-aware types are expressive enough to obtain these resource bounds without referring the implementation of the store interface. For instance, an important property of queues is that every insertion is more costly than the previous one. The cost of insertion depends on the size of the queue, which, in turn, increases with every insertion. Hence, the complexity of the queue system depends on the sequence in which inserts and deletes are performed. In particular, we can consider the efficiency of two different clients for the queue system, by solving the above system of equations.

For instance, consider two clients $Q_{\ell_1,n}$ and $Q_{\ell_2,n}$, with two different message lists $\ell_1 = [\text{ins}, \dots, \text{ins}, \text{del}, \dots, \text{del}]$, i.e. m insertions followed by m deletions, and $\ell_2 = [\text{ins}, \text{del}, \text{ins}, \text{del}, \dots, \text{ins}, \text{del}]$, i.e. m instances of alternate insertions and deletions. In both cases, we have the same number of insertions and deletions. However, the resource cost of the two systems are completely different. Solving the system of equations, we get that $\phi(\ell_1, n) = 2mn + m(m-1) + 2m$, while $\phi(\ell_2, n) = 2m(n+1)$, which shows that the second client is an order of magnitude more efficient than the first one.

B Proof of Soundness Theorem

We present the complete proof of the soundness theorem. We reiterate the rules of cost semantics and typing. Figure 4 presents the cost semantics for our session-typed language, while Figure 5 presents the full set of typing rules. Finally, Theorem 6.1 defines the soundness theorem establishing that the typing rules for a configuration presented in Figure 6 are sound w.r.t. the rules for cost semantics presented in Figure 4. We follow the complete proof of the soundness theorem.

Proof. The proof proceeds by case analysis on the cost semantics of our language, i.e. on the judgment $C \mapsto C'$. By the compose rule, we can split the configuration such that $C = (C_M \ \mathcal{D})$ and $C' = (C'_M \ \mathcal{D})$ and $C_M \mapsto C'_M$. Using the compose rule,

$$\frac{\Sigma ; \Omega \Vdash^{S_M} C_M :: \Omega' \quad \Sigma ; \Omega' \Vdash^{S_D} \mathcal{D} :: \Omega''}{\Sigma ; \Omega \Vdash^{S_M+S_D} (C_M \ \mathcal{D}) :: \Omega''} \text{compose} \\ \frac{\Sigma ; \Omega \Vdash^{S_M} C'_M :: \Omega' \quad \Sigma ; \Omega' \Vdash^{S_D} \mathcal{D} :: \Omega''}{\Sigma ; \Omega \Vdash^{S'_M+S_D} (C'_M \ \mathcal{D}) :: \Omega''} \text{compose}'$$

Hence, to show that $S' \leq S$, it suffices to show that $S'_M \leq S_M$. We proceed by case analysis on the $C_M \mapsto C'_M$ judgment.

$$\begin{array}{l}
\Omega(x : A) (s : \text{store}_A[n]) \text{ } \sharp^q C_{\text{ins}::\ell, n} :: (d : D) \\
d \leftarrow C_{\text{ins}::\ell, n} \leftarrow \Omega x s = \\
\text{s.ins ;} \\
\text{send } s x ; \\
d \leftarrow C_{\ell, n+1} \leftarrow \Omega s
\end{array}
\quad
\begin{array}{l}
\% \quad \Omega(x : A) (s : A \overset{a}{\dashv} \text{store}_A[n+1]) \text{ } \sharp^{q-i} d : D \\
\% \quad \Omega(s : \text{store}_A[n+1]) \text{ } \sharp^{q-i-a} d : D
\end{array}$$
Figure 7. Implementation and type derivation when head is ins
$$\begin{array}{l}
\Omega(s : \text{store}_A[n]) \text{ } \sharp^q C_{\text{del}::\ell} :: (d : D) \\
d \leftarrow C_{\text{del}::\ell, n} \leftarrow \Omega s = \\
\text{s.del ;} \\
\text{case } s (\\
\text{some } \Rightarrow x \leftarrow \text{recv } s ; \\
\quad d \leftarrow C_{\ell, n-1} \leftarrow \Omega x s \\
| \text{none } \Rightarrow \text{wait } s)
\end{array}
\quad
\begin{array}{l}
\% \quad \Omega(s : \oplus\{\text{none}^n : 1^e, \text{some}^s : A \overset{t}{\otimes} \text{store}_A[n-1]\}) \text{ } \sharp^{q-d} d : D \\
\% \quad \Omega(x : A) (s : \text{store}_A[n-1]) \text{ } \sharp^{q-d+s+t} d : D \\
\% \quad \Omega \text{ } \sharp^{q-d+p+e} d : D
\end{array}$$
Figure 8. Implementation and type derivation when head is del

- Case (spawn_c) : $C_M = \text{proc}(d, w, x \leftarrow P_x \leftarrow \bar{y} ; Q_x)$. Inverting the typing rule spawn on configuration C_M , we get

$$r \geq p + q \quad (4)$$

$$\Omega_1 \Omega_2 \text{ } \sharp^r x \leftarrow P_x \leftarrow \bar{y} ; Q_x :: (d : U)$$

and in $C'_M = \text{proc}(c, 0, P_c) \text{proc}(d, w, Q_c)$, we get (the premise due to inversion)

$$\Omega_1 \text{ } \sharp^p P_c :: (c : S) \quad \Omega_2 (c : S) \text{ } \sharp^q Q_c :: (d : U)$$

From the cost semantics rule spawn_c ,

$$\frac{\text{proc}(d, w, x \leftarrow P_x \leftarrow \bar{y} ; Q_x)}{\text{proc}(c, 0, P_c)} \text{spawn}_c \text{proc}(d, w, Q_c)$$

Since S is the sum of work and potential of each process in the configuration, using Equation 4, we get

$$\begin{aligned}
S'_M &= (pp + wp) + (pQ + wQ) \\
&= (p + 0) + (q + w) \\
&= p + q + w \\
&\leq r + w \\
&= S_M
\end{aligned}$$

- Case (fwd_s) : $C_M = \text{proc}(c, w, c \leftarrow d)$. Inverting the typing rule fwd on C_M ,

$$q \geq 0 \quad d : S \text{ } \sharp^q c \leftarrow d :: (c : S)$$

Inverting the same rule in C'_M ,

$$q' \geq 0 \quad d : S \text{ } \sharp'_{cf} c \leftarrow d :: (c : S)$$

Using the semantics rule fwd_s ,

$$\frac{\text{proc}(c, w, c \leftarrow d)}{\text{msg}(c, w, c \leftarrow d)} \text{fwd}_s$$

Since we are free to choose q' , we set $q' \leq q$. Computing S_M and S'_M , we get

$$\begin{aligned}
S'_M &= q' + w \\
&\leq q + w \\
&= S_M
\end{aligned}$$

- Case (fwd_r^+) : $C_M = \text{proc}(d, w, P) \text{msg}(c, w', c \leftarrow d)$. Inverting the fwd rule for C_M ,

$$q_1 \geq 0 \quad \Omega \text{ } \sharp^{q_1} P :: (d : S) \quad (5)$$

$$q_2 \geq 0 \quad d : S \text{ } \sharp^{q_2} c \leftarrow d :: (c : S)$$

Using Equation 5 and noting that c and d have the same type S , we get for C'_M ,

$$q'_1 \geq 0 \quad \Omega \text{ } \sharp^{q'_1} [c/d]P :: (c : S)$$

From the cost semantics rule fwd_r^+ , we get

$$\frac{\text{proc}(d, w, P) \quad \text{msg}(c, w', c \leftarrow d)}{\text{proc}(c, w + w', [c/d]P)} \text{fwd}_r^+$$

Since we are free to choose q'_1 , we set $q'_1 \leq q_1$. Computing S_M and S'_M , we get

$$\begin{aligned}
S'_M &= q'_1 + w + w' \\
&\leq q_1 + q_2 + w + w' \\
&= (q_1 + w) + (q_2 + w') \\
&= S_M
\end{aligned}$$

- Case (fwd_r^-) : $C_M = \text{proc}(e, w, P) \text{msg}(c, w', c \leftarrow d)$. Inverting the fwd rule for C_M ,

$$q_1 \geq 0 \quad \Omega (c : S) \text{ } \sharp^{q_1} P :: (e : U) \quad (6)$$

$$q_2 \geq 0 \quad d : S \text{ } \sharp^{q_2} c \leftarrow d :: (c : S)$$

Using Equation 6 and noting that c and d have the same type S , we get for C'_M ,

$$q'_1 \geq 0 \quad \Omega (d : S) \text{ } \sharp^{q'_1} [d/c]P :: (e : U)$$

From the cost semantics rule fwd_r^- , we get

$$\frac{\text{proc}(e, w, P) \quad \text{msg}(c, w', c \leftarrow d)}{\text{proc}(e, w + w', [d/c]P)} \text{fwd}_r^-$$

Since we are free to choose q'_1 , we set $q'_1 \leq q_1$.

$$\begin{aligned}
S'_M &= q'_1 + w + w' \\
&\leq q_1 + q_2 + w + w' \\
&= (q_1 + w) + (q_2 + w') \\
&= S_M
\end{aligned}$$

- Case ($\oplus C_s$): $C_M = \text{proc}(c, w, c.l_k ; P)$. Inverting the typing rule $\oplus R_k$ on C_M , we get

$$q_1 \geq p + r_k + M^{\text{label}} \quad (7)$$

$$\Omega \stackrel{\#1}{\vdash} (c.l_k ; P) :: (c : \oplus \{l_i^{r_i} : S_i\}_{i \in I})$$

and in C'_M , we get (the premise due to inversion)

$$\Omega \stackrel{\#}{\vdash} [c'/c]P :: (c' : S_k)$$

$$q'_1 \geq q_2 + r_k \quad (8)$$

$$c' : S_k \stackrel{\#1}{\vdash} (c.l_k ; c \leftarrow c') :: (c : \oplus \{l_i^{r_i} : S_i\}_{i \in I})$$

$$q_2 \geq 0 \quad c' : S_k \stackrel{\#2}{\vdash} c \leftarrow c' :: (c : S_k) \quad (9)$$

Using the cost semantics rule $\oplus C_s$, we get

$$\frac{\text{proc}(c, w, c.l_k ; P)}{\text{proc}(c', w + M^{\text{label}}, [c'/c]P) \quad \text{msg}(c, 0, c.l_k ; c \leftarrow c')} \oplus C_s$$

Again, $S'_M = (p + w + M^{\text{label}}) + (q'_1 + 0)$. Since, we need to prove that there exists such an S'_M , we can choose q'_1 and q_2 arbitrarily such that they satisfy Equations 8 and 9. We set $q_2 = 0$ and $q'_1 = r_k$. Hence, using Equation 7

$$\begin{aligned} S'_M &= (p + w + M^{\text{label}}) + r_k \\ &\leq q_1 + w \\ &\leq S_M \end{aligned}$$

- Case ($\oplus C_r$): In this case, $C_M = \text{msg}(c, w, c.l_k ; c \leftarrow c') \text{proc}(d, w', \text{case } c (l_i \Rightarrow Q_i)_{i \in I})$. Inverting the typing rule $\oplus L$ on C_M , we get

$$q_1 \geq q_2 + r_k \quad (10)$$

$$c' : S_k \stackrel{\#1}{\vdash} (c.l_k ; c \leftarrow c') :: (c : \oplus \{l_i^{r_i} : S_i\}_{i \in I})$$

$$q_2 \geq 0 \quad c' : S_k \stackrel{\#2}{\vdash} c \leftarrow c' :: (c : S_k) \quad (11)$$

$$q_3 + r_k \geq q_k \quad (12)$$

$$\Omega (c : \oplus \{l_i^{r_i} : S_i\}_{i \in I}) \stackrel{\#3}{\vdash} \text{case } c (l_i \Rightarrow Q_i)_{i \in I} :: (d : U)$$

and in C'_M (premise of the typing rule due to inversion), we get

$$\Omega (c' : S_k) \stackrel{\#k}{\vdash} [c'/c]Q_k :: (d : U)$$

Using the cost semantics rule $\oplus C_r$, we get

$$\frac{\text{msg}(c, w, c.l_k ; c \leftarrow c') \quad \text{proc}(d, w', \text{case } c (l_i \Rightarrow Q_i)_{i \in I})}{\text{proc}(d, w + w', [c'/c]Q_k)}$$

Again, using Equations 10, 11 and 12, we get

$$\begin{aligned} S'_M &= q_k + w + w' \\ &\leq q_3 + r_k + w + w' \\ &\leq q_3 + q_2 + r_k + w + w' \\ &\leq q_3 + q_1 + w + w' \\ &\leq (q_1 + w) + (q_3 + w') = S_M \end{aligned}$$

- Case ($\& C_s$): Analogous to $\rightarrow C_s$.
- Case ($\& C_r$): Analogous to $\rightarrow C_r$.
- Case ($\otimes C_s$): Analogous to $\oplus C_s$.
- Case ($\otimes C_r$): Analogous to $\oplus C_r$.

- Case ($\rightarrow C_s$): $C_M = \text{proc}(d, w, \text{send } c e ; P)$. Applying the rule $\rightarrow L$ on C_M , we get

$$q_1 \geq p + r + M^{\text{channel}}$$

$$\Omega (e : S) (c : S \xrightarrow{r} T) \stackrel{\#1}{\vdash} (\text{send } c e ; P) :: (d : U)$$

Inverting the same rule on C'_M , we get

$$\Omega (c' : T) \stackrel{\#}{\vdash} [c'/c]P :: (d : U)$$

$$q'_1 \geq q_2 + r$$

$$(e : S) (c : S \xrightarrow{r} T) \stackrel{\#1}{\vdash} \text{send } c e ; c' \leftarrow c :: (c' : T)$$

$$q_2 \geq 0 \quad c : T \stackrel{\#2}{\vdash} c' \leftarrow c :: (c' : T)$$

From the cost semantics rule $\rightarrow C_s$, we get

$$\frac{\text{proc}(d, w, \text{send } c e ; P)}{\text{proc}(d, w + M^{\text{channel}}, [c'/c]P) \quad \text{msg}(c', 0, \text{send } c e ; c' \leftarrow c)}$$

Since we can choose arbitrary values for q'_1 and q_2 satisfying the above inequalities, we set $q_2 = 0$ and $q'_1 = r$. Computing S_M and S'_M , we get

$$\begin{aligned} S'_M &= (p + w + M^{\text{channel}}) + (q'_1 + w') \\ &= (p + r + M^{\text{channel}} + w) + (q'_1 - r + w') \\ &\leq q_1 + w + w' \\ &= S_M \end{aligned}$$

- Case ($\rightarrow C_r$): In this case, $C_M = \text{msg}(c', w, \text{send } c e ; c' \leftarrow c) \text{proc}(c, w', x \leftarrow \text{recv } c ; Q_x)$. Applying the rule $\rightarrow L$ on the message in C_M , we get

$$q_1 \geq q_2 + r$$

$$(e : S) (c : S \xrightarrow{r} T) \stackrel{\#1}{\vdash} \text{send } c e ; c' \leftarrow c :: (c' : T)$$

$$q_2 \geq 0 \quad c : T \stackrel{\#2}{\vdash} c' \leftarrow c :: (c' : T)$$

Applying the rule $\rightarrow R$ on the process in C_M , we get

$$q_3 + r \geq p \quad \Omega \stackrel{\#3}{\vdash} (x \leftarrow \text{recv } c ; Q_x) :: (x : S \xrightarrow{r} T)$$

Inverting the $\rightarrow R$ rule, we get for C'_M ,

$$\Omega (e : S) \stackrel{\#}{\vdash} [c'/c]Q_e :: (c' : T)$$

From the cost semantics rule $\rightarrow C_r$, we get

$$\frac{\text{msg}(c', w, \text{send } c e ; c' \leftarrow c) \quad \text{proc}(c, w', x \leftarrow \text{recv } c ; Q_x)}{\text{proc}(c, w + w', [c'/c]Q_e)}$$

Computing S_M and S'_M , we get

$$\begin{aligned} S'_M &= p + w + w' \\ &\leq q_3 + r + w + w' \\ &\leq q_3 + q_2 + r + w + w' \\ &\leq q_3 + q_1 + w + w' \\ &= (q_1 + w) + (q_3 + w') \\ &= S_M \end{aligned}$$

- Case ($1C_s$): $C_M = \text{proc}(c, w, \text{close } c)$. Applying the $1R$ rule on C_M , we get

$$q \geq r + M^{\text{close}} \quad \cdot \stackrel{\#}{\vdash} \text{close } c :: (c : 1^r)$$

Inverting the same rule for C'_M , we get

$$q' \geq r \quad \cdot \stackrel{\#}{\vdash} \text{close } c :: (c : 1^r)$$

From the cost semantics rule $1C_s$, we get

$$\frac{\text{proc}(c, w, \text{close } c)}{\text{msg}(c, w + M^{\text{close}}, \text{close } c)} \quad 1C_s$$

Setting $q' = r$ and computing S_M and S'_M , we get

$$\begin{aligned} S'_M &= q' + w + M^{\text{close}} \\ &= r + w + M^{\text{close}} \\ &\leq q + w \\ &= S_M \end{aligned}$$

- Case ($1C_r$): $C_M = \text{msg}(c, w, \text{close } c) \text{ proc}(d, w', \text{wait } c ; Q)$. Applying the rule $1R$ on the message in C_M , we get

$$q_1 \geq r \quad \cdot \text{!}^{q_1} \text{close } c :: (c : \mathbf{1}^r)$$

Applying the $1L$ rule on the process in C_M , we get

$$q_2 + r \geq p \quad \Omega (c : \mathbf{1}^r) \text{!}^{q_2} \text{wait } c ; Q :: (d : U)$$

Inverting the $1L$ rule for C'_M , we get

$$\Omega \text{!}^p Q :: (d : U)$$

From the cost semantics rule $1C_r$, we get

$$\frac{\text{msg}(c, w, \text{close } c) \quad \text{proc}(d, w', \text{wait } c ; Q)}{\text{proc}(d, w + w', Q)} \quad 1C_r$$

Computing S_M and S'_M , we get

$$\begin{aligned} S'_M &= p + w + w' + M^{\text{close}} \\ &\leq q_2 + r + w + w' \\ &\leq q_2 + q_1 + w + w' \\ &\leq (q_1 + w) + (q_2 + w') \\ &= S_M \end{aligned}$$

Hence, in all of the above cases, $S'_M \leq S_M$ establishing that $S' \leq S$, thus showing that the potential type system is sound w.r.t. the cost semantics. \square

C More Examples

Our type system is quite expressive and can be used to derive bounds on many more examples. In this section, we will derive bounds on several list processes. We will start with simple examples, such as the *nil*, *cons* and *append* processes. We will then derive bounds on stacks and queues, first being implemented internally using a sequence of processes, and then using lists. Finally, we will conclude with some higher order processes such as *map* and *fold*. For each of the following examples, we assume the standard cost metric, where we count the number of messages exchanged, i.e. $M^{\text{label}} = M^{\text{channel}} = M^{\text{close}} = 1$.

C.1 Standard List Processes

This section presents the implementations of standard list processes, namely *nil*, *cons* and *append*. First, we consider the list protocol as a simple session type.

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

The type prescribes that a process providing service of type list_A will either send a label *cons* followed by an element of type A and recurse, or will send a *nil* label followed by a close message and then terminate. On the client side (i.e. a process that uses a channel

of type list_A in its context), the opposite behavior is observed, i.e. a client receives the messages that the provider sends (sequence of *cons* labels and elements terminated by a *nil* label and the close message).

C.1.1 *nil*

The *nil* process is used to create an empty list. Formally, a *nil* process uses an empty context, and provides an empty list along a channel $l : \text{list}_A$. Intuitively, this means it sends a *nil* label followed by a close message along l . Thus, the list l being offered by the *nil* process can be used as an empty list by some other process. First, we introduce the resource-aware session type for list_A^0 . (The superscript 0 on the type variable list_A^0 is insignificant and present only to distinguish different resource-aware types for the same base type list_A .)

$$\text{list}_A^0 = \oplus \{ \text{nil}^0 : \mathbf{1}^0, \text{cons}^0 : A \otimes \text{list}_A^0 \}$$

This resource-aware type decorates each label and type operator with 0 potential, implying that none of the messages carry any potential, and the process potential needs to pay only for the cost of sending the messages. We present the implementation followed by the type derivation for the *nil* process.

$$\begin{aligned} \cdot \text{!}^2 \text{nil} :: (l : \text{list}_A^0) \\ l \leftarrow \text{nil} = \\ \begin{array}{ll} l.\text{nil} ; & \% \cdot \text{!}^1 l : \mathbf{1}^0 \\ \text{close } l & \% \cdot \text{!}^0 \cdot \end{array} \end{aligned}$$

The type of *nil* process shows that it needs a potential of 2, which intuitively agrees with our cost model. The process sends two messages, each of them costing unit potential. We explain the type derivation briefly. The initial type of the process is $\cdot \text{!}^2 l : \text{list}_A$. For l to behave as an empty list, the *nil* process needs to send the *nil* label first. As the list_A type prescribes, the *nil* label carries no potential, this send only costs 1. Updating the type of l and the process potential, we get $\cdot \text{!}^1 l : \mathbf{1}^0$. Finally, the *nil* process needs to send the close message, which again costs 1 as the type $\mathbf{1}^0$ carries no potential. Our type system proves that the *nil* process needs a potential of 2, successfully verifying its resource usage as 2.

C.1.2 *cons*

The *cons* process takes an element x of type A and a tail t of type list_A in its context, and produces a list that *cons* x onto t . The *cons* process is implemented as follows.

$$\begin{aligned} (x : A) (t : \text{list}_A^0) \text{!}^2 \text{cons} :: (l : \text{list}_A^0) \\ l \leftarrow \text{cons} \leftarrow x t = \\ \begin{array}{ll} l.\text{cons} ; & \% (x : A) (t : \text{list}_A^0) \text{!}^1 l : A \otimes \text{list}_A^0 \\ \text{send } l x ; & \% (t : \text{list}_A^0) \text{!}^0 l : \text{list}_A^0 \\ l \leftarrow t & \end{array} \end{aligned}$$

The implementation dictates that the *cons* process first sends the *cons* message along l , followed by the element x . It then forwards the channel l along t . The resource cost of *cons* is 2 since it sends two messages before it forwards. This is verified by the potential of 2 in the type of the process.

Although the *cons* process has a unique simple session type, it can have several resource-aware session types. We demonstrate one such type below. First, we introduce the resource-aware type of list_A .

$$\text{list}_A^1 = \oplus \{ \text{nil}^0 : \mathbf{1}^0,$$

```

1:  $(l_1 : \text{list}_A^2) (l_2 : \text{list}_A^0) \stackrel{0}{\vdash} \text{append} :: (l : \text{list}_A^0)$ 
2:  $l \leftarrow \text{append} \leftarrow l_1 l_2 =$ 
3:   case  $l_1$  (cons  $\Rightarrow x \leftarrow \text{recv } l_1$  ;           %  $(x : A) (l_1 : \text{list}_A^2) (l_2 : \text{list}_A^0) \stackrel{2}{\vdash} l : \text{list}_A^0$ 
4:          $l.\text{cons}$  ;                                     %  $(x : A) (l_1 : \text{list}_A^2) (l_2 : \text{list}_A^0) \stackrel{1}{\vdash} l : \text{list}_A^0$ 
5:         send  $l x$  ;                                   %  $(l_1 : \text{list}_A^2) (l_2 : \text{list}_A^0) \stackrel{0}{\vdash} l : \text{list}_A^0$ 
6:          $l \leftarrow \text{append} \leftarrow l_1 l_2$ 
7:   | nil  $\Rightarrow$  wait  $l_1$  ;                               %  $(l_1 : \mathbf{1}) (l_2 : \text{list}_A^0) \stackrel{0}{\vdash} l : \text{list}_A^0$ 
8:          $l \leftarrow l_2$                                %  $(l_2 : \text{list}_A^0) \stackrel{0}{\vdash} l : \text{list}_A^0$ 

```

Figure 9. Implementation and type derivation of the *append* process

$$\text{cons}^1 : A \otimes \text{list}_A^1 \}$$

The type prescribes that the *cons* message carries a potential of 1, while all other messages carry no potential. We present the implementation along with the type derivation.

```

 $(x : A) (t : \text{list}_A^1) \stackrel{3}{\vdash} \text{cons} :: (l : \text{list}_A^1)$ 
 $l \leftarrow \text{cons} \leftarrow x t =$ 
   $l.\text{cons}$  ;           %  $(x : A) (t : \text{list}_A^1) \stackrel{1}{\vdash} l : A \otimes \text{list}_A^1$ 
  send  $l x$  ;         %  $(t : \text{list}_A^1) \stackrel{0}{\vdash} l : \text{list}_A^1$ 
   $l \leftarrow t$ 

```

The resource cost of this process is 3. Sending the *cons* message has a cost of 2, one goes to the cost of sending and the other goes to the potential that needs to be sent with the message. There is another cost of 1 for sending the element. This resource usage is indeed confirmed by the potential of 3 in the type of the *cons* process.

C.1.3 *append*

The *append* process takes two lists l_1 and l_2 of type list_A and produces a list $l : \text{list}_A$ which behaves as the list l_2 appended to the list l_1 . To derive the type for *append*, we need to introduce two resource-aware session types.

$$\text{list}_A^2 = \oplus \{ \text{nil}^0 : \mathbf{1}, \\ \text{cons}^2 : A \otimes \text{list}_A^2 \}$$

$$\text{list}_A^0 = \oplus \{ \text{nil}^0 : \mathbf{1}^0, \\ \text{cons}^0 : A \otimes \text{list}_A^0 \}$$

The two types differ only in the potential associated with the *cons* message. The types list_A^2 and list_A^0 carry a potential of 2 and 0 with the *cons* message, respectively.

Figure 9 shows the implementation of the *append* process along with its type derivation. The implementation shows that the *append* process first cases analyzes on list l_1 . Either it receives the *cons* message followed by the element $x : A$, which are then passed along l . Or it receives the *nil* message in which case, it forwards l_2 along l .

The resource usage of the *append* process is twice the number of elements in l_1 . This becomes clear from analyzing the *cons* branch of the implementation of the *append* process. For every *cons* message it receives, it sends 2 messages along l , first the *cons* message followed by the element $x : A$. This is verified by its type. The list l_1 has type list_A^2 , which implies that the *cons* message carries a potential of 2 which is used to send the two messages.

C.2 Stacks

A stack is implemented as a sequence of *elem* processes terminated by an *empty* process. The implementation and type derivation of *elem* is presented in Figure 10.

The recursive *elem* process stores an element of the stack. It uses channel $x : A$ (element being stored) and channel $t : \text{stack}_A$ (tail of the stack) and provides service along $s : \text{stack}_A$. The implementation demonstrates that if the *elem* process receives an *ins* message along s , it receives the element y (line 4), spawns a new *elem* process using its original element x (line 5), and continues with another instance of the *elem* process with the received element y (line 6). In this way, it adds the element y to the head of the sequence. Otherwise, *elem* receives a *del* message along s and responds with the some label (line 7), followed by the channel x it stores (line 8). It then forwards all communication along s to t .

Inserting an element has no resource cost, since no messages are sent by the *elem* process. Similarly, deleting an element has a cost of 2, which is used to send two messages: the some label and the element x . This is reflected by the type stack_A , which needs 0 and 2 potential units for insertion and deletion, respectively, as indicated by the resource annotations.

The sequence of *elem* processes ends with an *empty* process, providing service along channel s where it can receive the label *ins* or *del*. If it receives the label *ins*, it receives the element y to be inserted (line 12), spawns a new *empty* process (line 13), and continues execution as an *elem* process with the received element (line 14). On receiving the label *del*, it just sends the none label (line 15) followed by the close message (line 16), indicating that the stack is empty.

Inserting an element sends no messages and thus has cost 0. Deleting an element sends two messages and has cost 2, which is reflected in the resource annotations of the labels in the type stack_A . Note that deleting an element requires the system to send back two messages, either the none label followed by the close message, or the some label followed by the element. Therefore, an implementation of stacks will have a resource cost of at least 2 for deletion. This shows that the above implementation is the most efficient w.r.t. our cost semantics because insertion has no resource cost, and deletion has the least possible cost.

C.3 Stack using one list

The stack interface introduced in Section 7 can be internally implemented using a single list. The list essentially acts as a stack, and insertion is equivalent to the *cons* operation, while deletion is equivalent to removal of the head of the list. First, we introduce the resource-aware types, both for the list and the stack.

```

1:  $(x:A) (t:\text{stack}_A) \Vdash^0 \text{elem} :: (s : \text{stack}_A)$ 
2:    $s \leftarrow \text{elem} \leftarrow x t =$ 
3:     case  $s$ 
4:        $(\text{ins} \Rightarrow y \leftarrow \text{recv } s ;$            %  $(y:A) (x:A) (t:\text{stack}_A) \Vdash^0 s : \text{stack}_A$ 
5:          $s' \leftarrow \text{elem} \leftarrow x t ;$      %  $(y:A) (s' : \text{stack}_A) \Vdash^0 s : \text{stack}_A$ 
6:          $s \leftarrow \text{elem} \leftarrow y s'$ 
7:       |  $\text{del} \Rightarrow s.\text{some} ;$                  %  $(x:A) (t:\text{stack}_A) \Vdash^1 s : A \otimes \text{stack}_A$ 
8:          $\text{send } s x ;$                          %  $t:\text{stack}_A \Vdash^0 s : \text{stack}_A$ 
9:          $s \leftarrow t)$ 

10:  $\cdot \Vdash^0 \text{empty} :: (s : \text{stack}_A)$ 
11:    $s \leftarrow \text{empty} =$ 
12:     case  $s$   $(\text{ins} \Rightarrow y \leftarrow \text{recv } s ;$    %  $(y:A) \Vdash^0 s : \text{stack}_A$ 
13:        $e \leftarrow \text{empty} ;$                    %  $(y:A) (e : \text{stack}_A) \Vdash^0 s : \text{stack}_A$ 
14:        $s \leftarrow \text{elem} \leftarrow y e$ 
15:     |  $\text{del} \Rightarrow s.\text{none} ;$                  %  $\cdot \Vdash^1 s : 1$ 
16:       close  $s)$ 

```

Figure 10. Implementation and type derivation of *elem* and *empty* processes for stack

```

1:  $\cdot \Vdash^4 \text{stack\_new} :: s : (\text{stack}_A^4)$ 
2:    $s \leftarrow \text{stack\_new} =$ 
3:      $e \leftarrow \text{nil} ;$             $(e : \text{list}_A^2) \Vdash^0 (s : \text{stack}_A^4)$ 
4:      $s \leftarrow \text{stack} \leftarrow e$ 

5:  $l : \text{list}_A^2 \Vdash^0 \text{stack} :: s : (\text{stack}_A^4)$ 
6:    $s \leftarrow \text{stack} \leftarrow l$ 
7:     case  $s$ 
8:        $(\text{ins} \Rightarrow x \leftarrow \text{recv } s ;$        %  $(x:A)(l : \text{list}_A^2) \Vdash^4 s : \text{stack}_A^2$ 
9:          $l' \leftarrow \text{cons} \leftarrow x l$      %  $l' : \text{list}_A^2 \Vdash^0 s : \text{stack}_A^4$ 
10:         $s \leftarrow \text{stack} \leftarrow l'$ 
11:      del  $\Rightarrow$  case  $l$ 
12:         $(\text{cons} \Rightarrow x \leftarrow \text{recv } l ;$      %  $(x:A)(l : \text{list}_A^2) \Vdash^2 s : \oplus\{\text{some}^0 : A \otimes \text{stack}_A^4, \text{none}^0 : 1^0\}$ 
13:           $s.\text{some} ;$                          %  $(x:A)(l : \text{list}_A^2) \Vdash^1 s : A \otimes \text{stack}_A^4$ 
14:           $\text{send } s x ;$                        %  $(l : \text{list}_A^2) \Vdash^0 s : \text{stack}_A^4$ 
15:           $s \leftarrow \text{stack} \leftarrow l$ 
16:        nil  $\Rightarrow$   $s.\text{none}$                      %  $(l : 1^0) \Vdash^1 s : 1^0$ 
17:          wait  $l ;$                            %  $\cdot \Vdash^1 s : 1^0$ 
18:          close  $s)$ 

```

Figure 11. Implementation of the *stack_new* and *stack* processes
$$\text{list}_A^2 = \oplus\{\text{nil}^2 : 1^0,$$

$$\text{cons}^2 : A \otimes \text{list}_A^2\}$$

$$\text{stack}_A^4 = \&\{\text{ins}^4 : A \multimap \text{stack}_A^4,$$

$$\text{del}^0 : \oplus\{\text{some}^0 : A \otimes \text{stack}_A^4,$$

$$\text{none}^0 : 1^0\}$$

The type for the list prescribes that the cons and nil messages carries a potential of 2, while all other messages carry no potential. The type for stack dictates that insertion costs a potential of 4, while deletion is cost-free. Intuitively, this represents an amortized analysis, we pay for the deletion of an element when it is being inserted. Hence, insertion has a cost of 2, while a potential of 2 is stored with the element, which is used for deletion.

We need 2 helper processes, namely the *nil* and *cons* to implement the stack. The type for these helper processes is different from the one introduced in Section C.1, and we present their implementation below.

```

 $\cdot \Vdash^4 \text{nil} :: (l : \text{list}_A^2)$ 
 $l \leftarrow \text{nil} =$ 
   $l.\text{nil} ;$            %  $\cdot \Vdash^1 l : 1^0$ 
  close  $l$            %  $\cdot \Vdash^0$ 

```

The *nil* process requires a potential of 4, of which 2 is used to pay for the cost, while 2 is used to pay for the potential associated with the nil message.

```

 $(x:A) (t : \text{list}_A^2) \Vdash^4 \text{cons} :: (l : \text{list}_A^2)$ 
 $l \leftarrow \text{cons} \leftarrow x t =$ 

```

$$\begin{array}{ll}
l.\text{cons} ; & \% (x : A) (t : \text{list}_A^2) \dagger l : A \otimes \text{list}_A^2 \\
\text{send } l \ x ; & \% (t : \text{list}_A^2) \dagger l : \text{list}_A^2 \\
l \leftarrow t &
\end{array}$$

The *cons* process also needs a potential of 4, and has a similar reasoning as the *nil* process. Finally, we implement the stack interface using two processes, *stack_new* and *stack*. The implementations and the type derivations are shown in Figure 11.

The *stack_new* process creates an empty stack. It uses an empty context and provides along a channel $s : \text{stack}_A^4$. It simply spawns a new *nil* process with channel $e : \text{list}_A^2$ and then continues as the *stack* process.

The *stack* process uses a list $l : \text{list}_A^2$ in its context and provides along $s : \text{stack}_A^2$. The list l acts as the internal list implementation of the stack. When the *stack* process receives the *ins* message followed by the element, it simply spawns a new *cons* process with the channel l' and continues with the new internal list l' . If the *stack* process receives the *del* message, it case analyzes on its internal list l . If the internal list is empty, the *stack* process simply sends the *none* message along s followed by closing the channels l and s . If l is non-empty, it simply receives the element at the head of the list, sends it along s and recurses back to the *stack* process.

The resource usage of the *stack* process is 4 times the number of insertions. There is a cost of 2 during insertion for spawning the *cons* process. Similarly, there is a cost of 2 for deletion, either for sending the *none* and close messages or the *some* message and the element. This is confirmed by the type $s : \text{stack}_A^4$ where insertion needs a potential of 4.

C.4 Queues

Similar to a stack, a queue is also implemented by a sequence of *elem* processes, connected via channels, and terminated by the *empty* process. Their implementations are presented in Figure 12. Similar to the implementation of a stack, the *elem* process provides along $s : \text{queue}_A$, stores the element $x : A$, and uses the tail of the queue $t : \text{queue}_A$. When the *elem* process receives the *ins* message along s , it receives the element y (line 4), and passes the *ins* message (line 5) along with y (line 5) to t . Since the process at the other end of t is also implemented using *elem*, it passes along the element to its tail too. Thus the element travels to the end of the queue where it is finally inserted. The deletion is similar to that for stack.

For each insertion, the *ins* label along with the element travels to the end of the queue. Hence, the resource cost of each insertion is $2n$ where n is the size of the queue and this is reflected in the type queue_A in the potential annotation of *ins* as $2n$. Similar to the stack, deletion has a resource cost of 2 to get back the *some* label and the element.

The implementation of *empty* process is identical to that of stacks. Since insertion does not cause the process to send any messages, its resource cost is 0. On the other hand, deletion costs 2 units because the process sends back the *none* label followed by the close message. This is correctly reflected in the queue type. Since $s : \text{queue}_A[0]$, the annotation for *ins* is $2n = 2 \cdot 0 = 0$. Similarly, *del* is annotated with a potential of 2.

C.5 Queue using two lists

A standard implementation of a queue can be done using two stacks, or two lists. The queue is internally represented using two lists, *in* and *out*. Insertions are performed to the *in* list, while deletions are

performed from the *out* list. If the *out* list is empty during deletion, the *in* list is reversed and a new empty list is created. The empty list acts as the new *in* list, while the reversed list acts as the new *out* list. We present the implementation in Figure 13. First, we describe the resource-aware types.

$$\begin{array}{l}
\text{list}_A^{(2,2)} = \oplus \{ \text{nil}^2 : \mathbf{1}^0, \\
\text{cons}^2 : A \otimes \text{list}_A^{(2,2)} \}
\end{array}$$

The type prescribes that the *nil* and *cons* messages carry a potential of 2, while all other messages carry no potential.

$$\begin{array}{l}
\text{list}_A^4 = \oplus \{ \text{nil}^0 : \mathbf{1}^0, \\
\text{cons}^4 : A \otimes \text{list}_A^4 \}
\end{array}$$

This type implies that only the *cons* message carries a potential of 4, while all other messages are free of potential.

$$\begin{array}{l}
\text{queue}_A^{(6,2)} = \& \{ \text{ins}^6 : A \multimap \text{queue}_A^{(6,2)}, \\
\text{del}^2 : \oplus \{ \text{some}^0 : A \otimes \text{queue}_A^{(6,2)}, \\
\text{none}^0 : \mathbf{1}^0 \} \}
\end{array}$$

The type of the queue entails that insertion requires a potential of 6, while deletion has a cost of 2. The rest of the messages don't carry any potential.

Once again, we need to use 3 helper processes, *nil*, *cons* and *rev*. We present the implementations of *nil* and *cons* while *rev* is omitted.

$$\begin{array}{l}
\cdot \dagger \text{nil} :: (l : \text{list}_A^4) \\
l \leftarrow \text{nil} = \\
\begin{array}{ll}
l.\text{nil} ; & \% \cdot \dagger l : \mathbf{1}^0 \\
\text{close } l & \% \cdot \dagger \cdot
\end{array}
\end{array}$$

The *nil* process requires a potential of 2, since the *nil* and close messages carry no potential.

$$\begin{array}{l}
(x : A) (t : \text{list}_A^4) \dagger \text{cons} :: (l : \text{list}_A^4) \\
l \leftarrow \text{cons} \leftarrow x \ t =
\end{array}$$

$$\begin{array}{ll}
l.\text{cons} ; & \% (x : A) (t : \text{list}_A^4) \dagger l : A \otimes \text{list}_A^4 \\
\text{send } l \ x ; & \% (t : \text{list}_A^4) \dagger l : \text{list}_A^4 \\
l \leftarrow t &
\end{array}$$

The *cons* process requires a potential of 6, of which 2 is used for sending the two messages, and 4 is required to pay for the potential of the *cons* message. The *rev* process we use in our implementation has the following type.

$$(\text{in} : \text{list}_A^4) \dagger \text{rev} :: (\text{out} : \text{list}_A^{(2,2)})$$

The main process *queue2* is implemented in Figure 13. If the process receives an *ins* message, it spawns a new process using the received element $x : A$ and $\text{in} : \text{list}_A^4$ and offers along $\text{in}' : \text{list}_A^4$. The *queue2* process then uses in' as its internal list. On the other hand, if the *del* message is received, it case analyzes on *out*. If the *out* list contains an element at its head, the *queue2* process removes this element from *out* and sends it along $s : \text{queue}_A^{(6,2)}$. If *out* is empty, then the *in* list is reversed by calling the *rev* process using $\text{in} : \text{list}_A^4$ and returning $\text{out}' : \text{list}_A^{(2,2)}$. If out' contains an element $x : A$ at its head, this element is removed and sent along $s : \text{queue}_A^{(6,2)}$ with the *some* message. If out' is empty, this means that the queue is empty, and *queue2* sends the *none* and close messages along channel s .

The amortized resource usage for this implementation is constant. Although the list can potentially be reversed during deletion, and reverse has a linear cost, the amortized cost for deletion is still

```

1: (x:A) (t:queueA[n-1]) ⊢0 elem :: (s : queueA[n])
2: s ← elem ← x t =
3:   case s (
4:     ins ⇒ y ← recv s ;           % (y:A) (x:A) (t:queueA[n-1]) ⊢2n s:queueA[n+1]
5:         t.ins ;                   % (y:A)(x:A)(t:A →0 queueA[n]) ⊢1 s:queueA[n+1]
6:         send t y ;               % (x : A) (t : queueA[n]) ⊢0 s : queueA[n+1]
7:         s ← elem ← x t

8:   | del ⇒ s.some ;              % (x:A)(t:queueA[n-1]) ⊢1 s:A ⊗ queueA[n-1]
9:         send s x ;              % t:queueA[n-1] ⊢0 s : queueA[n-1]
10:        s ← t)

11: · ⊢0 empty :: (s : queueA[0])
12: s ← empty =
13:   case s (
14:     ins ⇒ y ← recv s ;           % (y:A) ⊢0 s : queueA[1]
15:         e ← empty ;             % (y:A) (e : queueA[0]) ⊢0 s : queueA[1]
16:         s ← elem ← y e
17:   | del ⇒ s.none ;             % · ⊢1 s : 1
18:   close s)

```

Figure 12. Implementations and type derivations of *elem* and *empty* processes for queues

```

1: (in : listA4) (out : listA(2,2)) ⊢0 queue2 :: s : (queueA(6,2))
2: s ← queue2 ← l
3:   case s
4:     (ins ⇒ x ← recv s ;           % (x : A)(in : listA4) (out : listA(2,2)) ⊢6 s : queueA(6,2)
5:         in' ← cons ← x in        % (in' : listA4) (out : listA(2,2)) ⊢0 s : queueA(6,2)
6:         s ← queue2 ← in' out
7:     del ⇒ case out
8:         (cons ⇒ x ← recv out ;    % (x : A) (in : listA4) (out : listA(2,2)) ⊢2
9:         s : ⊕{some0 : A ⊗ queueA(6,2), none0 : 10}
10:        s.some ;                 % (x : A) (in : listA4) (out : listA(2,2)) ⊢1 s : A ⊗ queueA(6,2)
11:        send s x ;               % (in : listA4) (out : listA(2,2)) ⊢0 s : queueA(6,2)
12:        s ← queue2 ← in out
13:    nil ⇒ wait out ;             % (in : listA4) ⊢4 s : queueA(6,2)
14:        out' ← rev ← in ;        % (out' : listA(2,2)) ⊢2 s : queueA4
15:        case out'
16:          (cons ⇒ x ← recv out' ;  % (x : A) (out' : listA(2,2)) ⊢4
17:          s : ⊕{some0 : A ⊗ queueA(6,2), none0 : 10}
18:          s.some ;               % (x : A) (out' : listA(2,2)) ⊢3 s : A ⊗ queueA(6,2)
19:          send s x ;             % (out' : listA(2,2)) ⊢2 s : queueA4
20:          in' ← nil ;            % (in' : listA4) (out' : listA0) ⊢0 s : queueA(6,2)
21:          s ← queue2 ← in' out'

22:        nil ⇒ wait out' ;        % · ⊢4 s : ⊕{some0 : A ⊗ queueA(6,2), none0 : 10}
23:        s.none ;                % · ⊢3 s : 10
24:        close s)))

```

Figure 13. Implementation and type derivation of the *queue2* process

constant. This can be observed by providing extra potential during insertion, which is stored in the *in* list, and used to pay for reversal during deletion. Our type system supports amortized analysis, and indeed confirms the constant amortized cost since the insertion requires a potential of 6, while deletion requires a potential of 2, both being constants.

C.6 Higher Order Processes

We present the implementations and type derivations of the session-typed equivalent of two higher order processes, *map* and *fold*. We demonstrate that our type system can derive bounds for both these functions.

C.6.1 *map*

The *map* process is a session-typed implementation of the standard map function for lists. It takes a list of elements and a mapping function, and produces a list of mapped elements. First, we introduce the resource-aware session type of a mapping function.

$$\text{mapper}_{AB}^2 = \&\{ \text{next}^0 : A \overset{0}{\multimap} B \overset{0}{\multimap} \text{mapper}_{AB}^2, \\ \text{done}^0 : 1^2 \}$$

The mapper_{AB} is an external choice between *next* and *done*. If it receives the *next* message, it receives the element of type *A*. It then sends back an element of type *B* and recurses back to mapper_{AB} . Else, it receives the *done* message, and sends back the close message. To map an element $x : A$, we send the *next* label followed by x on the channel of type mapper_{AB} . It will then send back the element $y : B$. Regarding the potential, only the element of type *B* and close messages carry potential of 2 each, while all other messages carry no potential. We also present the types for the two lists.

$$\text{list}_A^2 = \oplus\{ \text{nil}^1 : 1^0, \\ \text{cons}^2 : A \overset{0}{\otimes} \text{list}_A^2 \}$$

$$\text{list}_B^0 = \oplus\{ \text{nil}^0 : 1^0, \\ \text{cons}^0 : A \overset{0}{\otimes} \text{list}_B^0 \}$$

The input list list_A contains a potential of 2 on the *cons* label and a potential of 1 on the *nil* label. The output list list_B contains no potential on any message.

Figure 14 presents the implementation of the *map* process. It uses $l : \text{list}_A^2$ and $m : \text{mapper}_A^2$ in its context, and offers along $k : \text{list}_A^0$, where l is the input list and k is the output list. It starts with a case analysis on l . If it receives the *cons* message, it receives the element $x : A$, sends the *next* message and the element x along channel m . It then receives the element $y : B$ along m . Finally, it sends the *cons* label and element y along channel k and recurses back to the *map* process. If it receives the *nil* message along l , it waits for the channel l to close. It then sends the *done* message to m and waits for it to close. Finally, it sends the *nil* and close messages along k and terminates. This indicates that the input list l has been completely processed.

The resource usage of *map* function is $4n + 3$, where n is the number of elements in the input list l . This can be observed by analyzing the two branches in the implementation. In the *cons* branch, the process sends 4 messages, two along m and two along k . In the *nil* branch, it sends 3 messages, one along m and two along k . This is indeed certified by the type of the *map* process, which receives a potential of 2 from the list l with every *cons* message, and a potential of 2 from the mapper m with the element $y : B$.

Also, it receives a potential of 1 from l with the *nil* message and a potential of 2 from the mapper m . Hence, the potential needed is exactly equal to the work done by the *map* process.

C.6.2 *fold*

The *fold* process is the session-typed equivalent of the standard fold function. Again, we introduce the resource-aware type for the folding function.

$$\text{folder}_{AB}^0 = \&\{ \text{next}^0 : A \overset{0}{\multimap} B \overset{0}{\multimap} \text{folder}_{AB}^0, \\ \text{done}^0 : 1^0 \}$$

Similar to the mapper, folder_{AB} expects to receive either *next* or *done* labels. If it receives the *next* label, it receives the element of type *A* and two elements of type *B*. It then produces an element of type *B* and recurses back to folder_{AB} . If it receives the *done* label, it sends back the close message and terminates. We present the resource-aware type for the input list.

$$\text{list}_A^3 = \oplus\{ \text{nil}^1 : 1^0, \\ \text{cons}^3 : A \overset{0}{\otimes} \text{list}_A^2 \}$$

The *cons* message carries a potential of 3 and the *nil* message carries a potential of 1. All other messages carry no potential.

Figure 15 presents the implementation of the *fold* process. It takes the list $l : \text{list}_A^3$, the folder $m : \text{folder}_{AB}^0$ and the initial input $b : B$ in its context, and provides service along $r : B$. It starts with a case analysis on l . If it receives the *cons* message, it receives the element $x : A$. It then sends the *next* message, followed by $x : A$ and $b : B$ to the channel $m : \text{folder}_{AB}^0$. It then receives the element $y : B$ along m and recurses back to the *fold* process. If it receives the *nil* label along l , it waits for the channel to close, it then sends the *done* label to m and waits for it to close. Finally, it simply forwards b along r . This indicates that the list has been folded and the result is stored in $b : B$.

The resource usage of the *fold* process is $3n + 1$, where n is the number of elements in the list l . This can be observed from the implementation where the process sends 3 messages in the *cons* branch, and sends one message in the *nil* branch. This resource cost is affirmed by the type system, which associates a potential of 3 with the *cons* label, and a potential of 1 in the *nil* label.

```

1:  $(l : \text{list}_A^2) (m : \text{mapper}_{AB}^2) \vdash \text{map} :: (k : \text{list}_A^0)$ 
2:  $k \leftarrow \text{map} \leftarrow l \ m =$ 
3:   case  $l$ 
4:      $(\text{cons} \Rightarrow x \leftarrow \text{recv } l ;$            %  $(x : A) (l : \text{list}_A^2) (m : \text{mapper}_{AB}^2) \vdash^2 (k : \text{list}_B^0)$ 
5:          $m.\text{next} ;$                          %  $(x : A) (l : \text{list}_A^2) (m : A \multimap B \otimes \text{mapper}_{AB}^2) \vdash^1 (k : \text{list}_B^0)$ 
6:          $\text{send } m \ x ;$                        %  $(l : \text{list}_A^2) (m : B \otimes \text{mapper}_{AB}^2) \vdash^0 (k : \text{list}_B^0)$ 
7:          $y \leftarrow \text{recv } m ;$              %  $(l : \text{list}_A^2) (y : B) (m : \text{mapper}_{AB}^2) \vdash^2 (k : \text{list}_B^0)$ 
8:          $k.\text{cons} ;$                          %  $(l : \text{list}_A^2) (y : B) (m : \text{mapper}_{AB}^2) \vdash^1 (k : B \otimes \text{list}_B^0)$ 
9:          $\text{send } k \ y ;$                      %  $(l : \text{list}_A^2) (m : \text{mapper}_{AB}^2) \vdash^0 (k : \text{list}_B^0)$ 
10:         $k \leftarrow \text{map} \leftarrow l \ m$ 
11:   nil  $\Rightarrow$   $\text{wait } l ;$                        %  $(m : \text{mapper}_{AB}^2) \vdash^1 (k : \text{list}_B^0)$ 
12:         $m.\text{done} ;$                          %  $(m : \mathbf{1}^2) \vdash^0 (k : \text{list}_B^0)$ 
13:         $\text{wait } m ;$                          %  $\cdot \vdash^2 (k : \text{list}_B^0)$ 
14:         $k.\text{nil} ;$                           %  $\cdot \vdash^1 (k : \mathbf{1}^0)$ 
15:        close  $k$ )

```

Figure 14. Implementation and type derivation of *map*

```

1:  $(l : \text{list}_A^3) (m : \text{folder}_{AB}^0) (b : B) \vdash \text{fold} :: (r : B)$ 
2:  $r \leftarrow \text{fold} \leftarrow l \ m \ b =$ 
3:   case  $l$ 
4:      $(\text{cons} \Rightarrow x \leftarrow \text{recv } l ;$            %  $(x : A) (l : \text{list}_A^2) (m : \text{folder}_{AB}^0) (b : B) \vdash^3 (r : B)$ 
5:          $m.\text{next} ;$                          %  $(x : A) (l : \text{list}_A^3) (m : A \multimap B \multimap B \otimes \text{folder}_{AB}^0) (b : B) \vdash^2 (r : B)$ 
6:          $\text{send } m \ x ;$                        %  $(l : \text{list}_A^3) (m : B \multimap B \otimes \text{folder}_{AB}^0) (b : B) \vdash^1 (r : B)$ 
7:          $\text{send } m \ b ;$                        %  $(l : \text{list}_A^3) (m : B \otimes \text{folder}_{AB}^0) \vdash^0 (r : B)$ 
8:          $y \leftarrow \text{recv } m ;$              %  $(l : \text{list}_A^2) (y : B) (m : \text{folder}_{AB}^0) \vdash^2 (r : B)$ 
9:          $r \leftarrow \text{fold} \leftarrow l \ m \ y$ 
10:   nil  $\Rightarrow$   $\text{wait } l ;$                        %  $(m : \text{folder}_{AB}^0) (b : B) \vdash^1 (r : B)$ 
11:         $m.\text{done} ;$                          %  $(m : \mathbf{1}^0) (b : B) \vdash^0 (r : B)$ 
12:         $\text{wait } m ;$                          %  $(b : B) \vdash^0 (r : B)$ 
13:         $r \leftarrow b$ )

```

Figure 15. Implementation and type derivation of *fold*