

# Polarized Substructural Session Types

Frank Pfenning<sup>1</sup> and Dennis Griffith<sup>2</sup>

<sup>1</sup> Carnegie Mellon University, Pittsburgh, PA 15213, USA  
fp@cs.cmu.edu

<sup>2</sup> University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA  
dgriffi3@illinois.edu

**Abstract.** The deep connection between session-typed concurrency and linear logic is embodied in the language SILL that integrates functional and message-passing concurrent programming. The exacting nature of linear typing provides strong guarantees, such as global progress, absence of deadlock, and race freedom, but it also requires explicit resource management by the programmer. This burden is alleviated in an affine type system where resources need not be used, relying on a simple form of garbage collection.

In this paper we show how to effectively support both linear and affine typing in a single language, in addition to the already present unrestricted (intuitionistic) types. The approach, based on Benton's adjoint construction, suggests that the usual distinction between synchronous and asynchronous communication can be viewed through the lens of modal logic. We show how polarizing the propositions into positive and negative connectives allows us to elegantly express synchronization in the type instead of encoding it by extra-logical means.

## 1 Introduction

Session types prescribe the communication behavior of concurrent message-passing processes [13,14]. Anticipated with some analogies for some time [11,23], session types have recently been placed upon the firm foundation of linear logic via a Curry-Howard interpretation of linear propositions as types, proofs as processes, and cut reduction as communication. Variations apply for both intuitionistic [5,6] and classical [24] linear logic. This has enabled the application of proof-theoretic techniques in this domain, for example, developing logical relations [17], corecursion [22], and parametricity and behavioral polymorphism [4]. It has also given rise to the design of SILL, a modular extension of an underlying functional language with session-typed concurrency [21].

Practical experience with a SILL prototype has led to a number of new questions. For example, should the type system really be *linear*, where all resources must be fully accounted for by the programmer, or should it be *affine* [16], where resources may be reclaimed by a form of garbage collection? Another question concerns the underlying model of communication: should it be synchronous or asynchronous? The proof theory does not provide a definitive answer to this question, supporting both. The purpose of this paper is to show that we can

have our cake and eat it, too, in both cases. First, we combine linear and affine types in an elegant and proof-theoretically justified way, slightly reformulating unrestricted types along the way. Second, we show how to support synchronous and asynchronous communication patterns in a single language, again taking our inspiration from proof theory.

The central idea behind the first step is to generalize Benton’s LNL [3] in the spirit of Reed’s adjoint logic [18]. This stratifies the propositions into linear, affine, and unrestricted ones, with modal operators *shifting* between the strata. For example, the familiar exponential of linear logic  $!A$  is decomposed into two shifting modalities, one going from  $A$  (which is linear) into the unrestricted layer, and one going from the unrestricted layer back to the linear one. Similar modalities connect the linear and affine layers of the language.

The main idea behind the second step is to *polarize* the presentation of linear logic [15], segregating positive (sending) connectives from negative (receiving) connectives. Surprisingly, the two sublanguages of propositions can be connected by new versions of the shift modalities, fully consistent with the adjoint construction, leading to a pleasantly coherent language.

In the rest of this note we walk through these steps, taking small liberties with previously published notations for the sake of consistency.

## 2 Linear Logic and Session Types

We give here only the briefest review of linear logic and its deep connection to session types. The interested reader is referred to [5,6,21] for further background.

The key idea of linear logic [12] is to view logical propositions as resources: they must be used exactly once in a proof. We adopt the intuitionistic version [2], which is defined via a *linear hypothetical judgment* [8]

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

where the hypotheses  $A_1, \dots, A_n$  must be used exactly once in the proof of the conclusion  $A$ . We do not care about the order of the assumptions, treating them like a multiset, and use  $\Delta$  to denote such a multiset. The judgmental rules (sometimes called structural rules) explain the meaning of the hypothetical judgment itself and are independent of any particular propositions. In a sequent calculus, there are two such rules: cut, which states that if we can prove  $A$  we are justified to use  $A$  as a resource, and identity, which says that we can use a resource  $A$  to prove  $A$ .

$$\frac{\Delta \vdash A \quad \Delta', A \vdash C}{\Delta, \Delta' \vdash C} \text{ cut} \qquad \frac{}{A \vdash A} \text{ id}$$

Under the Curry-Howard isomorphism for intuitionistic logic, propositions are related to types, proofs to programs, and proof reduction to computation. Here, linear logic propositions are related to session types, proofs to concurrent programs, and cut reduction in proofs to computation. For this correspondence,

each hypothesis is labeled by a *channel* (rather than a variable). In addition, we also label the conclusion by a channel. This is because, unlike functional programming, we do not reduce a process to a value but we interact with it. For such interaction to take place in the concurrent setting, we need a channel to communicate along.

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

Here,  $x_1, \dots, x_n$  and  $x$  are distinct channels, and  $A_1, \dots, A_n$  and  $A$  are their respective session types. We say that process  $P$  *provides*  $A$  along channel  $x$  and *uses* channels  $x_1, \dots, x_n$ .

The rule of cut now is a form of *process composition*, connecting a client (here  $Q$ ) to a provider (here  $P$ ).

$$\frac{\Delta \vdash P_x :: (x : A) \quad \Delta, x:A \vdash Q_x :: (z : C)}{\Delta, \Delta' \vdash (x \leftarrow P_x ; Q_x) :: (z : C)} \text{ cut}$$

We use syntactic forms for processes, rather than  $\pi$ -calculus terms, to emphasize the interpretation of proofs as programs. Because every (well-typed) process  $P$  offers a session along exactly one channel, and each channel is provided by exactly one process, we can think of channels as unique process identifiers. Under this interpretation, suggested by the intuitionistic formulation of linear logic, we can see that the cut rule *spawns*  $P$  as a new process. More precisely, the process identified by  $z$  executing  $(x \leftarrow P_x ; Q_x)$  creates a fresh channel  $a$ , spawns a process executing  $P_a$  that provides session  $A$  along  $a$ , and continues as  $Q_a$ . Because  $a$  is fresh, this channel will be a private channel between  $P_a$  and  $Q_a$ .

We can express this in a *substructural operational semantics* [19] which is based on *multiset rewriting* [7]. The notation is again borrowed from linear logic, but it should not be confused with the use of linear logic propositions as session types.

$$\text{cut} : \text{proc}_c(x \leftarrow P_x ; Q_x) \multimap \{\exists a. \text{proc}_a(P_a) \otimes \text{proc}_c(Q_a)\}$$

In this formalization  $\text{proc}_c(P)$  is the state of a process executing program  $P$ , offering along channel  $c$ . The multiplicative conjunction ( $\otimes$ ) combines processes in the same state, linear implication ( $\multimap$ ) expresses a state transition from left to right, and the existential quantification corresponds to generation of a fresh channel. The curly braces  $\{\dots\}$  indicate a monad which essentially forces the rule above to be interpreted as a multiset rewriting rule.

The identity rule instead *forwards* between its client and the process that it uses, which must be of the same type.

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

There are several ways to describe this action operationally. A straightforward one globally identifies the channels  $x$  and  $y$ , while the forwarding process itself terminates.

$$\text{id} : \text{proc}_c(c \leftarrow d) \multimap \{c = d\}$$

This could be implemented in the substrate of the network or operating system. Or it could be implemented more explicitly by sending a message along  $c$  asking the client to use  $d$  for subsequent interactions. For now, we abstract over such lower level details.

Assigning process expressions to each rule of linear logic yields the following interpretation of propositions.

$A, B, C ::= \mathbf{1}$	send end and terminate
$A \otimes B$	send channel of type $A$ and continue as $B$
$A \oplus B$	send inl or inr and continue as $A$ or $B$ , respectively
$\tau \wedge B$	send value $v$ of type $\tau$ and continue as $B$
$A \multimap B$	receive channel of type $A$
$A \& B$	receive inl or inr and continue as $A$ or $B$ , respectively
$\tau \supset B$	receive value $V$ of type $\tau$ and continue as $B$

Here, we wrote  $\tau \wedge B$  as a special case of  $\exists x:\tau. B$  where  $x$  does not appear in  $B$ , and  $\tau \supset B$  is a special case of  $\forall x:\tau. B$ . The syntactic simplification is justified because in this paper we do not consider propositions that depend on terms.

Below is a summary of the process expressions, with the sending construct followed by the matching receiving construct. For the purpose of the examples we generalize the binary choice constructs  $A \& B$  and  $A \oplus B$  to n-ary choice  $\&\{lab_i : A_i\}_i$  and  $\oplus\{lab_i : A_i\}_i$ , respectively. We have as a special case  $A \& B = \&\{inl : A, inr : B\}$  and  $A \oplus B = \oplus\{inl : A, inr : B\}$ .

$P, Q, R ::= x \leftarrow P_x ; Q_x$	cut (spawn)
$c \leftarrow d$	id (forward)
close $c$   wait $c$	$\mathbf{1}$
send $c$ ( $y \leftarrow P_y$ ) ; $Q$   $x \leftarrow \text{recv } c ; R_x$	$A \otimes B, A \multimap B$
send $c$ $d$	derived form $A \otimes B, A \multimap B$
send $c$ $M$ ; $P$   $n \leftarrow \text{recv } c ; Q_n$	$A \wedge B, A \supset B$
$c.lab ; P$   case $c$ $\{lab_i \rightarrow Q_i\}_i$	$\&\{lab_i : A_i\}_i, \oplus\{lab_i : A_i\}_i$

As a running example in this paper we will use variations of an implementation of polymorphic queues. We begin with the purely linear version. The interface specifies that a queue presents an external choice between enqueue and dequeue operations. When the client selects to enqueue, we input a channel of type  $A$  (to be stored in the queue), and recurse. When the client selects to dequeue, we either indicate that the queue is empty and terminate, or we indicate that there is some element in the queue, send the first element (removing it in the process), and recurse. The “recursion” here is an instance of an *equirecursive session type* [10]; some logical underpinnings are available for *coinductive types* [22]. We also use polymorphism intuitively; a formal development can be found in [4].

First, the specification of the queue interface.

$$\text{queue } A = \&\{\text{enq} : A \multimap \text{queue } A, \text{deq} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \otimes \text{queue } A\}\}$$

We implement queues with two forms of recursive processes, *empty* for the empty queue and *elem* for a process holding exactly one element. For processes with

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

we write  $P : \{A \leftarrow A_1, \dots, A_n\}$  to specify its typing and  $x \leftarrow P \leftarrow x_1, \dots, x_n$  to provide its interface.

$$\begin{array}{ll} \text{empty} : \{\text{queue } A\} & \text{elem} : \{\text{queue } A \leftarrow A, \text{queue } A\} \\ c \leftarrow \text{empty} = & c \leftarrow \text{elem} \leftarrow x, d = \\ \text{case } c \text{ of} & \text{case } c \text{ of} \\ | \text{enq} \rightarrow x \leftarrow \text{recv } c ; & | \text{enq} \rightarrow y \leftarrow \text{recv } c ; \\ \quad e \leftarrow \text{empty} ; & \quad d.\text{enq} ; \text{send } d y ; \\ \quad c \leftarrow \text{elem} \leftarrow x, e & \quad c \leftarrow \text{elem} \leftarrow x, d \\ | \text{deq} \rightarrow c.\text{none} ; & | \text{deq} \rightarrow c.\text{some} ; \\ \quad \text{close } c & \quad \text{send } c x ; \\ & \quad c \leftarrow d \end{array}$$

From the perspective of the client, this implementation has constant time enqueue and dequeue operations. For dequeue this is obvious. For enqueue, the process at the front of the queue passes the element down the queue and is immediately available to serve another request while the element travels to the end of the queue.

### 3 Categorical Truth

The linear logic proposition  $!A$  allows  $A$  to be used arbitrarily often in a proof—it functions as an unrestricted resource. In the intuitionistic reconstruction of linear logic [8],  $!A$  internalizes a *categorical judgment*. We say that  $A$  is *valid* if it is true, and its proof does not depend on any assumptions about the truth of other propositions. Since we are working with a linear hypothetical judgment, this means that the proof of  $A$  does not depend on any resources. We further allow hypotheses  $\Gamma$  that are assumed to be valid (rather than merely true), and these are allowed in a proof  $A$  *valid*.

$$\Gamma ; \Delta \vdash C$$

The meaning of validity is captured in the following two judgmental rules, where ‘ $\cdot$ ’ stands for an empty context:

$$\frac{\Gamma ; \cdot \vdash A \quad (\Gamma, A) ; \Delta \vdash C}{\Gamma ; \Delta \vdash C} \text{cut!} \qquad \frac{(\Gamma, A) ; \Delta, A \vdash C}{(\Gamma, A) ; \Delta \vdash C} \text{copy}$$

The first, *cut!*, states that we are justified in assuming that  $A$  is valid if we can prove it without using any resources. The second, *copy*, states that we are justified in assuming a copy of the resource  $A$  if  $A$  is known to be valid. All the

purely linear rules are generalized by adding an unrestricted context  $\Gamma$  which is propagated to all premises.

How do we think of these in terms of processes? We introduce a new form of channel, called a *shared channel* (denoted by  $u, w$ ) which can be used arbitrarily often in a client, and by arbitrarily many clients. It is offered by a *persistent process*. Operationally, a persistent process offering along  $w : A$  inputs a fresh linear channel  $c$  and spawns a new process  $P$  that offers  $A$  along  $c$ .

We have the following typing rules, first at the level of judgments.

$$\frac{\Gamma, u:A ; \Delta, x:A \vdash P_x :: (z:C)}{\Gamma, u:A ; \Delta \vdash (x \leftarrow \text{send } u ; P_x) :: (z:C)} \text{ copy}$$

$$\frac{\Gamma ; \cdot \vdash P_y :: (y:A) \quad \Gamma, u:A ; \Delta \vdash Q_u :: (z:C)}{\Gamma ; \Delta \vdash (u \leftarrow !(y \leftarrow \text{recv } u ; P_y) ; Q_u) :: (z:C)} \text{ cut!}$$

The copy rule has a slightly strange process expression,

$$x \leftarrow \text{send } u ; P_x$$

It expresses that we send a *new* channel  $x$  along  $u$ . The continuation  $P$  refers to  $x$  so it can communicate along this new channel. This pattern will be common for sending fresh channels in a variety of constructs in this paper.

We see that the cut! rule incorporates two steps: creating a new shared channel  $u$  and then immediately receiving a linear channel  $y$  along  $u$ . There is no simple way to avoid this, since  $P$  in the first premise offers along a linear channel  $y$ . We will see alternatives in later sections.

In the operational semantics we write  $!\text{proc}_w(P)$  for a persistent process, offering along shared channel  $w$ . In the language of substructural specification,  $!\text{proc}_w(P)$  on the left-hand side of a rule means that it has to match a persistent proposition. We therefore do not need to repeat it on the right-hand side: it will continue to appear in the state. In this notation, the operational semantics is as follows:

$$\begin{aligned} \text{copy} : & \text{!proc}_w(y \leftarrow \text{recv } w ; P_y) \otimes \text{proc}_c(x \leftarrow \text{send } w ; Q_x) \\ & \multimap \{ \exists a. \text{proc}_a(P_a) \otimes \text{proc}_c(Q_a) \} \\ \text{cut!} : & \text{proc}_c(u \leftarrow !(y \leftarrow \text{recv } u ; P_y) ; Q_u) \\ & \multimap \{ \exists w. \text{!proc}_w(y \leftarrow \text{recv } w ; P_y) \otimes \text{proc}_c(Q_w) \} \end{aligned}$$

The validity judgment realized by persistent processes offering along unrestricted channels can be internalized as a proposition  $!A$  with the following rules. Note that the linear context must be empty in the  $!R$  rule, since validity is a categorical judgment. Allowing dependence on linear channels would violate their linearity.

$$\frac{\Gamma ; \cdot \vdash P_y :: (y:A)}{\Gamma ; \cdot \vdash (u \leftarrow \text{send } x ; !(y \leftarrow \text{recv } u ; P_y)) :: (x:!A)} !R$$

$$\frac{\Gamma, u:A ; \Delta \vdash Q_u :: (z:C)}{\Gamma ; \Delta, x:!A \vdash (u \leftarrow \text{recv } x ; Q_u) :: (z:C)} !L$$

Again the !R rule combines two steps: sending a new persistent channel  $u$  along  $x$  and then receiving a linear channel  $y$  along  $u$ . Operationally:

$$\text{bang} : \text{proc}_c(u \leftarrow \text{recv } a ; Q_u) \otimes \text{proc}_a(u \leftarrow \text{send } a ; !(y \leftarrow \text{recv } u ; P_y)) \\ \rightarrow \{\exists w. \text{proc}_c(Q_w) \otimes !\text{proc}_w(y \leftarrow \text{recv } w ; P_y)\}$$

As expected, the persistent process spawned by the `bang` computation rule has exactly the same form as the one spawned by `cut!`, because a linear cut for a proposition  $!A$  becomes a persistent cut for a proposition  $A$ .

Let's analyze the two-step rule in more detail.

$$\frac{\Gamma ; \cdot \vdash P_y :: (y:A)}{\Gamma ; \cdot \vdash (u \leftarrow \text{send } x !(y \leftarrow \text{recv } u ; P_y)) :: (x:!A)} \text{!R}$$

The judgment  $A$  *valid* (corresponding to an unrestricted hypothesis  $u:A$ ) is elided on the right-hand side: we jump directly from the truth of  $!A$  to the truth of  $A$ . Writing it out as an intermediate step appears entirely reasonable. We do not even mention the linear hypotheses in the intermediate step, since the validity of  $A$  depends only on assumptions of validity in  $\Gamma$ .

$$\frac{\frac{\Gamma ; \cdot \vdash P_y :: (y:A)}{\Gamma \vdash (y \leftarrow \text{recv } u ; P_y) :: (u:A)} \text{valid}}{\Gamma ; \cdot \vdash (u \leftarrow \text{send } x ; !(y \leftarrow \text{recv } u ; P_y)) :: (x:!A)} \text{!R}$$

We emphasize that  $!A$  is positive (in the sense of polarized logic), so it corresponds to a `send`, while  $A$  *valid* is negative as a judgment, so it correspond to a `receive`. In the next section we elevate this from a judgmental to a first-class logical step.

Revisiting the example, recall that if we are the client of a channel  $c : \text{queue } A$ , we must use this channel. This means we have to explicitly dequeue all its elements. In fact, we have to explicitly consume each of the elements as well, since they are also linear. However, if we know that each element in the queue is in fact unrestricted, we can destroy it recursively with the following program.

`destroy` :  $\{\mathbf{1} \leftarrow \text{queue } (!A)\}$

$c \leftarrow \text{destroy} \leftarrow q =$

```

q.deq ;
case q of
| none → wait q ; close c
| some → x ← recv q ;           % obtain element x
        u ← recv x ;           % receive shared channel u, using x
        c ← destroy ← q       % recurse, ignoring u

```

## 4 Adjoint Logic

Adjoint logic is based on the idea that instead of a modality like  $!A$  that remains within a given language of propositions, we have two mutually dependent languages and *two* modalities going back and forth between them. For this to make

sense, the operators have to satisfy certain properties that pertain to the semantics of the two languages. We have in fact three language layers, which we call *linear propositions*  $A_L$ , *affine propositions*  $A_F$ , and *unrestricted propositions*  $A_U$ . They are characterized by the structural properties they satisfy: linear propositions are subject to none (they must be used exactly once), affine proposition can be weakened (they can be used at most once), and unrestricted propositions can be contracted and weakened (they can be used arbitrarily often). The order of propositions in the context matters for none of them. The hierarchy of structural properties is reflected in a hierarchy of modes of truth:

$$U > F > L$$

$U$  is *stronger* than  $F$  in the sense that unrestricted hypotheses can be used to prove affine conclusions, but not vice versa, and similarly for the other relations. Contexts  $\Psi$  combine assumptions with all modes. We write  $\geq$  for the reflexive and transitive closure of  $>$  and define

$$\Psi \geq k \quad \text{if } m \geq k \text{ for every } B_m \text{ in } \Psi$$

and

$$\Psi \vdash A_k \quad \text{presupposes } \Psi \geq k$$

We use the notation  $\uparrow_k^m A_k$  for an operator going from mode  $k$  *up* to mode  $m$ , and  $\downarrow_k^m A_m$  for an operator going *down* from mode  $m$  to mode  $k$ . In both cases we presuppose  $m > k$ .

Taking this approach we obtain the following language:

$$\begin{array}{ll} \text{Modes} & m, k, r ::= U \mid F \mid L \\ \text{Propositions} & A_m, B_m ::= \mathbf{1}_m \mid A_m \otimes_m B_m \mid A_m \oplus_m B_m \mid \tau \wedge_m B_m \\ & \mid A_m \multimap_m B_m \mid A_m \&_m B_m \mid \tau \supset_m B_m \\ & \mid \uparrow_k^m A_k \quad (m > k) \\ & \mid \downarrow_m^r A_r \quad (r > m) \end{array}$$

Because both  $!A$  and  $A$  are linear propositions, the exponential  $!A$  decomposes into two modalities:

$$!A = \downarrow_L^U \uparrow_L^U A_L$$

Because linear and affine propositions behave essentially the same way except that affine channels need not be used, we reuse all the same syntax (both for propositions and for process expressions) at these two layers. Unrestricted propositions would behave quite differently in ways that are outside the scope of this note, so we specify that there are no unrestricted propositions besides  $\uparrow_L^U A_L$  and  $\uparrow_F^U A_F$ .

In the following logical rules we always presuppose that the sequent in the conclusion is well-formed and add enough conditions to verify the presupposition



in the premises.

$$\frac{\Psi \vdash A_k}{\Psi \vdash \uparrow_k^m A_k} \uparrow R \qquad \frac{k \geq r \quad \Psi, A_k \vdash C_r}{\Psi, \uparrow_k^m A_k \vdash C_r} \uparrow L$$

$$\frac{\Psi_{\geq m} \vdash A_m}{\Psi \vdash \downarrow_k^m A_m} \downarrow R \qquad \frac{\Psi, A_m \vdash C_r}{\Psi, \downarrow_k^m A_m \vdash C_r} \downarrow L$$

Here  $\Psi_{\geq m}$  is the restriction of  $\Psi$  to propositions  $A_k$  with  $k \geq m$ . The rule does not apply if this would erase a linear proposition  $A_L$  since only affine and unrestricted propositions are subject to weakening.

The rules with no condition on the modes are invertible, while the others are not invertible. This means  $\uparrow A$  is *negative* while  $\downarrow A$  is *positive* (in the terminology of polarized logic [15]). We already noted that processes offering a negative type receive, while processes offering a positive type send. But what do we send or receive? Thinking of channels as intrinsically linear, affine, or shared suggests that we should send and receive fresh channels of different modes. Following this reasoning we obtain:

$$\frac{\Psi \vdash P_{x_k} :: (x_k : A_k)}{\Psi \vdash (x_k \leftarrow \text{recv } x_m ; P_{x_k}) :: (x_m : \uparrow_k^m A_k)} \uparrow R$$

$$\frac{k \geq r \quad \Psi, x_k : A_k \vdash Q_{x_k} :: (z_r : C_r)}{\Psi, x_m : \uparrow_k^m A_k \vdash (x_k \leftarrow \text{send } x_m ; Q_{x_k}) :: (z_r : C_r)} \uparrow L$$

For clarity, we annotate each channel with its mode, although it may not be strictly necessary. Operationally:

$$\begin{aligned} \text{up}_k^m : \text{proc}_{a_r}(x_k \leftarrow \text{send } c_m ; Q_{x_k}) \otimes \text{proc}_{c_m}(y_k \leftarrow \text{recv } c_m ; P_{y_k}) \\ \multimap \{ \exists c_k. \text{proc}_{a_r}(Q_{c_k}) \otimes \text{proc}_{c_k}(P_{c_k}) \} \end{aligned}$$

And for the other modality:

$$\frac{\Psi \geq m \quad \Psi \vdash Q_{x_m} :: (x_m : A_m)}{\Psi \vdash (x_m \leftarrow \text{send } x_k ; Q_{x_m}) :: (x_k : \downarrow_k^m A_m)} \downarrow R$$

$$\frac{\Psi, x_m : A_m \vdash P_{x_m} :: (z_r : C_r)}{\Psi, x_k : \downarrow_k^m A_m \vdash (x_m \leftarrow \text{recv } x_k ; P_{x_m}) :: (z_r : C_r)} \downarrow L$$

Operationally:

$$\begin{aligned} \text{down}_k^m : \text{proc}_{a_r}(y_m \leftarrow \text{recv } c_k ; P_{y_m}) \otimes \text{proc}_{c_k}(x_m \leftarrow \text{send } c_k ; Q_{x_m}) \\ \multimap \{ \exists c_m. \text{proc}_{a_r}(P_{c_m}) \otimes \text{proc}_{c_m}(Q_{c_m}) \} \end{aligned}$$

Since processes offering along unrestricted channels are persistent, we use here the (admittedly dangerous) notational convention that all processes offering along unrestricted channels  $c_U$  are implicitly marked persistent. In particular,

we should read  $\text{up}_k^U$  and  $\text{down}_k^U$  as

$$\begin{aligned} \text{up}_k^U &: \text{proc}_{a_r}(x_k \leftarrow \text{send } c_U ; Q_{x_k}) \otimes !\text{proc}_{c_U}(x_k \leftarrow \text{recv } c_U ; P_{x_k}) \\ &\quad \multimap \{ \exists c_k. \text{proc}_{a_r}(Q_{c_k}) \otimes \text{proc}_{c_k}(P_{c_k}) \} \\ \text{down}_k^U &: \text{proc}_{a_r}(x_U \leftarrow \text{recv } c_k ; P_{x_U}) \otimes \text{proc}_{c_k}(x_U \leftarrow \text{send } c_k ; Q_{x_U}) \\ &\quad \multimap \{ \exists c_U. \text{proc}_{a_r}(P_{c_U}) \otimes !\text{proc}_{c_U}(Q_{c_U}) \} \end{aligned}$$

At this point we have achieved that every logical connective, including the up and down modalities, correspond to exactly one matching send and receive action. Moreover, as we can check, the compound rules for  $!A$  decompose into individual steps.

Returning to our example, we can now specify that our queue is supposed to be affine, that is, that we can decide to ignore it. We annotate defined types and type variables with their mode (U, F, or L), but we overload the logical connectives since their meanings, when defined, are consistent. The elements of an affine queue should also be affine. If we make them linear, as in

$$\begin{aligned} \text{queue}_F A_L &= \&\{ \text{enq} : \uparrow_L^F A_L \multimap \text{queue}_F A_L, \\ &\quad \text{deq} : \{ \text{none} : \mathbf{1}, \text{some} : \uparrow_L^F A_L \otimes \text{queue}_F A_L \} \} \end{aligned}$$

then we could never use  $x : \uparrow_L^F A_L$  in a process offering an affine service (rule  $\uparrow L$ ) since  $L \not\geq F$ . So instead we should define an affine queue as

$$\begin{aligned} \text{queue}_F A_F &= \&\{ \text{enq} : A_F \multimap \text{queue}_F A_F, \\ &\quad \text{deq} : \{ \text{none} : \mathbf{1}, \text{some} : A_F \otimes \text{queue}_F A_F \} \} \end{aligned}$$

so that all types in the definition (including  $A_F$ ) are affine. Now we no longer need to explicitly destroy a queue, we can just abandon it and the runtime system will deallocate it by a form of garbage collection.

If we want to enforce a linear discipline, destroying a queue with linear elements will have to rely on a consumer for the elements of the queue. This consumer must be unrestricted because it is used for each element. Channels are linear by default, so in the example we only annotate affine and unrestricted channels with their mode.

$$\text{destroy} : \{ \mathbf{1} \leftarrow \text{queue}_L A_L, \uparrow_L^U(A_L \multimap \mathbf{1}) \}$$

$c \leftarrow \text{destroy} \leftarrow q, u_U =$

```

q.deq ;
case q of
| none → wait q ; close c
| some → x ← recv q ;
          d ← send u_U ;           % obtain instance d of u_U
          send d x ; wait d ;     % use d to consume x
          c ← destroy ← q, u_U   % recurse, reusing u_U

```

## 5 Polarized Logic

We now take a step in a different direction by introducing asynchronous communication, postponing discussion of the modalities for now. In asynchronous communication each linear channel contains a message queue [11], which can be related directly to the proof system via continuation channels [9]. Sending adds to the queue on one end and receiving takes from the other. Because session-based communication goes in both directions, the queue switches direction at certain times. Moreover, the queue must maintain some information on the direction of the queue so that a process that performs a send followed by a receive does not incorrectly read its own message. Fortunately, session typing guarantees that there is no send/receive mismatch.

A simple way to maintain the direction of a queue is to set a flag when enqueueing a message. We write just  $q$  when the direction of  $q$  does not matter, and  $\overleftarrow{q}$  and  $\overrightarrow{q}$  for the two directions. Our convention is that  $\overleftarrow{q}$  corresponds to messages from a provider to its client, and  $\overrightarrow{q}$  for messages from a client to the provider. The reasons for this convention is that in  $\text{proc}_c(P)$ , the channel  $c$  is to the left of  $P$ , which is in turn derived from  $c \leftarrow P$  for a process expression  $P$  offering a service along  $c$ .

We have a predicate  $\text{queue}(c, q, d)$  for a queue  $q$  connecting a process  $Q$  using  $c$  with one providing  $d$ . Here are two example rules for sending and receiving data values.

$$\begin{aligned} \text{and\_s} &: \text{queue}(c, q, d) \otimes \underline{\text{proc}}_d(\text{send } d \ v ; P) \\ &\quad \multimap \{ \text{queue}(c, \overleftarrow{q} \cdot v, d) \otimes \text{proc}_d(P) \} \\ \text{and\_r} &: \text{proc}_c(x \leftarrow \text{rcv } c ; Q_x) \otimes \text{queue}(c, \overleftarrow{v} \cdot q, d) \\ &\quad \multimap \{ \text{proc}_c(Q_v) \otimes \text{queue}(c, q, d) \} \end{aligned}$$

We see some difficulty in the second rule, where the direction of  $q$  is unclear. It should be  $\overleftarrow{q}$  unless  $q$  is empty, in which case it is unknown. This ambiguity is also present in forwarding.

$$\begin{aligned} \text{fwd} &: \text{queue}(c, p, d) \otimes \text{proc}_d(d \leftarrow e) \otimes \text{queue}(d, q, e) \\ &\quad \multimap \{ \text{queue}(c, p \cdot q, e) \} \end{aligned}$$

We won't go into detail why there are some difficulties implementing this, but we see that there are multiple possibilities for  $p$  and  $q$  pointing left, right, or being empty.

Next we note that the *polarity* of each connective determines the direction of communication. From the perspective of the service provider, if we have  $P :: (x:A)$  for a *positive*  $A$  then the action of  $P$  along  $x$  will be a send, if  $A$  is *negative* it will be receive. Intuitively this is because the right rules for negative connectives are invertible and therefore carry no information: any information has to come from the outside. Conversely, the right rules for positive connectives involve some choice and can therefore communicate the essence of that information. We can make this explicit by *polarizing* the logic, dividing the propositions into

positive and negative propositions with explicit *shift* operators connecting them. Omitting other modalities, the syntax of polarized logic is:

Positive propositions $A^+, B^+ ::=$	$\mathbf{1}$	send <b>end</b> and terminate
	$  A^+ \otimes B^+$	send channel of type $A^+$
	$  A^+ \oplus B^+$	send <b>inl</b> or <b>inr</b>
	$  \tau \wedge B^+$	send value of type $\tau$
	$  \downarrow A^-$	send <b>shift</b> , then receive
Negative propositions $A^-, B^- ::=$	$A^+ \multimap B^-$	receive channel of type $A^+$
	$  A^- \& B^-$	receive <b>inl</b> or <b>inr</b>
	$  \tau \supset B^-$	receive value of type $\tau$
	$  \uparrow A^+$	receive <b>shift</b> , then send

Note that a process that sends along a channel will continue to do so until it sends a **shift** and then it starts receiving. Conversely, a process that receives continues to do so until it receives a **shift** after which it starts sending. The new constructs are:

$P, Q, R ::=$	$\text{send } c \text{ shift} ; P$	send <b>shift</b> , then receive along $c$ in $P$
	$  \text{shift} \leftarrow \text{recv } c ; Q$	receive <b>shift</b> , then send along $c$ in $Q$

We have already annotated the shifts with their expected operational semantics. Queues now always have a definite direction and there can be no further messages following a **shift**. We write  $m$  for messages other than **shift**, such as data values, labels, and channels and treat  $\cdot$  as an associative concatenation operator with the empty queue as its unit.

Queue filled by provider	$\overleftarrow{q} ::= \cdot \mid \overleftarrow{m \cdot q} \mid \overleftarrow{\text{end}} \mid \overleftarrow{\text{shift}}$
Queue filled by client	$\overrightarrow{q} ::= \overrightarrow{\text{shift}} \mid \overrightarrow{q \cdot m} \mid \cdot$

In the polarized setting, we just need to initialize the direction correctly when a new channel is created, after which the direction is maintained correctly throughout. When receiving, the direction needs to be checked. When sending, the direction will always be correct by invariant.

$\text{and\_s} :$	$\text{queue}(c, \overleftarrow{q}, d) \otimes \text{proc}_d(\text{send } d \text{ } v ; P)$
	$\multimap \{ \text{queue}(c, \overleftarrow{q \cdot v}, d) \otimes \text{proc}_d(P) \}$
$\text{and\_r} :$	$\text{proc}_c(n \leftarrow \text{recv } c ; Q_n) \otimes \text{queue}(c, \overleftarrow{v \cdot q}, d)$
	$\multimap \{ \text{proc}_c(Q_v) \otimes \text{queue}(c, \overleftarrow{q}, d) \}$

The shift reverses direction when received.

$\text{shift\_s} :$	$\text{queue}(c, \overleftarrow{q}, d) \otimes \text{proc}_d(\text{send } d \text{ } \text{shift} ; P)$
	$\multimap \{ \text{queue}(c, \overleftarrow{q \cdot \text{shift}}, d) \otimes \text{proc}_d(P) \}$
$\text{shift\_r} :$	$\text{proc}_a(\text{shift} \leftarrow \text{recv } c ; Q) \otimes \text{queue}(c, \overleftarrow{\text{shift}}, d)$
	$\multimap \{ \text{proc}_a(Q) \otimes \text{queue}(c, \overrightarrow{\cdot}, d) \}$

There are symmetric rules for  $\overrightarrow{\text{shift}}$ , which we elide here.

In our running example, the natural polarization would interpret *queue* as a negative type, since it offers an external choice. We have to switch to positive when we send a response to the dequeue request, and then switch again before we recurse. The type parameter  $A$  is most naturally positive, since both occurrences in the type are in fact positive.

$$\text{queue}^- A^+ = \&\{ \text{enq} : A^+ \multimap \text{queue}^- A^+, \\ \text{deq} : \uparrow \oplus \{ \text{none} : \mathbf{1}, \text{some} : A^+ \otimes \downarrow \text{queue}^- A^+ \} \}$$

The code requires some minimal changes: we have to insert three shift operators.

$$\begin{array}{ll} \text{empty} : \{ \text{queue}^- A^+ \} & \text{elem} : \{ \text{queue}^- A^+ \leftarrow A^+, \text{queue}^- A^+ \} \\ c \leftarrow \text{empty} = & c \leftarrow \text{elem} \leftarrow x, d = \\ \text{case } c \text{ of} & \text{case } c \text{ of} \\ | \text{enq} \rightarrow x \leftarrow \text{recv } c ; & | \text{enq} \rightarrow y \leftarrow \text{recv } c ; \\ \quad e \leftarrow \text{empty} ; & \quad d.\text{enq} ; \text{send } d \ y ; \\ \quad c \leftarrow \text{elem} \leftarrow x, e & \quad c \leftarrow \text{elem} \leftarrow x, d \\ | \text{deq} \rightarrow \text{shift} \leftarrow \text{recv } c ; & | \text{deq} \rightarrow \text{shift} \leftarrow \text{recv } c ; \quad \% \text{ shift } c \text{ to send} \\ \quad c.\text{none} ; & \quad c.\text{some} ; \text{send } c \ x ; \\ \quad \text{close } c & \quad \text{send } c \ \text{shift} ; \quad \% \text{ shift } c \text{ to recv} \\ & \quad c \leftarrow d \end{array}$$

## 6 Recovering Synchronous Communication

We obtain maximally asynchronous communication by inserting shifts in a bare (unpolarized) session type only where necessary.

$$\begin{array}{ll} (1)^+ & = 1 \\ (A \otimes B)^+ & = (A)^+ \otimes (B)^+ \\ (A \oplus B)^+ & = (A)^+ \oplus (B)^+ \\ (\tau \wedge B)^+ & = \tau \wedge (B)^+ \\ (A)^+ & = \downarrow(A)^- \quad \text{for other propositions } A \\ (A \multimap B)^- & = (A)^+ \multimap (B)^- \\ (A \& B)^- & = (A)^- \& (B)^- \\ (\tau \supset B)^- & = \tau \supset (B)^- \\ (A)^- & = \uparrow(A)^+ \quad \text{for other propositions } A \end{array}$$

As a provider, we can send asynchronously at a positive session type until we shift explicitly to perform an input because we are now at a negative proposition. A client behaves dually.

In order to simulate synchronous communication, we insert additional shifts to prevent two consecutive send operations on the same channel. Here, the down shift after a send switches to a mode where we wait for an acknowledgment,

which is implicit in the next receive. If this is another shift, it acts as a pure acknowledgment, otherwise it is already the next message.

$$\begin{aligned}
(1)^+ &= \mathbf{1} \\
(A \otimes B)^+ &= (A)^+ \otimes \downarrow(B)^- \\
(A \oplus B)^+ &= \downarrow(A)^- \oplus \downarrow(B)^- \\
(\tau \wedge B)^+ &= \tau \wedge \downarrow(B)^- \\
(A)^+ &= \downarrow(A)^- && \text{for other propositions } A \\
(A \multimap B)^- &= (A)^+ \multimap \uparrow(B)^+ \\
(A \& B)^- &= \uparrow(A)^+ \& \uparrow(B)^+ \\
(\tau \supset B)^- &= \tau \supset \uparrow(B)^+ \\
(A)^- &= \uparrow(A)^+ && \text{for other propositions } A
\end{aligned}$$

If we want to bound the size of message queues then we can insert shift in session types which would otherwise allow an unbounded number of consecutive sends.

In our running example, a client of a queue can perform an unbounded number of enqueue operations in the asynchronous operational semantics before the queue implementation must react. This is because this portion of the queue type is entirely negative. In order to force synchronization, we can change the type of the enqueue operation before we recurse.

$$\begin{aligned}
\text{queue}^- A^+ &= \&\{ \text{enq} : A^+ \multimap \uparrow \downarrow \text{queue}^- A^+, \\
&\quad \text{deq} : \uparrow \oplus \{ \text{none} : \mathbf{1}, \text{some} : A^+ \otimes \downarrow \text{queue}^- A^+ \} \}
\end{aligned}$$

Now the maximal size of the queue will be 3 in one direction ( $\text{shift} \cdot x \cdot \text{enq}$ ) and also 3 in the other direction ( $\text{some} \cdot x \cdot \text{shift}$ ). In a slightly different language, boundedness calculations for queues in asynchronous session-typed communication can be found in [11], so we do not repeat a more formal analysis here.

## 7 Synthesis in Polarized Adjoint Logic

Now we are ready to combine the ideas from adjoint logic in Sec. 4 with polarization in Sec. 5. Amazingly, they are fully consistent. The two differences to the polarized presentation are that (a) the modalities go between positive and negative propositions (already anticipated by the fact that  $\downarrow$  is positive and  $\uparrow$  is negative), and (b) the modalities  $\downarrow_k^m A$  and  $\uparrow_k^m$  allow  $m \geq k$  rather than presupposing  $m > k$  as before. We no longer index the connectives, overloading their meaning at the different layers.

Pos. propositions $A_m^+, B_m^+ ::= \mathbf{1}$	$\downarrow A_m^+ \otimes B_m^+$ send end and terminate $\downarrow A_m^+ \oplus B_m^+$ send channel of type $A_m^+$ $\downarrow \tau \wedge B_m^+$ send inl or inr $\downarrow_m^r A_r^+$ send value of type $\tau$ $\downarrow_m^r A_r^+$ ( $r \geq m$ ), send <b>shift</b> , then receive
Neg. propositions $A_m^-, B_m^- ::= A_m^+ \multimap B_m^-$	$\downarrow A_m^- \& B_m^-$ receive channel of type $A_m^+$ $\downarrow \tau \supset B_m^-$ receive inl or inr $\uparrow_k^m A_k^+$ receive value of type $\tau$ $\uparrow_k^m A_k^+$ ( $m \geq k$ ), receive <b>shift</b> , then send

A shift staying at the same level just changes the polarity but is otherwise not subject to any restrictions. We can see this from the rules, now annotated with a polarity: if  $m = k$  in  $\uparrow\mathbf{L}$ , then  $k \geq r$  by presupposition since  $(\Psi, \uparrow_k^m A_k^+) \geq r$ . Similarly, in  $\downarrow\mathbf{R}$ ,  $\Psi \geq m$  by presupposition if  $m = k$ .

$$\frac{\Psi \vdash A_k^+}{\Psi \vdash \uparrow_k^m A_k^+} \uparrow\mathbf{R} \qquad \frac{k \geq r \quad \Psi, A_k^+ \vdash C_r}{\Psi, \uparrow_k^m A_k^+ \vdash C_r} \uparrow\mathbf{L}$$

$$\frac{\Psi_{\geq m} \vdash A_m^-}{\Psi \vdash \downarrow_k^m A_m^-} \downarrow\mathbf{R} \qquad \frac{\Psi, A_m^- \vdash C_r}{\Psi, \downarrow_k^m A_m^- \vdash C_r} \downarrow\mathbf{L}$$

Adding process expressions in a straightforward manner generalizes the shift to carry a fresh channel because there may now be a change in modes associated with the shift. We have the following new syntax

$$P, Q ::= \text{shift } x_k \leftarrow \text{send } c_m ; P_{x_k} \quad \text{send fresh shift } x_k, \text{ then recv. along } x_k \text{ in } P$$

$$| \text{shift } x_k \leftarrow \text{rcv } c_m ; Q_{x_k} \quad \text{receive shift } x_k, \text{ then send along } x_k \text{ in } Q$$

and the modified rules

$$\frac{\Psi \vdash P_{x_k} :: (x_k : A_k^+)}{\Psi \vdash (\text{shift } x_k \leftarrow \text{rcv } x_m ; P_{x_k}) :: (x_m : \uparrow_k^m A_k^+)} \uparrow\mathbf{R}$$

$$\frac{k \geq r \quad \Psi, x_k : A_k^+ \vdash Q_{x_k} :: (z_r : C_r)}{\Psi, x_m : \uparrow_k^m A_k^+ \vdash (\text{shift } x_k \leftarrow \text{send } x_m ; Q_{x_k}) :: (z_r : C_r)} \uparrow\mathbf{L}$$

$$\frac{\Psi_{\geq m} \vdash Q_{x_m} :: (x_m : A_m^-)}{\Psi \vdash (\text{shift } x_m \leftarrow \text{send } x_k ; Q_{x_m}) :: (x_k : \downarrow_k^m A_m^-)} \downarrow\mathbf{R}$$

$$\frac{\Psi, x_m : A_m^- \vdash P_{x_m} :: (z_r : C_r)}{\Psi, x_k : \downarrow_k^m A_m^- \vdash (\text{shift } x_m \leftarrow \text{rcv } x_k ; P_{x_m}) :: (z_r : C_r)} \downarrow\mathbf{L}$$

Operationally:

$$\text{up}_k^m \text{ } \_s \quad : \text{proc}_{a_r}(\text{shift } x_k \leftarrow \text{send } c_m ; Q_{x_k}) \otimes \text{queue}(c_m, \overrightarrow{q}, d_m)$$

$$\quad \rightarrow \{ \exists c_k. \exists d_k. \text{proc}_{a_r}(Q_{c_k}) \otimes \text{queue}(c_k, \text{shift } d_k \cdot \overrightarrow{q}, d_m) \}$$

$$\text{up}_k^m \text{ } \_r \quad : \text{queue}(c_k, \overrightarrow{d_k}, d_m) \otimes \text{proc}_{d_m}(\text{shift } x_k \leftarrow \text{rcv } d_m ; P_{x_k})$$

$$\quad \rightarrow \{ \text{queue}(c_k, \overleftarrow{\cdot}, d_k) \otimes \text{proc}_{d_k}(P_{d_k}) \}$$

$$\text{down}_k^m \text{ } \_s \quad : \text{queue}(c_k, \overleftarrow{q}, d_k) \otimes \text{proc}_{d_k}(\text{shift } x_m \leftarrow \text{send } d_k ; Q_{x_m})$$

$$\quad \rightarrow \{ \exists c_m. \exists d_m. \text{queue}(c_k, \overleftarrow{q} \cdot \text{shift } c_m, d_m) \otimes \text{proc}_{d_m}(Q_{d_m}) \}$$

$$\text{down}_k^m \text{ } \_r \quad : \text{proc}_{a_r}(\text{shift } x_m \leftarrow \text{rcv } c_k ; P_{x_m}) \otimes \text{queue}(c_k, \overleftarrow{\text{shift } c_m}, d_m)$$

$$\quad \rightarrow \{ \text{proc}_{a_r}(P_{c_m}) \otimes \text{queue}(c_m, \overrightarrow{\cdot}, d_m) \}$$

As pointed out in Sec. 4, we have to assume that processes that offer along an unrestricted channel  $c_U$  are persistent. Also, this formulation introduces a new

channel even when  $m = k$ , a slight redundancy best avoided in the syntax and semantics of a real implementation. Even when going between linear and affine channels, creating new channels might be avoided in favor of just changing some channel property.

Returning to forwarding, the earlier agnostic formulation will work more elegantly, since both queues to be appended are guaranteed to go into the same direction.

$$\begin{aligned} \text{fwd} : \text{queue}(c, p, d) \otimes \text{proc}_d(d \leftarrow e) \otimes \text{queue}(d, q, e) \\ \multimap \{ \text{queue}(c, p \cdot q, e) \} \end{aligned}$$

If implementation or other considerations suggest forwarding as an explicit message, we can also implement this, taking advantage of the direction information that is always available. Here we write  $x \leftarrow \text{recv } c$  as a generic receive operation along channel  $c$ , which is turned into a receive along the forwarded channel  $e$ .

$$\begin{aligned} \text{fwd}_s : \text{queue}(c, \overleftarrow{p}, d) \otimes \text{proc}_d(d \leftarrow e) \\ \multimap \{ \text{queue}(c, \overleftarrow{p \cdot \text{fwd}}, e) \} \\ \text{fwd}_r : \text{proc}_a(x \leftarrow \text{recv } c ; P_x) \otimes \text{queue}(c, \overleftarrow{\text{fwd}}, e) \\ \multimap \{ \text{proc}_a(x \leftarrow \text{recv } e ; P_x) \} \end{aligned}$$

We elide the symmetric version of the rules pointing to the right. The reason we would forward in the direction of the current communication is so that `send` remains fully asynchronous and does not have to check if a forwarding message may be present on the channel.

Once again rewriting the linear version of the example, forcing synchronization.

$$\begin{aligned} \text{queue}^- A^+ = \&\{ \text{enq} : A^+ \multimap \uparrow \downarrow \text{queue}^- A^+, \\ \text{deq} : \uparrow \oplus \{ \text{none} : \mathbf{1}, \text{some} : A^+ \otimes \downarrow \text{queue}^- A^+ \} \} \end{aligned}$$

$\begin{aligned} \text{empty} : \{ \text{queue}^- A^+ \} \\ c \leftarrow \text{empty} = \\ \text{case } c \text{ of} \\   \text{enq} \rightarrow x \leftarrow \text{recv } c ; \\ \quad \text{shift } c \leftarrow \text{recv } c \\ \quad \text{shift } c \leftarrow \text{send } c \\ \quad e \leftarrow \text{empty} ; \\ \quad c \leftarrow \text{elem} \leftarrow x, e \\   \text{deq} \rightarrow \text{shift } c \leftarrow \text{recv } c \\ \quad c.\text{none} ; \\ \quad \text{close } c \end{aligned}$	$\begin{aligned} \text{elem} : \{ \text{queue}^- A^+ \leftarrow A^+, \text{queue}^- A^+ \} \\ c \leftarrow \text{elem} \leftarrow x, d = \\ \text{case } c \text{ of} \\   \text{enq} \rightarrow y \leftarrow \text{recv } c ; \\ \quad \text{shift } c \leftarrow \text{recv } c ; \quad \% \text{ shift to send} \\ \quad \text{shift } c \leftarrow \text{send } c ; \quad \% \text{ send ack} \\ \quad d.\text{enq} ; \text{send } d y ; \\ \quad \text{shift } d \leftarrow \text{send } d ; \quad \% \text{ shift to recv} \\ \quad \text{shift } d \leftarrow \text{recv } d ; \quad \% \text{ recv ack} \\ \quad c \leftarrow \text{elem} \leftarrow x, d \\   \text{deq} \rightarrow \text{shift } c \leftarrow \text{recv } c ; \quad \% \text{ shift to send} \\ \quad c.\text{some} ; \text{send } c x ; \\ \quad \text{shift } c \leftarrow \text{send } c ; \quad \% \text{ shift to recv} \\ \quad c \leftarrow d \end{aligned}$
--	--

And destroying a linear queue with affine elements:



```

destroy : {  $\mathbf{1} \leftarrow \text{queue} (\downarrow_{\mathbb{L}}^{\mathbb{F}} A_{\mathbb{F}})$  }
c  $\leftarrow$  destroy  $\leftarrow$  q =
  q.deq ;
  shift q  $\leftarrow$  send q ;           % shift to recv
  case q of
  | none  $\rightarrow$  wait q ; close c
  | some  $\rightarrow$  x  $\leftarrow$  recv q ;   % obtain element x
    shift  $a_{\mathbb{F}}$   $\leftarrow$  recv x ;     % obtain affine  $a_{\mathbb{F}}$ , consuming x
    shift q  $\leftarrow$  send q ;       % shift to recv
    c  $\leftarrow$  destroy  $\leftarrow$  q     % recurse, ignoring  $a_{\mathbb{F}}$ 

```

## 8 Sequent Calculus for Polarized Adjoint Logic

We summarize the sequent calculus rules for polarized adjoint logic in Fig. 1, omitting the uninteresting rules for existential and universal quantification. However, we have added in atomic propositions  $p_m^+$  and  $p_m^-$  (corresponding to session type variables) and removed the stipulation that the only unrestricted propositions are  $\uparrow_m^{\mathbb{U}} A_m^+$ , thereby making our theorem slightly more general at the expense of a nonstandard notation for intuitionistic connectives such as  $A_{\mathbb{U}} \multimap_{\mathbb{U}} B_{\mathbb{U}}$  for  $A \supset B$ .

We have the following theorem.

### Theorem 1.

1. *Cut is admissible in the system without cut.*
2. *Identity is admissible for arbitrary propositions in the system with the identity restricted to atomic propositions and without cut.*

*Proof.* The admissibility of cut follows by a nested structural induction, first on the cut formula  $A$ , second simultaneously on the proofs of the left and right premise. We liberally use a lemma which states that we can weaken a proof with affine and unrestricted hypotheses without changing its structure and we exploit the transitivity of  $\geq$ . See [8,18] for analogous proofs.

The admissibility of identity at  $A$  follows by a simple structural induction on the proposition  $A$ , exploiting the reflexivity of  $\geq$  in one critical case.  $\square$

A simple corollary is *cut elimination*, stating that every provable sequent has a cut-free proof. Cut elimination of the logic is the central reason why the session-typed processes assigned to these rules satisfy the by now expected properties of *session fidelity* (processes are guaranteed to follow the behavior prescribed by the session type) and *global progress* (a closed process network of type  $c_0 : \mathbf{1}$  can either take a step will send end along  $c_0$ ). In addition, we also have *productivity* (processes will eventually perform the action prescribed by the session type) and *termination* if recursive processes are appropriately restricted. The proofs of these properties closely follow those in the literature for related systems [6,22,20], so we do not formally state or prove them here.

$$\begin{aligned}
m, k, r & ::= \mathbf{U} \mid \mathbf{F} \mid \mathbf{L} \quad \text{with } \mathbf{U} > \mathbf{F} > \mathbf{L} \\
A_m^+, B_m^+ & ::= p_m^+ \mid \mathbf{1}_m \mid A_m^+ \otimes_m B_m^+ \mid A_m^+ \oplus_m B_m^+ \mid \downarrow_m^r A_r^- \quad (r \geq m) \\
A_m^-, B_m^- & ::= p_m^- \mid A_m^+ \multimap_m B_m^- \mid A_m^- \&_m B_m^- \mid \uparrow_k^m A_k^+ \quad (m \geq k) \\
A_m, B_m, C_m & ::= A_m^+ \mid A_m^-
\end{aligned}$$

$$\begin{array}{c}
\frac{\Psi \geq \mathbf{F}}{\Psi, A_m \vdash A_m} \text{id} \qquad \frac{\Psi \geq m \geq r \quad \Psi \vdash A_m \quad \Psi', A_m \vdash C_r}{\Psi, \Psi' \vdash C_r} \text{cut} \\
\\
\frac{\Psi \vdash A_k^+}{\Psi \vdash \uparrow_k^m A_k^+} \uparrow R \qquad \frac{k \geq r \quad \Psi, A_k^+ \vdash C_r}{\Psi, \uparrow_k^m A_k^+ \vdash C_r} \uparrow L \\
\\
\frac{\Psi_{\geq m} \vdash A_m^-}{\Psi \vdash \downarrow_k^m A_m^-} \downarrow R \qquad \frac{\Psi, A_m^- \vdash C_r}{\Psi, \downarrow_k^m A_m^- \vdash C_r} \downarrow L \\
\\
\frac{\Psi \geq \mathbf{F}}{\Psi \vdash \mathbf{1}_m} \mathbf{1}R \qquad \frac{\Psi \vdash C_r}{\Psi, \mathbf{1}_m \vdash C_r} \mathbf{1}L \\
\\
\frac{\Psi \vdash A_m^+ \quad \Psi' \vdash B_m^+}{\Psi, \Psi' \vdash A_m^+ \otimes_m B_m^+} \otimes R \qquad \frac{\Psi, A_m^+, B_m^+ \vdash C_r}{\Psi, A_m^+ \otimes_m B_m^+ \vdash C_r} \otimes L \\
\\
\frac{\Psi, A_m^+ \vdash B_m^-}{\Psi \vdash A_m^+ \multimap_m B_m^-} \multimap R \qquad \frac{\Psi \geq m \quad \Psi \vdash A_m^+ \quad \Psi', B_m^- \vdash C_r}{\Psi, \Psi', A_m^+ \multimap_m B_m^- \vdash C_r} \multimap L \\
\\
\frac{\Psi \vdash A_m^- \quad \Psi \vdash B_m^-}{\Psi \vdash A_m^- \&_m B_m^-} \& R \qquad \frac{\Psi, A_m^- \vdash C_r}{\Psi, A_m^- \&_m B_m^- \vdash C_r} \& L_1 \qquad \frac{\Psi, B_m^- \vdash C_r}{\Psi, A_m^- \&_m B_m^- \vdash C_r} \& L_2 \\
\\
\frac{\Psi \vdash A_m^+}{\Psi \vdash A_m^+ \oplus_m B_m^+} \oplus R_1 \qquad \frac{\Psi \vdash B_m^+}{\Psi \vdash A_m^+ \oplus_m B_m^+} \oplus R_2 \qquad \frac{\Psi, A_m^+ \vdash C_r \quad \Psi, B_m^+ \vdash C_r}{\Psi, A_m^+ \oplus_m B_m^+ \vdash C_r} \oplus L
\end{array}$$

All judgments  $\Psi \vdash A_m$  presuppose  $\Psi \geq m$ .

$\Psi, \Psi'$  allows contraction of unrestricted  $A_U$  shared between  $\Psi$  and  $\Psi'$

**Fig. 1.** Polarized Adjoint Logic

## 9 Conclusion

We have developed a language which uniformly integrates linear, affine, and unrestricted types, allowing the programmer to vary the degree of precision with which resources are managed. At the same time, the programmer has fine-grained control over which communications are synchronous or asynchronous, and these decisions are reflected in the type in a logically motivated manner.

On the pragmatic side, we should decide to what extent the constructs here are exposed to the programmer or inferred during type checking, and develop a concise and intuitive concrete syntax for those that are explicitly available in types and process expressions.

Finally, our language is polarized, but deductions are not focused [1]. This is perhaps somewhat unexpected since the two are closely connected and historically tied to each other. It suggests that some further benefits from proof-theoretic concepts are still to be discovered, continuing the current line of investigation into the foundation of session-typed concurrency.

**Acknowledgments.** This material is based upon work supported by the National Science Foundation under Grant No. CNS-1423168 and by the Department of Energy under Grant No. SNL 1488651. The authors would like to acknowledge helpful discussions with Stephanie Balzer, Luís Caires, Henry DeYoung, Hannah Gommerstadt, Elsa Gunter, Limin Jia, Rokhini Prabhu, Bernardo Toninho, and Max Willsey.

## References

1. Andreoli, J.M.: Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation* 2(3), 197–347 (1992)
2. Barber, A.: Dual intuitionistic linear logic. Tech. Rep. ECS-LFCS-96-347, Department of Computer Science, University of Edinburgh (September 1996)
3. Benton, N.: A mixed linear and non-linear logic: Proofs, terms and models. In: Pacholski, L., Tiuryn, J. (eds.) *CSL 1994. LNCS*, vol. 933, pp. 121–135. Springer, Heidelberg (1995), an extended version appears as Technical Report UCAM-CL-TR-352, University of Cambridge
4. Caires, L., Pérez, J.A., Pfenning, F., Toninho, B.: Behavioral polymorphism and parametricity in session-based communication. In: Felleisen, M., Gardner, P. (eds.) *ESOP 2013. LNCS*, vol. 7792, pp. 330–349. Springer, Heidelberg (2013)
5. Caires, L., Pfenning, F.: Session types as intuitionistic linear propositions. In: Gastin, P., Laroussinie, F. (eds.) *CONCUR 2010. LNCS*, vol. 6269, pp. 222–236. Springer, Heidelberg (2010)
6. Caires, L., Pfenning, F., Toninho, B.: Linear logic propositions as session types. *Mathematical Structures in Computer Science* (2013) (to appear ) Special Issue on Behavioural Types
7. Cervesato, I., Scedrov, A.: Relating state-based and process-based concurrency through linear logic. *Information and Computation* 207(10), 1044–1077 (2009)
8. Chang, B.Y.E., Chaudhuri, K., Pfenning, F.: A judgmental analysis of linear logic. Tech. Rep. CMU-CS-03-131R, Carnegie Mellon University, Department of Computer Science (December 2003)
9. DeYoung, H., Caires, L., Pfenning, F., Toninho, B.: Cut reduction in linear logic as asynchronous session-typed communication. In: Cégielski, P., Durand, A. (eds.) *Proceedings of the 21st Conference on Computer Science Logic, CSL 2012, Leibniz International Proceedings in Informatics, Fontainebleau, France*, pp. 228–242 (September 2012)
10. Gay, S.J., Hole, M.: Subtyping for session types in the  $\pi$ -calculus. *Acta Informatica* 42(2-3), 191–225 (2005)
11. Gay, S.J., Vasconcelos, V.T.: Linear type theory for asynchronous session types. *Journal of Functional Programming* 20(1), 19–50 (2010)
12. Girard, J.Y.: Linear logic. *Theoretical Computer Science* 50, 1–102 (1987)

13. Honda, K.: Types for dyadic interaction. In: Best, E. (ed.) CONCUR 1993. LNCS, vol. 715, pp. 509–523. Springer, Heidelberg (1993)
14. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998)
15. Laurent, O.: Polarized proof-nets: Proof-nets for LC. In: Girard, J.-Y. (ed.) TLCA 1999. LNCS, vol. 1581, pp. 213–227. Springer, Heidelberg (1999)
16. Mostrous, D., Vasconcelos, V.: Affine sessions. In: Kühn, E., Pugliese, R. (eds.) COORDINATION 2014. LNCS, vol. 8459, pp. 115–130. Springer, Heidelberg (2014)
17. Pérez, J.A., Caires, L., Pfenning, F., Toninho, B.: Linear logical relations and observational equivalences for session-based concurrency. *Information and Computation* 239, 254–302 (2014)
18. Reed, J.: A judgmental deconstruction of modal logic (2009) (unpublished manuscript)
19. Simmons, R.J.: Substructural Logical Specifications. Ph.D. thesis, Carnegie Mellon University, available as Technical Report CMU-CS-12-142 (November 2012)
20. Toninho, B.: A Logical Foundation for Session-based Concurrent Computation. Ph.D. thesis, Carnegie Mellon University and New University of Lisbon (2015) (in preparation)
21. Toninho, B., Caires, L., Pfenning, F.: Higher-order processes, functions, and sessions: A monadic integration. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 350–369. Springer, Heidelberg (2013)
22. Toninho, B., Caires, L., Pfenning, F.: Corecursion and non-divergence in session-typed processes. In: Proceedings of the 9th International Symposium on Trustworthy Global Computing (TGC 2014), Rome, Italy (September 2014) (to appear)
23. Vasconcelos, V.T.: Fundamentals of session types. *Information and Computation* 217, 52–70 (2012)
24. Wadler, P.: Propositions as sessions. In: Proceedings of the 17th International Conference on Functional Programming, ICFP 2012, pp. 273–286. ACM Press, Copenhagen (2012)