

# Asynchronous Multistructural Session Types

FRANK PFENNING, Carnegie Mellon University

KLAAS PRUIKSMA, Carnegie Mellon University

We present a system of multistructural session types that uniformly integrates layers with linear and non-linear types based on Benton's LNL. The underlying computational mechanism derives from traditional principal cut reductions and therefore is, by nature, synchronous. We then reformulate the sequent calculus in order to enforce *asynchronous* communication in the sense that a sender can proceed before a message is received, but the messages must still be received in order.

In addition to the usual point-to-point message passing primitives supported by binary session types, we also obtain logical explanations of multicast (sending a message to multiple recipients) and copy-on-receive (instantiating a new copy of a service upon receipt of a message). Our language also allows a write-once shared memory interpretation suitable for implementation. Surprisingly, this shared memory interpretation does *not* map sends to writes and receives to reads, but cuts to allocation, right rules to writes, and left rules to reads.

Additional Key Words and Phrases: Linear logic, session types, concurrency

## 1 INTRODUCTION

Binary session types [Honda 1993; Honda et al. 1998] arise in a natural way from a Curry-Howard interpretation of linear sequent calculus [Caires and Pfenning 2010; Wadler 2012], where a principal cut reduction is interpreted as an act of communication. This leads to a *synchronous* model of communication since a cut reduction simultaneously affects both premises of the cut. While theoretically elegant, synchronous message passing is not immediately implementable, so *asynchronous* models of message exchange have also been investigated [Carbone et al. 2008; Gay and Vasconcelos 2010]. Unlike in the  $\pi$ -calculus, in the setting of linear session types, synchronous and asynchronous models of computation can simulate each other [Balzer and Pfenning 2017; DeYoung et al. 2012; Pfenning and Griffith 2015].

A first contribution of this paper is the design of a variant of the sequent calculus where the cut reduction steps correspond directly to *asynchronous* communication. This calculus does not satisfy traditional cut elimination, but proofs have an analytic normal form that can be obtained via translations to and from the traditional sequent calculus. In addition, computation satisfies the expected session fidelity and global progress properties even when extended with recursive types and processes.

Our second contribution is a generalization LNL<sup>†</sup> of the modified sequent calculus to encompass LNL [Benton 1994], which uniformly integrates linear and intuitionistic layers. While it is difficult to see a feasible synchronous concurrent semantics for the (traditional) sequent calculus, its asynchronous formulation generalizes rather easily, modeling multicast (simultaneously sending one message to multiple recipients) and copy-on-receive (instantiating a new copy of a service upon message receipt).

As our third contribution we observe that the asynchronous sequent calculus also admits a *shared memory semantics*. One might expect that send actions become memory writes and receive actions become memory reads, but instead all right rules write and all left rules read. On the negative connectives this means that instead of receiving a message, a process will write a continuation

---

Authors' addresses: Frank Pfenning, Computer Science Department, Carnegie Mellon University, fp@cs.cmu.edu; Klaas Pruiksma, Computer Science Department, Carnegie Mellon University, kpruiksm@andrew.cmu.edu.

2018. 2475-1421/2018/1-ART1 \$15.00

<https://doi.org/>

to memory and instead of sending a message, a process will jump to a stored continuation. This comes close to being a typed version of futures [Halstead 1985] except that we still have too much parallelism: we synchronize in the style of futures, but other language constructs are also concurrent. The notion of shared memory, however, remains pure in the sense that we allocate uninitialized memory but it becomes immutable once it has been written to.

Both operational interpretations leverage the expressive power of *substructural operational semantics* [Pfenning 2004; Pfenning and Simmons 2009; Simmons 2012] in an interesting way: The structural properties of semantic objects under their multiset rewriting interpretation [Cervesato and Scedrov 2009] mirror the structural properties of the types they interpret.

The remainder of the paper is organized as follows. In Sec. 2 we present a minor variant of Benton’s sequent calculus for LNL, followed by an analysis of cut reduction as communication for the purely linear fragment in Sec. 3. We then formalize our operational intuition in the style of substructural operational semantics in Sec. 4, still for the linear fragment. This is generalized to encompass non-linear types in Sec. 5. We analyze the metatheory, proving session fidelity and freedom from deadlock in Sec. 6. A shared memory semantics is then presented in Sec. 7. We conclude with some additional related and future work.

## 2 A SEQUENT CALCULUS FOR LNL

In this section we provide a slightly streamlined formulation of LNL [Benton 1994] with an eye towards future generalization to Adjoint Logic [Pruikma et al. 2018a; Reed 2009]. We posit two modes of truth, *linear* and *non-linear*, where linear propositions satisfy only exchange while non-linear propositions also satisfy weakening and contraction. We write L for the linear mode and U for the non-linear mode and index every proposition and connective with a mode.

$$\begin{array}{lcl}
 A_m, B_m, C_m & ::= & \oplus_m \{ \ell : A_m^\ell \}_{\ell \in L} \mid A_m \otimes_m B_m \mid \mathbf{1}_m & \text{positives} \\
 & & \mid \&_m \{ \ell : A_m^\ell \}_{\ell \in L} \mid A_m \multimap_m B_m & \text{negatives} \\
 & & \mid (\uparrow A_L)_U \mid (\downarrow A_U)_L & \text{shifts}
 \end{array}$$

All connectives except for the shift modalities have linear as well as non-linear versions. For example, the usual intuitionistic implication  $A \rightarrow B$  would be written as  $A_U \multimap_U B_U$  for non-linear propositions  $A$  and  $B$ . The shift modalities have the restriction that  $\uparrow A_L$  is of mode U and  $\downarrow A_U$  is of mode L. In Benton’s notation,  $\downarrow A_U = F A_U$  and  $\uparrow A_L = G A_L$ . It is possible to polarize the proposition via the shift operators [Pfenning and Griffith 2015] but this does not appear important for the results in the paper. We generally omit the mode subscript on the connectives since it follows from the mode of the constituent propositions. We also permit propositional variables  $a_m$  and proofs can be parametric in such variables.

We write  $\Psi$  for antecedents with implicit exchange that consist of mixed linear and non-linear propositions. We write  $\Psi \geq m$  if all propositions in  $\Psi$  have a mode greater or equal to  $m$ , where  $U > L$ .  $\Psi_U$  stands for a collection of antecedents of mode U. All our sequents  $\Psi \vdash A_m$  must syntactically obey the *independence principle* stating that  $\Psi \geq m$ . We enforce this in the rules when read bottom-up with appropriate premises. Just as in Benton’s formulation, we think of this is being a syntactic presupposition necessary for the well-formedness of a sequent and never consider sequents that violate independence. LNL is a generalization of intuitionistic linear logic in that we can recover the exponential  $!A$  as  $\downarrow \uparrow A_L$ .

The inference rules for LNL can be found in Figure 1. The independence principle entails, for example, that there are three versions of the rule of cut ( $m = U$  and  $k$  either U or L or  $m = k = L$ ). This calculus satisfies the usual expected property of cut and identity elimination.

**THEOREM 2.1 (CUT AND IDENTITY [BENTON 1994]).** *If  $\Psi \vdash A_k$  then there is a proof of the same sequent without the rule of cut and using identity only on propositional variables.*

$$\begin{array}{c}
\frac{}{A_m \vdash A_m} \text{id} \qquad \frac{\Psi_1 \geq m \geq k \quad \Psi_1 \vdash A_m \quad \Psi_2, A_m \vdash C_k}{\Psi_1, \Psi_2 \vdash C_k} \text{cut} \\
\frac{\Psi \vdash C}{\Psi, A_u \vdash C_k} \text{weaken} \qquad \frac{\Psi, A_u, A_u \vdash C_k}{\Psi, A_u \vdash C_k} \text{contract} \\
\frac{i \in L \quad \Psi \vdash A_m^i}{\Psi \vdash \oplus\{\ell : A_m^\ell\}_{\ell \in L}} \oplus R_i \qquad \frac{(\text{for all } \ell \in L) \quad \Psi, A_m^\ell \vdash C_k}{\Psi, \oplus\{\ell : A_m^\ell\}_{\ell \in L} \vdash C_k} \oplus L \\
\frac{}{\cdot \vdash \mathbf{1}} \mathbf{1}R \qquad \frac{\Psi \vdash C_k}{\Psi, \mathbf{1} \vdash C_k} \mathbf{1}L \\
\frac{\Psi_1 \vdash A_m \quad \Psi_2 \vdash B_m}{\Psi_1, \Psi_2 \vdash A_m \otimes B_m} \otimes R \qquad \frac{\Psi, A_m, B_m \vdash C_k}{\Psi, A_m \otimes B_m \vdash C_k} \otimes L \\
\frac{(\text{for all } \ell \in L) \quad \Psi \vdash A_m^\ell}{\Psi \vdash \&\{\ell : A_m^\ell\}_{\ell \in L}} \&R \qquad \frac{i \in L \quad \Psi, A_m^i \vdash C_k}{\Psi, \&\{\ell : A_m^\ell\}_{\ell \in L} \vdash C_k} \&L_i \\
\frac{\Psi, A_m \vdash B_m}{\Psi \vdash A_m \multimap B_m} \multimap R \qquad \frac{\Psi_1 \geq m \quad \Psi_1 \vdash A_m \quad \Psi_2, B_m \vdash C_k}{\Psi_1, \Psi_2, A_m \multimap B_m \vdash C_k} \multimap L \\
\frac{\Psi_u \vdash A_u}{\Psi_u \vdash \downarrow A_u} \downarrow R \qquad \frac{\Psi, A_u \vdash C_l}{\Psi, \downarrow A_u \vdash C_l} \downarrow L \\
\frac{\Psi \vdash A_l}{\Psi \vdash \uparrow A_l} \uparrow R \qquad \frac{\Psi, \uparrow A_l, A_l \vdash C_l}{\Psi, \uparrow A_l \vdash C_l} \uparrow L
\end{array}$$

Fig. 1. LNL with explicit structural rules

### 3 CUT REDUCTION AS COMMUNICATION

In this section we consider only the operational interpretation of the linear fragment of LNL, to be generalized in Sec. 5. Under the Curry-Howard isomorphism, linear propositions correspond to types and proofs to processes. Equally significant is that in the sequent calculus, the rules of the operational semantics are derived from cut reduction [Caires and Pfenning 2010].

Fundamentally, we interpret the proof of a sequent  $A_1, \dots, A_n \vdash C$  as a process  $P$  and write

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

where  $x_i$  and  $z$  are distinct *channels*. We say the  $P$  is the *client* to  $x_i$  and *provides*  $z$ . The propositions  $A_i$  and  $C$  are interpreted as *session types* prescribing the form of interactions of  $P$  along these channels. Before we get to the interpretation of specific types, we review the linear form of cut, annotated with process expressions. Here we write  $\Delta$  instead of  $\Psi$  to emphasize the linearity of all propositions. Where bound variables are present we will generally indicate a dependence of  $P$  on one or more variables by writing  $P[x_1, \dots, x_n]$ .

$$\frac{\Delta \vdash P :: (x : A) \quad \Delta', x : A \vdash Q :: (z : C)}{\Delta, \Delta' \vdash (x \leftarrow P[x] ; Q[x]) :: (z : C)} \text{cut}$$

We see that cut corresponds to the parallel composition of  $P$  and  $Q$  with a private channel  $x$  provided by  $P$  and used by  $Q$ .

As an example of a logical connective we now consider the binary version of internal choice,  $A \oplus B = \oplus\{\pi_1 : A, \pi_2 : B\}$ . The following cut reduction (plus the symmetric one for  $\oplus R_2$ ) should be the guide:

$$\frac{\frac{\frac{P}{\Delta \vdash A}}{\Delta \vdash A \oplus B} \oplus R_1 \quad \frac{\frac{\frac{Q_1}{\Delta', A \vdash C} \quad \frac{Q_2}{\Delta', B \vdash C}}{\Delta', A \oplus B \vdash C} \oplus L}{\Delta, \Delta' \vdash C} \text{cut}_{A \oplus B}}{\Delta, \Delta' \vdash C} \text{cut}_A \quad \Rightarrow \quad \frac{\frac{P}{\Delta \vdash A} \quad \frac{Q_1}{\Delta', A \vdash C}}{\Delta, \Delta' \vdash C} \text{cut}_A$$

We see that the proof of the first premise has some information (namely: which disjunct is actually proved) while the proof of the second premise branches on this information. In other words, the process implementing the first premise has to communicate either  $\pi_1$  or  $\pi_2$  to the process implementing the second premise. We can also see that communication according to this rule is *synchronous*: the processes in both premises step forward in unison during the reduction.

Extrapolating from this example, we can see that, from the provider's perspective, positive connectives ( $\oplus$ ,  $\mathbf{1}$ ,  $\otimes$ ) send information while negative connectives ( $\&$ ,  $\multimap$ ) receive. Of course, the client performs complementary actions.

### 3.1 Enforcing Asynchronous Communication

As we have seen in the preceding section, cut reduction in a standard sequent calculus implies a synchronous model of communication under the Curry-Howard correspondence. While synchronous communication is not directly implementable, it has previously been observed that we can *encode* asynchronous communication [Balzer and Pfenning 2017; DeYoung et al. 2012]. Nevertheless, from a foundational perspective this is somewhat unsatisfactory because either (a) the implementer of a session-typed language must somehow code synchronous communication at a lower level of abstraction, for example, with acknowledgments [Pfenning and Griffith 2015], or (b) the implementer uses an asynchronous semantics, departing from the Curry-Howard foundation of the calculus.

The *asynchronous  $\pi$ -calculus* replaces the usual action prefix for output  $x(y).P$  by a process expression  $x(y)$  *without a continuation*, thereby ensuring that communication is asynchronous. Such a process represents the message  $y$  sent along channel  $x$ . Under our interpretation, the continuation process corresponds to the proof of the premise of a rule. Therefore, if we can restructure the sequent calculus so that the rules that send ( $\oplus R$ ,  $\mathbf{1}R$ ,  $\otimes R$ ,  $\downarrow R$ ,  $\&L$ ,  $\multimap L$ ,  $\uparrow L$ ) have zero premises, then we may achieve a corresponding effect.

As an example, we consider the two right rules for  $\oplus$ . Reformulated as axioms, they become

$$\frac{}{A \vdash A \oplus B} \oplus R_1^0 \quad \frac{}{B \vdash A \oplus B} \oplus R_2^0$$

In the presence of cut, these two rules together produce the same theorems as the usual two right rules. In one direction, we use cut

$$\frac{\Delta \vdash A \quad \frac{}{A \vdash A \oplus B} \oplus R_1^0}{\Delta \vdash A \oplus B} \text{cut}_A \quad \frac{\Delta \vdash B \quad \frac{}{B \vdash A \oplus B} \oplus R_2^0}{\Delta \vdash A \oplus B} \text{cut}_B$$

and in the other direction we use identity

$$\frac{\frac{}{A \vdash A} \text{id}_A}{A \vdash A \oplus B} \oplus R_1 \quad \frac{\frac{}{B \vdash B} \text{id}_B}{B \vdash A \oplus B} \oplus R_2$$

to derive the other rules.

Returning to the  $\pi$ -calculus, instead of explicitly *sending* a message  $a\langle b \rangle.P$  we *spawn* a new process in parallel  $a\langle b \rangle \mid P$ . This use of parallel composition corresponds to a cut; receiving a message is achieved by cut reduction:

$$\frac{\frac{\frac{}{A \vdash A \oplus B} \oplus R_1^0 \quad \frac{\frac{Q_1 \quad \Delta', A \vdash C \quad \Delta', B \vdash C}{\Delta', A \oplus B \vdash C} \oplus L}{\Delta', A \vdash C} \text{cut}_{A \oplus B}}{\Delta', A \vdash C} \implies \Delta', A \vdash C$$

We see the cut reduction completely eliminates the cut in one step, which corresponds precisely to receiving a message. In this example the message would be  $\pi_1$  since the axiom  $\oplus R_1^0$  was used; for  $\oplus R_1^0$  it would be  $\pi_2$ .

In summary, if we restructure the sequent calculus so that the non-invertible rules (those that send) have zero premises, then (1) messages are proofs of axioms, (2) message sends are modeled by cut, and (3) message receives are a new form of cut reduction with a single continuation.

In the process we give something up, namely the traditional cut elimination theorem. For example, the sequent  $\cdot \vdash 1 \oplus 1$  has no cut-free proof since no rule matches this conclusion. The saving grace is that we can reach a normal form where each cut just simulates the usual rules of the sequent calculus. This can be shown by translation to the ordinary sequent calculus, applying cut elimination, and translating the result back. Proofs in this normal form have a strong subformula property. Perhaps more importantly, we have session fidelity and deadlock freedom for the corresponding process calculus even in the presence of recursive types and processes, which is ultimately what we care about for the resulting concurrent programming language.

In addition to replacing some of the rules with their axiomatic form, we also make the structural rules implicit in the standard manner: antecedents  $A_U$  are propagated to all premises of rule (implicit contraction) and are allowed in zero premise rules (implicit weakening). The rules for the resulting system  $\text{LNL}^\dagger$  are summarized in Fig. 2. We obtain the following theorem.

**THEOREM 3.1 (ADEQUACY OF  $\text{LNL}^\dagger$ ).**

- (i) *Weakening and contraction are admissible in  $\text{LNL}^\dagger$ .*
- (ii) *If  $\Psi \vdash A$  in  $\text{LNL}$  then  $\Psi \vdash A$  in  $\text{LNL}^\dagger$ .*
- (iii) *If  $\Psi \vdash A$  in  $\text{LNL}^\dagger$  then  $\Psi \vdash A$  in  $\text{LNL}$ .*

**PROOF.** By straightforward inductions over the given deductions. For part (ii) we use cuts with the zero-premise rules. For part (iii) we use the  $\text{LNL}$  proofs of the zero-premise sequents.  $\square$

A simple variant of this system omits the premises marked  $B|_U$  in the  $\oplus L$ ,  $1L$ , and  $\otimes L$  rules. These premises are logically redundant and, moreover, contraction can be derived straightforwardly as an instance of cut:

$$\frac{A_U \vdash A_U \quad \Psi, A_U, A_U \vdash C_k}{\Psi, A_U \vdash C_k} \text{cut}$$

We prefer to keep them for two reasons: logically we adhere to the principle of implicit structural rules and computationally these additional premises allow the programmer to reuse channels if they wish to do so.

#### 4 AN ASYNCHRONOUS OPERATIONAL SEMANTICS

We now assign process expressions to  $\text{LNL}^\dagger$  and give an asynchronous operational semantics. We consider some examples of our process notation, starting with  $\oplus\{\ell : A^\ell\}_{\ell \in L}$ .

$$\frac{i \in L}{y : A^i \vdash x.i(y) :: (x : \oplus\{\ell : A^\ell\}_{\ell \in L})} \oplus R_i^0$$

$$\begin{array}{c}
\frac{}{\Psi_U, A_m \vdash A_m} \text{id} \qquad \frac{\Psi_1 \geq m \geq k \quad \Psi_U, \Psi_1 \vdash A_m \quad \Psi_U, \Psi_2, A_m \vdash C_k}{\Psi_U, \Psi_1, \Psi_2 \vdash C_k} \text{cut} \\
\\
\frac{i \in L}{\Psi_U, A_i \vdash \{\ell : A_m^\ell\}_{\ell \in L}} \oplus R_i^0 \qquad \frac{(\text{for all } \ell \in L) \quad \Psi, (\oplus\{\ell : A_m^\ell\}_{\ell \in L})|_U, A_m^\ell \vdash C_k}{\Psi, \oplus\{\ell : A_m^\ell\}_{\ell \in L} \vdash C_k} \oplus L \\
\\
\frac{}{\Psi_U \vdash \mathbf{1}} \mathbf{1}R^0 \qquad \frac{\Psi, \mathbf{1}_m|_U \vdash C_k}{\Psi, \mathbf{1}_m \vdash C_k} \mathbf{1}L \\
\\
\frac{}{\Psi_U, A_m, B_m \vdash A_m \otimes B_m} \otimes R^0 \qquad \frac{\Psi, (A_m \otimes B_m)|_U, A_m, B_m \vdash C_k}{\Psi, A_m \otimes B_m \vdash C_k} \otimes L \\
\\
\frac{(\text{for all } \ell \in L) \quad \Psi \vdash A_m^\ell}{\Psi \vdash \&\{\ell : A_m^\ell\}_{\ell \in L}} \&R \qquad \frac{i \in L}{\Psi_U, \&\{\ell : A_m^\ell\}_{\ell \in L} \vdash A_m^i} \&L_i^0 \\
\\
\frac{\Psi, A_m \vdash B_m}{\Psi \vdash A_m \multimap B_m} \multimap R \qquad \frac{}{\Psi_U, A_m, A_m \multimap B_m \vdash B_m} \multimap L^0 \\
\\
\frac{}{\Psi_U, A_U \vdash \downarrow A_U} \downarrow R^0 \qquad \frac{\Psi, A_U \vdash C_L}{\Psi, \downarrow A_U \vdash C_L} \downarrow L \\
\\
\frac{\Psi \vdash A_L}{\Psi \vdash \uparrow A_L} \uparrow R \qquad \frac{}{\Psi_U, \uparrow A_L \vdash A_L} \uparrow L^0
\end{array}$$

Fig. 2. LNL<sup>†</sup> with implicit structural rules  
 $B|_U$  in rules  $\oplus L$ ,  $\mathbf{1}L$ , and  $\otimes L$  means  $B$  is present iff  $m = U$

We read  $x.i(y)$  as “send label  $i$  along  $x$  with continuation channel  $y$ ”. Conversely, the recipient remains unchanged from the usual session type interpretation.

$$\frac{\text{for all } \ell \in L \quad \Delta, y : A^\ell \vdash Q_\ell[y] :: (z : C)}{\Delta, x : \oplus\{\ell : A^\ell\}_{\ell \in L} \vdash (\text{case } x(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}) :: (z : C)} \oplus L$$

Now, a provider  $c.i(d)$  (representing a single, asynchronous message) should interact with a client (case  $c(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}$ ) by transitioning to  $Q_i[d]$ , selecting the branch  $i$  and substituting the continuation channel  $d$ . We present this semantics in the style of *substructural operational semantics* [Cervesato et al. 2002; Pfenning 2004; Simmons 2012] which uses *multiset rewriting* [Cervesato and Scedrov 2009]. A *process configuration* is a collection of semantic objects of the form  $\text{proc}(c, P)$ , which indicates that process  $P$  is executing and providing along channel  $c$ . We do not explicitly record the channels that  $P$  uses, because under the benign restriction that internal and external choice are never empty, these consist just of the channels free in  $P$ . For simplicity, we will make this restriction in the remainder of this paper. In a configuration, all channels  $c$  must be distinct, and the order of the objects is irrelevant. A multiset rewriting rule for semantic objects  $\phi_j$ ,  $\psi_k$  and channels  $a_l$  has the form

$$\phi_1, \dots, \phi_m \longrightarrow \psi_1, \dots, \psi_p \quad (a_1, \dots, a_n \text{ fresh})$$

It matches the objects  $\phi_j$  against objects in the configuration, ignoring order but not allowing duplication, and replaces the matching objects by  $\psi_k$  after creating fresh channels  $a_l$ . Note that all

channels free in  $\psi_k$  must either occur free in  $\phi_j$  or be among  $a_l$ , where the  $a_l$  chosen must be fresh and may not occur elsewhere in the configuration.

The first rule then has the form

$$\text{proc}(c, c.i(d)), \text{proc}(e, \text{case } c(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}) \longrightarrow \text{proc}(e, Q_i[d])$$

The complete process assignment can be found in Fig. 3 and the operational rules for *the linear fragment* in Fig. 4. A few further remarks on these rules. We notice that only cuts create fresh channels. In part, this is because there is no sending constructs: all “sending” of messages is accomplished by spawning an appropriate process using cut. Another observation is the difference to the asynchronous  $\pi$ -calculus where messages may arrive out of order. In the typed setting this would be disastrous if consecutive messages had different type, so we employ use-once channels (as already proposed by Kobayashi et al. [1996]) together with continuation channels. The continuation channels in effect form a queue that guarantees in-order arrival and thereby, ultimately, session fidelity. Here, however, this structure arises entirely from the logical side and reflects the cut reductions in LNL<sup>†</sup>. We have generally used  $c$  for the principal channel of communication and  $d$  for its intended continuation.

Finally, consider the identity rule.

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

We read the process expression  $(x \leftarrow y)$  as “ $x$  is implemented by  $y$ ” and call it *forwarding*. Intuitively, the process  $(x \leftarrow y)$  is supposed to provide  $x$  and fulfills that responsibility by forwarding to  $y$ . This is sound because both channels have the same type  $A$ .

In cut elimination, there are two reductions when cut meets identity.

$$\frac{\frac{P}{\Delta \vdash A} \quad \frac{}{A \vdash A} \text{id}}{\Delta \vdash A} \text{cut} \implies \frac{P}{\Delta \vdash A} \quad \frac{\frac{}{A \vdash A} \text{id} \quad \frac{Q}{\Delta', A \vdash C}}{\Delta', A \vdash C} \text{cut} \implies \frac{Q}{\Delta', A \vdash C}$$

These give rise to the two rules

$$\begin{aligned} (\text{id}C) \quad & \text{proc}(d, P[d]), \text{proc}(c, c \leftarrow d) \longrightarrow \text{proc}(c, P[c]) \\ (\text{id}C') \quad & \text{proc}(c, c \leftarrow d), \text{proc}(e, Q[c]) \longrightarrow \text{proc}(e, Q[d]) \end{aligned}$$

In a closed system where a whole configuration promises interaction with the outside world along a single channel  $c_0$ , the rule  $\text{id}C$  is preferable because with  $\text{id}C'$  a forwarding process  $\text{proc}(c_0, c_0 \leftarrow d)$  at the interface might never be eliminated. However, lower-level considerations might influence the choice of whether to use just one or maybe even both of these rules. Fortunately, they are confluent in the sense that applying one or the other leads to configurations that are related by renaming an internal channel and would therefore be considered identical. Other reasons for preferring  $\text{id}C$  will emerge when we consider shared channels and the shared memory semantics.

#### 4.1 Recursive Types and Processes

For the examples we add recursive types and recursively defined processes. Without further restrictions, they break the Curry-Howard isomorphism; appropriate investigation of least and greatest fixed point interpretations are the subject of ongoing work.

We encode recursive types via (potentially mutually recursive) type definitions stored in a global signature. Since our types are equirecursive we stipulate that they must be *contractive* [Gay and Hole 2005]. This allows us to silently replace type names by their definitions. In addition, we allow

$$\begin{array}{c}
\frac{\Psi_1 \geq m \geq k \quad \Psi_U, \Psi_1 \vdash P[x] :: (x : A_m) \quad \Psi_U, \Psi_2, x : A_m \vdash Q[x] :: (z : C_k)}{\Psi_U, \Psi_1, \Psi_2 \vdash (x \leftarrow P[x] ; Q[x]) :: (z : C_k)} \text{ cut} \\
\\
\frac{}{\Psi_U, y : A_m \vdash (x \leftarrow y) :: (x : A_m)} \text{ id} \\
\\
\frac{i \in L}{\Psi_U, y : A_i \vdash x.i(y) :: (x : \oplus\{\ell : A_m^\ell\}_{\ell \in L})} \oplus R_i^0 \\
\\
\frac{(\text{for all } \ell \in L) \quad \Psi, x : (\oplus\{\ell : A_m^\ell\}_{\ell \in L})|_U, y : A_m^\ell \vdash Q[y] :: (z : C_k)}{\Psi, x : \oplus\{\ell : A_m^\ell\}_{\ell \in L} \vdash (\text{case } x(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}) :: (z : C_k)} \oplus L \\
\\
\frac{}{\Psi_U \vdash (x.\langle \rangle) :: (x : \mathbf{1})} \mathbf{1}R^0 \qquad \frac{\Psi, x : \mathbf{1}_m|_U \vdash Q :: (z : C_k)}{\Psi, x : \mathbf{1}_m \vdash (\langle \rangle \leftarrow x ; Q) :: (z : C_k)} \mathbf{1}L \\
\\
\frac{}{\Psi_U, w : A_m, y : B_m \vdash (x.\langle w, y \rangle) :: (x : A_m \otimes B_m)} \otimes R^0 \\
\\
\frac{\Psi, x : (A_m \otimes B_m)|_U, w : A_m, y : B_m \vdash Q[w, y] :: (z : C_k)}{\Psi, x : A_m \otimes B_m \vdash (\langle w, y \rangle \leftarrow x ; Q[w, y]) :: (z : C_k)} \otimes L \\
\\
\frac{(\text{for all } \ell \in L) \quad \Psi \vdash Q_\ell[y] :: (y : A_m^\ell)}{\Psi \vdash (\text{case } x(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}) :: (x : \&\{\ell : A_m^\ell\}_{\ell \in L})} \&R \\
\\
\frac{(i \in L)}{\Psi_U, x : \&\{\ell : A_m^\ell\}_{\ell \in L} \vdash x.i(y) :: (y : A_m^i)} \&L_i^0 \\
\\
\frac{\Psi, w : A_m \vdash P[w, y] :: (y : B_m)}{\Psi \vdash (\langle w, y \rangle \leftarrow x ; P[w, y]) :: (x : A_m \multimap B_m)} \multimap R \\
\\
\frac{}{\Psi_U, w : A_m, x : A_m \multimap B_m \vdash x.\langle w, y \rangle :: (y : B_m)} \multimap L^0 \\
\\
\frac{}{\Psi_U, y : A_U \vdash x.\text{shift}(y) :: (x : \downarrow A_U)} \downarrow R^0 \qquad \frac{\Psi, y : A_U \vdash Q[y] :: (z : C_l)}{\Psi, x : \downarrow A_U \vdash (\text{shift}(y) \leftarrow x ; Q[y]) :: (z : C_l)} \downarrow L \\
\\
\frac{\Psi \vdash P[y] :: (y : A_l)}{\Psi \vdash (\text{shift}(y) \leftarrow x ; P[y]) :: (x : \uparrow A_l)} \uparrow R \qquad \frac{}{\Psi_U, x : \uparrow A_l \vdash (x.\text{shift}(y)) :: (y : A_l)} \uparrow L^0
\end{array}$$

Fig. 3. LNL<sup>†</sup> with process expressions

recursive process definitions of the form

$$\begin{array}{l}
y_1 : B_1, \dots, y_n : B_n \vdash f : (x : A) \\
x \leftarrow f \leftarrow y_1, \dots, y_n = P[x, y_1, \dots, y_n]
\end{array}$$

The first line declares the type of  $f$ , while the second defines  $f$  in terms of a process expression  $P$ . In a program we invoke it as a special form of cut

$$x \leftarrow f \leftarrow y_1, \dots, y_n ; Q[x]$$



(cutC)	$\text{proc}(c, x \leftarrow P[x]; Q[x])$	$\longrightarrow$	$\text{proc}(a, P[a]), \text{proc}(c, Q[a])$ (a fresh)
(idC)	$\text{proc}(d, P[d]), \text{proc}(c, c \leftarrow d)$	$\longrightarrow$	$\text{proc}(c, P[c])$
( $\oplus$ C)	$\text{proc}(c, c.i(d)), \text{proc}(e, \text{case } c(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L})$	$\longrightarrow$	$\text{proc}(e, Q_i[d])$
(1C)	$\text{proc}(c, c.\langle \rangle), \text{proc}(e, \langle \rangle \leftarrow c; Q)$	$\longrightarrow$	$\text{proc}(e, Q)$
( $\otimes$ C)	$\text{proc}(c, c.\langle e, d \rangle), \text{proc}(e, \langle w, y \rangle \leftarrow c; Q[w, y])$	$\longrightarrow$	$\text{proc}(e, Q[e, d])$
( $\&$ C)	$\text{proc}(c, \text{case } c(\ell(y) \Rightarrow P_\ell[y])), \text{proc}(e, c.i(d))$	$\longrightarrow$	$\text{proc}(d, P_i[d])$
( $- \circ$ C)	$\text{proc}(c, \langle w, y \rangle \leftarrow c; P[w, y]), \text{proc}(d, c.\langle e, d \rangle)$	$\longrightarrow$	$\text{proc}(d, P[e, d])$

Fig. 4. LNL<sup>†</sup> operational semantics, linear fragment

which evolves to an ordinary cut with the following rule:

$$\text{proc}(c, x \leftarrow f \leftarrow d_1, \dots, d_n; Q[x]) \longrightarrow \text{proc}(a, x \leftarrow P[x, d_1, \dots, d_n]; Q[x])$$

If there is no continuation  $Q[x]$  we have the special form of a tail call  $x \leftarrow f \leftarrow y_1, \dots, y_n$  which is a syntactic abbreviation for a cut with an identity:

$$(x \leftarrow f \leftarrow y_1, \dots, y_n) = (x' \leftarrow f \leftarrow y_1, \dots, y_n; x \leftarrow x')$$

for a fresh  $x'$ . For brevity, we abbreviate  $y_1, \dots, y_n = \bar{y}$  in some transition rules.

As a last syntactic extension we consider a version of cut where the premises are reversed.

$$\frac{\Psi_1 \geq m \geq k \quad \Psi_0, \Psi_2, x : A_m \vdash Q[x] :: (z : C_k) \quad \Psi_0, \Psi_1 \vdash P[x] :: (x : A_m)}{\Psi_0, \Psi_1, \Psi_2 \vdash (x \leftarrow Q[x]; P[x]) :: (z : C_k)} \text{cut}^{\text{rev}}$$

From a logical perspective this is of course completely redundant, but it will allow us to write programs syntactically in a more natural order. We overload the notation since it can easily be disambiguated.

## 4.2 Example: Bit Streams

Consider the recursive type

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

Any process providing a channel  $x : \text{bits}$  should send a potentially infinite sequence of messages b0 and b1. If it is finite, after the last bit it should send the label \$. The continuation is then of type 1, which means it should close the channel by sending  $\langle \rangle$ . As an example, we consider how to send the number  $6 = (110)_2$ . The representation of numbers is in “little endian” form, that is, we actually need to send the following sequence: b0, b1, b1, \$,  $\langle \rangle$ , where the last message signifies the closure of the channel. Using the reverse cut notation, we can implement the process *six*

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

as follows

```

x ← six = x1 ← x.b0(x1) ;      % send b0 along x, continue as x1
      x2 ← x1.b1(x2) ;      % send b1 along x1, continue as x2
      x3 ← x2.b1(x3) ;      % send b1 along x2, continue as x3
      x4 ← x3.$(x4) ;      % send $ along x3, continue as x4
      x4.⟨ ⟩                % close x4

```

When the comment says “send” this is of course implemented as a cut (spawn) that creates a new channel representing the continuation. Executing this program providing along some initial channel  $c_0 : \text{bits}$  yields the following configuration:

```

proc(c0, c0 ← six) →* proc(c4, c4.⟨⟩),
                        proc(c3, c3.$(c4)),
                        proc(c2, c2.b1(c3)),
                        proc(c1, c1.b1(c2)),
                        proc(c0, c0.b0(c1))

```

The resulting structure of messages can easily be seen to represent the desired message queue where the channels act as a kind of pointer. We will exploit this analogy in our shared memory semantics.

As a second example consider a program *plus1* for incrementing a bit stream. It is a transducer that takes a stream representing the number  $n$  to a stream that represents the number  $n + 1$ . The transducer *plus1* is a client to a bit stream along a channel  $y$  and provides a channel  $x$ . When  $y$  is b0 it outputs b1 along  $x$  and is finished: from then on, the output stream along  $x$  behaves exactly like the remaining input stream  $y$ . We implement this by forwarding  $x \leftarrow y$ . When  $y$  is b1 we have to output b0, but we also have to increment the rest of the stream (the “carry”) which we accomplish by a recursive call to *plus1*.

$y : \text{bits} \vdash \text{plus1} :: (x : \text{bits})$

```

x ← plus1 ← y =
  case y ( b0(y') ⇒ x' ← x.b1(x') ;      % send b1 along x, continue as x'
           x' ← y'                          % implement x' by y'
    | b1(y') ⇒ x' ← x.b0(x') ;            % send b0 along x, continue as x'
           x' ← plus1 ← y'                % continue to increment, modeling the carry bit
    | $(y') ⇒ x' ← x.b1(x') ;             % send b1 along x, continue as x'
           x'' ← x'.$(x'') ;              % send $ along x', continue as x''
           ⟨⟩ ← y' ;                       % wait for y' to close
           x''.⟨⟩                          % close x''

```

We can use the *plus1* process, for example, to compute the number 8.

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

```

x ← eight = x6 ← six ;
           x7 ← plus1 ← x6 ;
           x ← plus1 ← x7

```

We leave it to the reader to verify that this calculates the correct configuration representing the number 8. Moreover, this pipeline exhibits some parallelism as bits flow through the two *plus1* processes.

### 4.3 Example: A Binary Counter

A counter is process that can accept messages *inc* or *val*, in the form of an *external choice*. When receiving *inc* it increments its value and behaves again like a counter; when receiving *val* it turns into a stream of bits representing its current value.

```

bits = ⊕{b0 : bits, b1 : bits, $ : 1}
ctr = &{inc : ctr, val : bits}

```

The type *ctr* provides a behavioral abstraction to the client: internal representations are completely invisible. In that sense a process providing  $x : \text{ctr}$  can be thought of as an object with *methods* *inc* and *val* [Balzer and Pfenning 2015].

We implement a counter as a linear network of processes that behave like a bit 1, bit 0, or the counter with value zero. The first two provide the lowest order bit and are *clients* to the process representing the higher order bits, which is once again a little endian representation.

```

y : ctr ⊢ bit0 :: (x : ctr)
y : ctr ⊢ bit1 :: (x : ctr)
· ⊢ zero :: (x : ctr)

```

We present the implementations with short explanations in the comments. Note that here we use the ordinary rather than the reverse syntax for cut in the case of external choice.

```

x ← bit0 ← y =
  case x (inc(x') ⇒ x' ← bit1 ← y           % transition to bit1 in tail call
         | val(x') ⇒ x'' ← x'.b0(x'') ;      % respond with b0, continue as x''
           y' ← y.val(y') ;                  % request higher bits from y
           x'' ← y')                          % forward bits

x ← bit1 ← y =
  case x (inc(x') ⇒ y' ← y.inc(y') ;         % send on carry bit along y, continue as y'
         | val(x') ⇒ x'' ← x'.b1(x'') ;      % respond with b1, continue as x''
           y' ← y.val(y') ;                  % request higher bits from y
           x'' ← y')                          % forward bits

x ← zero =
  case x (inc(x') ⇒ z ← zero ;                % spawn new zero process
         | val(x') ⇒ x'' ← x'.$(x'') ;       % respond with $, continue as x''
           x''.< >)                            % close channel and terminate

```

If we start with zero and increment twice, we obtain a process network with three processes: *zero*, *bit1*, *bit0*, properly plugged together in this order.

```

· ⊢ two :: (x : ctr)

x ← two =
  z ← zero ;           % start a counter z at zero
  z' ← z.inc(z') ;    % increment z, continue as z'
  z'' ← z'.inc(z'') ; % increment z', continue as z''
  x ← z''              % implement x as z''

```

We can now execute this as expected

```

proc(c0, c0 ← two) →* proc(c2, c2 ← zero),
                      proc(c1, c1 ← bit1 ← c2),
                      proc(c0, c0 ← bit0 ← c1)

```

We see here the distinction between positive types (as in the case of bit streams), where we computed processes that can be interpreted as a message queue, and negatives types, where the configuration evolved to a network of processes that accept further messages in the style of concurrent object-oriented programming. We can convert between these two representations.

```

y : ctr ⊢ c2b :: (x : bits)

x ← c2b ← y = y' ← y.val(y') ; x ← y'

y : bits ⊢ b2c :: (x : ctr)

x ← b2c ← y' =
  case y (b0(y') ⇒ z ← b2c ← y' ;          % convert remaining stream, concurrently
         | x ← bit0 ← z)                    % lowest order bit is 0

```

$b1(y') \Rightarrow z \leftarrow b2c \leftarrow y'$ ;	% convert remaining stream, concurrently
$x \leftarrow bit1 \leftarrow z$	% lowest order bit is 1
$\$(y') \Rightarrow \langle \rangle \leftarrow y'$ ;	% wait for $y'$ to terminate
$x \leftarrow zero$	% convert empty bit stream to zero

## 5 INCORPORATING NON-LINEAR TYPES

We have presented the logical and typing rules in Figures 2 and 3 for a mixed linear and non-linear system, but we have presented the operational semantics only for the linear fragment. In this section we fill this gap and provide some examples.

A key generalization in the operational semantics is that now a process may have multiple clients. How does this come about? Consider the following example of cut in  $LNL^\dagger$ , where the modes of  $C$  and  $D$  are insignificant:

$$\frac{\frac{}{A_U \vdash A_U \oplus B_U} \oplus R_1^0 \quad \frac{\frac{Q \quad R}{A_U \oplus B_U \vdash C \quad A_U \oplus B_U, C \vdash D} \text{cut}}{A_U \oplus B_U \vdash D} \text{cut}}{A_U \vdash C} \text{cut}}$$

We should be able to commute this cut upwards so the proof in the first premise, let's call it  $P$  can interact with both  $Q$  and  $R$ . There are two approaches to permitting this. The first is to use explicit rules for contraction and weakening. In this approach, the most natural semantics duplicates or discards  $P$ . This approach has been followed by Pruikma et al. [2018b]. A second approach is to use implicit weakening and contraction as in Fig. 2 and *share*  $P$  between  $Q$  and  $R$ . We follow this here since it appears more amenable to a shared memory interpretation (see Sec. 7).

There are two key insights to model sharing elegantly in our formal semantics. The first is that we should make “messages” *persistent semantic objects*  $!\phi$ . Fortunately, this is a standard concept in substructural operational semantics and multiset rewriting. A generalized multiset rewriting rules allows *persistent premises*  $!\phi$  in the antecedents which we match against persistent semantic objects. However, the persistent objects are not removed from the configuration. We can also add new persistent semantic objects by using  $!\psi$  in the succedent. When we need to distinguish explicitly we will refer to the ordinary objects in the semantics as *ephemeral*.

The second insight is that we cannot simply make all processes  $\text{proc}(c_U, P)$  that provide along a shared (that is, non-linear) channel persistent. For example, say that  $P$  starts with a cut. If this object were persistent, the cut would spawn unboundedly many processes since the operational rule would continue to be applicable. We solve this problem by introducing three forms of semantic objects

- (1)  $\text{proc}(c_m, P)$  which is always ephemeral and evolves according the operational semantics,
- (2)  $\text{msg}(c_m, P)$  which represents a message and is persistent if  $m = U$ , otherwise ephemeral,
- (3)  $\text{srv}(c_m, P)$  which represents a service and is persistent if  $m = U$ , otherwise ephemeral.

The revised and generalized operational semantics is presented in Fig. 5. We write generically  $c.M$  for a process sending along  $c$ , and  $P\langle c \rangle$  for a process receiving along  $c$ . The precise syntax is summarized at the beginning of Sec. 6. We also use  $!_m\phi$  as an abbreviation for an ephemeral  $\phi$  if  $m = L$  and a persistent  $!\phi$  if  $m = U$ . In the rules for the shift the modes are predetermined by the typing rules, so we have made them explicit there.

An object  $!\text{msg}(c_U, c_U.M)$  represents a *multicast*, since this message is persistent will therefore be able to interact with all clients of  $c_U$ . The sender may not even be aware of all of its clients. Conversely, an object  $!\text{srv}(c_U, P\langle c_U \rangle)$  corresponds to a *copy-on-recv* since a fresh copy of  $P$ 's continuation will be spawned when this server interacts with one of its clients.

- (cutC)  $\text{proc}(c, x_m \leftarrow P[x_m] ; Q[x_m]) \longrightarrow \text{proc}(a_m, P[a_m]), \text{proc}(c, Q[a_m])$  ( $a_m$  fresh)
- (idC<sub>1</sub>)  $!_m \text{msg}(d_m, d_m.M), \text{proc}(c_m, c_m \leftarrow d_m) \longrightarrow !_m \text{msg}(c_m, c_m.M)$
- (idC<sub>2</sub>)  $!_m \text{srv}(d_m, P\langle d_m \rangle), \text{proc}(c_m, c_m \leftarrow d_m) \longrightarrow !_m \text{srv}(c_m, P\langle c_m \rangle)$
- (call)  $\text{proc}(c, x \leftarrow f \leftarrow \bar{d} ; Q[x]) \longrightarrow \text{proc}(a, x \leftarrow P[x, \bar{d}] ; Q[x])$   
 when  $z \leftarrow f \leftarrow \bar{y} = P[z, \bar{y}]$
- (msg)  $\text{proc}(c_m, c.M) \longrightarrow !_m \text{msg}(c_m, c.M)$
- (srv)  $\text{proc}(c_m, P\langle c \rangle) \longrightarrow !_m \text{srv}(c_m, P\langle c \rangle)$
- ( $\oplus$ C)  $!_m \text{msg}(c_m, c.i(d)), \text{proc}(e, \text{case } c_m(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}) \longrightarrow \text{proc}(e, Q_i[d])$
- (1C)  $!_m \text{msg}(c_m, c.\langle \rangle), \text{proc}(e, \langle \rangle \leftarrow c ; Q) \longrightarrow \text{proc}(e, Q)$
- ( $\otimes$ C)  $!_m \text{msg}(c_m, c.\langle b, d \rangle), \text{proc}(e, \langle w, y \rangle \leftarrow c ; Q[w, y]) \longrightarrow \text{proc}(e, Q[b, d])$
- ( $\&$ C)  $!_m \text{srv}(c_m, \text{case } c(\ell(y) \Rightarrow P_\ell[y])), \text{proc}(d, c.i(d)) \longrightarrow \text{proc}(d, P_i[d])$
- ( $\multimap$ C)  $!_m \text{srv}(c_m, \langle w, y \rangle \leftarrow c ; P[w, y]), \text{proc}(d, c.\langle b, d \rangle) \longrightarrow \text{proc}(d, P[b, d])$
- ( $\downarrow$ C)  $\text{msg}(c_\downarrow, c_\downarrow.\text{shift}(d_\downarrow)), \text{proc}(e, \text{shift}(y_\downarrow) \leftarrow c_\downarrow ; Q[y_\downarrow]) \longrightarrow \text{proc}(e, Q[d_\downarrow])$
- ( $\uparrow$ C)  $!_\downarrow \text{srv}(c_\downarrow, \text{shift}(y_\downarrow) \leftarrow c_\downarrow ; P[y_\downarrow]), \text{proc}(d_\downarrow, c_\downarrow.\text{shift}(d_\downarrow)) \longrightarrow \text{proc}(d_\downarrow, P[d_\downarrow])$

Fig. 5. LNL<sup>†</sup> operational semantics  
 $!_m \phi = \phi$  if  $m = L$  and  $!_m \phi = !\phi$  if  $m = U$

The identity is also interesting. It seems possible to convert an object  $\text{proc}(c_m, c_m \leftarrow d_m)$  into a message or service (depending on the polarity of the type of  $c_m$ ). This slightly awkward approach would generalize the (idC') rule we discarded. The (idC) rules seems more amenable for modification for mixed linear and non-linear channels in that it could either interact with a message (receive a message) or a server. A complication the revised rules need to account for is that both  $d_m$  and  $c_m$  may have multiple clients if  $m = U$ . When  $m = L$  it specializes essentially to the previous rule, where the message/service distinction is redundant.

- (idC<sub>1</sub>)  $!_m \text{msg}(d_m, d_m.M), \text{proc}(c_m, c_m \leftarrow d_m) \longrightarrow !_m \text{msg}(c_m, c_m.M)$
- (idC<sub>2</sub>)  $!_m \text{srv}(d_m, P\langle d_m \rangle), \text{proc}(c_m, c_m \leftarrow d_m) \longrightarrow !_m \text{srv}(c_m, P\langle c_m \rangle)$

In idC<sub>2</sub> it looks as if we have to copy  $P$  if  $m = U$ , but in an implementation presumably a single process can remember it needs to serve two (or, in general,  $n$ ) channels. Copying messages (which are always small) to send along  $c_m$  in idC<sub>1</sub> is more straightforward.

### 5.1 Example: Circuits

We call channels  $c_\downarrow$  that are subject to weakening and contraction *shared channels*. As an example that requires shared channels we use circuits. We start by programming a nor gate that processes infinite streams of zeros and ones.

$$\text{bits}_\downarrow^\infty = \oplus \{b_0 : \text{bits}_\downarrow^\infty, b_1 : \text{bits}_\downarrow^\infty\}$$

$$x : \text{bits}_\downarrow^\infty, y : \text{bits}_\downarrow^\infty \vdash \text{nor} :: (z : \text{bits}_\downarrow^\infty)$$

$$z \leftarrow \text{nor} \leftarrow x, y =$$

$$\begin{aligned} &\text{case } x \text{ ( } b_0(x') \Rightarrow \text{case } y \text{ ( } b_0(y') \Rightarrow z' \leftarrow z.b_1(z') ; \\ &\quad \quad \quad z' \leftarrow \text{nor} \leftarrow x', y' \\ &\quad \quad \quad | } b_1(y') \Rightarrow z' \leftarrow z.b_0(z') ; \end{aligned}$$

$$\begin{aligned}
& z' \leftarrow \text{nor} \leftarrow x', y' \\
| \text{b1}(x') \Rightarrow & \text{case } y \text{ (b0}(y') \Rightarrow z' \leftarrow z.\text{b0}(z') ; \\
& z' \leftarrow \text{nor} \leftarrow x', y' \\
| \text{b1}(y') \Rightarrow & z' \leftarrow z.\text{b0}(z') ; \\
& z' \leftarrow \text{nor} \leftarrow x', y' )
\end{aligned}$$

This is somewhat verbose, but note that all channels here are shared. For this particular gate they could also be linear because neither is reused, but in this paper we do not treat this mode polymorphism. This illustrates that programming can be uniform at different modes, which is a significant advantage of LNL over linear logic with an exponential  $!A$ . Our implementation of *nor* has the property that for bits  $A$ ,  $B$ , and  $C$  with  $C = \neg(A \vee B)$ , the following transitions are possible and characterize *nor*:

$$\begin{aligned}
& !\text{msg}(a, a.A(a')), !\text{msg}(b, b.B(b')), \text{proc}(c, c \leftarrow \text{nor} \leftarrow a, b) \\
\longrightarrow^* & \text{proc}(c', c' \leftarrow \text{nor} \leftarrow a', b'), !\text{msg}(c, c.C(c')) \quad (c' \text{ fresh})
\end{aligned}$$

The persistent messages on shared channels  $a$  and  $b$  continue to be available after the transitions: as shared channels they may be used elsewhere. A form of distributed garbage collection, reference counting, or a more precise semantics (see [Pruiksmā et al. \[2018b\]](#)) would be necessary to eliminate unreachable persistent semantic objects.

When we build an or-gate out of a nor-gate we need to exploit sharing to implement simple negation.

$$x : \text{bits}_U^\infty, y : \text{bits}_U^\infty \vdash \text{or} :: (z : \text{bits}_U^\infty)$$

$$\begin{aligned}
z \leftarrow \text{or} \leftarrow x, y = \\
u \leftarrow \text{nor} \leftarrow x, y \\
z \leftarrow \text{nor} \leftarrow u, u
\end{aligned}$$

An analogous computation to the above is possible, except that a shared intermediate channel  $d$  with  $!\text{msg}(d, d.D(d'))$  will also be created with  $D = \neg(A \vee B)$  and  $C = \neg D = A \vee B$ .

$$\begin{aligned}
& !\text{msg}(a, a.A(a')), !\text{msg}(b, b.B(b')), \\
& \text{proc}(c, c \leftarrow \text{or} \leftarrow a, b) \\
\longrightarrow^* & !\text{msg}(d, d.D(d')) \\
& \text{proc}(d', d' \leftarrow \text{nor} \leftarrow a', b'), \\
& \text{proc}(c', c' \leftarrow \text{nor} \leftarrow d', d'), \\
& !\text{msg}(c, c.C(c')) \quad (c', d' \text{ fresh})
\end{aligned}$$

## 5.2 Example: Map

Mapping a process over a list exemplifies several new ideas: a mixed linear/non-linear program, parametric definitions, and multiplicatives. We define a whole family of types indexed by a type  $A$ , which is not formally part of the language but is expressed at the metalevel.

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

Such a list should not be viewed as a data structure in memory. Instead, it is a behavioral description of a stream of messages. A process that maps a channel of type  $A$  to one of type  $B$  will itself have type  $A \multimap B$ . However, this process must be shared since it needs to be applied to every element. We therefore obtain the following type and definition, where all channels not annotated are linear.

$$f_U : \uparrow(A_L \multimap B_L), l : \text{list}_A \vdash \text{map} :: (k : \text{list}_B)$$

$$\begin{aligned}
k \leftarrow \text{map} \leftarrow f_U, l = \\
\text{case } l \text{ (cons}(l') \Rightarrow \langle x, l'' \rangle \leftarrow l' ; & \quad \% \text{ receive element } x : A \text{ with continuation } l'' \\
& f' \leftarrow f_U.\text{shift}(f') ; & \% \text{ obtain a fresh linear instance } f' \text{ of } f_U \\
& y \leftarrow f'.\langle x, y \rangle ; & \% \text{ send } x \text{ to } f', \text{ response will be along fresh } y
\end{aligned}$$

```

      k' ← k.cons(k') ;    % select cons
      k'' ← k'.⟨y, k''⟩ ; % send y with continuation k''
      k'' ← map ← fU, l'' % recurse with continuation channels
| nil(l') ⇒ k' ← k.nil(k') ;    % select nil
      ⟨⟩ ← l' ;                % wait for l' to close
      k'.⟨⟩                    % close k' and terminate

```

There is some potential parallelism here: each instance of  $f$  can run concurrently with mapping over the remainder of the list since only the result channel  $y$  (which is locally created as a destination in  $map$ ) is communicated to the client of  $map$ .

As a simple parametric use of this, consider a process  $id_A$  that behaves as the identity.

$$\vdash id :: (f_U : \uparrow(A \multimap A))$$

```

fU ← id =
  shift(f') ← fU ;          % obtain fresh linear channel f'
  ⟨x, f''⟩ ← f' ;           % receive x and continuation f''
  f'' ← x                    % forward x to f'' and terminate this instance of id

```

$$k : list_A \vdash copy :: (l : list_A)$$

```

l ← copy ← k =
  fU ← id                % start persistent identity process
  l ← map ← fU, k       % map fU over k to obtain l

```

## 6 METATHEORY

Since our LNL<sup>†</sup> is very directly based on a variant of the sequent calculus the metatheoretical properties we care about are not particularly difficult to establish. We define two equivalent notions of configuration typing, one more suitable for session fidelity, the other more suitable for deadlock freedom.

We summarize the grammar of messages and services. Recall that message objects have the form  $!_m \text{msg}(c_m, c_m.M)$  (representing the right rule of a positive type  $\oplus, 1, \otimes, \downarrow$ ) and service objects  $!_m \text{srv}(c_m, P\langle c \rangle)$  (representing the right rule of a negative type  $\&, \multimap, \uparrow$ ).

```

Messages  M    ::= i(d) | ⟨⟩ | ⟨b, d⟩ | shift(d)
Services  P⟨c⟩ ::= case c(ℓ(y) ⇒ Qℓ[y])ℓ∈L | ⟨z, y⟩ ← c ; Q[z, y] | shift(y) ← c ; Q[y]

```

### 6.1 Session Fidelity

We define  $\Psi' \vdash C :: \Psi$  to express that the collection of semantic objects in  $C$  provides all the channels in  $\Psi$  and uses all the channels in  $\Psi'$ . While configurations  $C$  are unordered (subject to exchange) their typing derivations analyze them in an order where each provider  $\phi(c, P)$  precedes

all clients of  $c$ .

$$\begin{array}{c} \overline{\Psi \vdash (\cdot) :: \Psi} \text{ id} \quad \frac{\Psi_0 \vdash C_1 :: \Psi_1 \quad \Psi_1 \vdash C_2 :: \Psi_2}{\Psi_0 \vdash (C_1, C_2) :: \Psi_2} \text{ compose} \\ \\ \frac{\Psi \geq m \quad \Psi_U, \Psi \vdash P :: (c_m : A_m)}{\Psi_U, \Psi', \Psi \vdash \text{proc}(c, P) :: (\Psi_U, \Psi', c_m : A_m)} \text{ proc} \\ \\ \frac{\Psi \geq m \quad \Psi_U, \Psi \vdash c_m.M :: (c_m : A_m)}{\Psi_U, \Psi', \Psi \vdash !_m \text{msg}(c_m, c_m.M) :: (\Psi_U, \Psi', c : A_m)} \text{ msg} \\ \\ \frac{\Psi \geq m \quad \Psi_U, \Psi \vdash P\langle c_m \rangle :: (c_m : A_m)}{\Psi_U, \Psi', \Psi \vdash !_m \text{srv}(c_m, P\langle c_m \rangle) :: (\Psi_U, \Psi', c_m : A_m)} \text{ srv} \end{array}$$

One key aspect of the rules is that persistent semantic objects  $! \text{msg}$  and  $! \text{srv}$  may only depend on shared channels  $d_U : B_U$ . The other is that shared channels in  $\Psi_U$  may be used in  $P$  but also by other clients because it also appears in the right-hand context.

Recall the presupposition that in a configuration  $C$  all channels  $c$  with  $\phi(c, P)$  are distinct. This ensures that there is no confusion among the typing of channels, including linear channels that do not appear in the external interface to a configuration  $C$  because they are provided by one object in  $C$  and used by another.

**THEOREM 6.1 (SESSION FIDELITY).** *If  $\Psi' \vdash C :: \Psi$  and  $C \longrightarrow \mathcal{D}$  then  $\Psi' \vdash \mathcal{D} :: (\Psi, \Psi_U)$  where  $\Psi_U$  may be empty or consist of a fresh channel  $a_U$  not in  $\Psi$  or  $\Psi'$ .*

**PROOF.** By cases over the possible transitions. In each case we find the corresponding objects in  $C$  and apply inversion to the typing of the object to infer the typing of the continuations. These are then sufficient to type the objects on the right-hand sides of the transitions and insert them into the evidently correct place into the typing derivation. With linear channels and ephemeral objects, the process, message, or service is removed together with the client. With shared channels, the persistent message or service remains in place so it can type the remaining clients. When executing a cut that creates a new shared channel  $a_U$ , this channel persists to the very right-hand side of the configuration and populates  $\Psi_U$  in the theorem statement.  $\square$

In order to prove freedom from deadlocks, that is, global progress, we introduce another typing for closed configurations that provides a suitable basis for the induction.

$$\begin{array}{c} \overline{\models (\cdot) :: (\cdot)} \text{ id} \\ \\ \frac{\models C :: (\Psi_U, \Psi', \Psi) \quad \Psi \geq m \quad \Psi_U, \Psi \vdash P :: (c_m : A_m)}{\models (C, \text{proc}(c, P)) :: (\Psi_U, \Psi', c_m : A_m)} \text{ proc} \\ \\ \frac{\models C :: (\Psi_U, \Psi', \Psi) \quad \Psi \geq m \quad \Psi_U, \Psi \vdash c_m.M :: (c_m : A_m)}{\models (C, !_m \text{msg}(c_m, c_m.M)) :: (\Psi_U, \Psi', c_m : A_m)} \text{ msg} \\ \\ \frac{\models C :: (\Psi_U, \Psi', \Psi) \quad \Psi \geq m \quad \Psi_U, \Psi \vdash P\langle c_m \rangle :: (c_m : A_m)}{\models (C, !_m \text{srv}(c_m, P\langle c_m \rangle)) :: (\Psi_U, \Psi', c_m : A_m)} \text{ srv} \end{array}$$

**LEMMA 6.2 (TYPING EQUIVALENCE).**  $\cdot \vdash C :: \Psi$  iff  $\models C :: \Psi$

**PROOF.** By simple inductions, re-associating the configuration typing to the left.  $\square$



Messages and services are blocked on the channel that they provide, waiting for a process to interact with. They therefore play somewhat the role of “values” in functional languages.

**THEOREM 6.3 (DEADLOCK FREEDOM).** *If  $\models C :: \Psi$  then*

- (i) *either  $C \longrightarrow \mathcal{D}$  for some  $\mathcal{D}$ , or*
- (ii)  *$C$  consists entirely of objects  $!_m \text{msg}(c_m, c_m.M)$  or  $!_m \text{srv}(c_m, P(c_m))$ .*

**PROOF.** By induction on the structure of the given derivation. For nonempty configurations  $C$  we single out the rightmost process, message, or service  $C = (C', \phi)$  and apply the induction hypothesis to  $C'$ . If  $C'$  or  $\phi$  can step by themselves, so can  $C$ . Otherwise  $C'$  consists entirely of messages and services. At this point, we apply inversion to the typing of  $\phi = \text{proc}(c, P)$ , which must be trying to communicate along a channel  $d$  it uses. Again by inversion, we find the provider of  $d$  is a matching message (positive type) or service (negative type) and the interaction can take place.  $\square$

Session fidelity and deadlock freedom support some simple corollaries. We call a configuration  $\mathcal{F}$  *final* if  $\mathcal{F}$  cannot make a transition. Then, if  $\text{proc}(c_0, P)$  with  $\cdot \vdash P :: (c_0 : \mathbf{1}_L)$  and  $\text{proc}(c_0, P) \longrightarrow^* \mathcal{F}$  for a final  $\mathcal{F}$ , then  $\mathcal{F} = (!\mathcal{F}', \text{msg}(c_0, c_0.\langle \rangle))$  where every object in  $\mathcal{F}$  is a persistent message or service.

## 7 A SHARED MEMORY SEMANTICS

The semantics of  $\text{LNL}^\dagger$  was carefully designed so that we did not need to send any code objects across the network, only concrete data such as labels or, for indirect access, channels. With a slight shift of perspective we can reuse almost the exact semantics to obtain a feasible shared memory implementation of  $\text{LNL}^\dagger$ .

The first, and perhaps obvious idea is to interpret channels as memory addresses. Since only cut allocates new channels, only cut allocates memory. That is perhaps already not entirely expected.

As for the operational semantics, based on the information contents of proofs, we have so far thought of the right rules for positive connectives ( $\oplus$ ,  $\mathbf{1}$ ,  $\otimes$ ,  $\downarrow$ ) and the left rules for negative connectives ( $\otimes$ ,  $-\circ$ ,  $\uparrow$ ) as sending. A reasonable expectation might be that sending rules *write* a message to memory and, conversely, receiving rules *read* a message from memory. This approach, however, while feasible in the purely linear fragment breaks down in the presence of channels with multiple clients. Instead we follow these simple principles:

- (1) *Cut allocates*
- (2) *Right rules write*
- (3) *Left rules read*
- (4) *Identity moves or copies*

We begin by unifying the notions of message and service into the notion of a *memory cell*. We have new semantic objects of the form  $!_m \text{cell}(c_m, V^?)$  which may be ephemeral (for  $m = L$ ) or persistent (for  $m = U$ ). The contents of the cell can be ‘ $\_$ ’ (it is empty) or  $V$  (it has been filled by a value  $V$ ). The continuation channel in many process expressions becomes the address of the continuation of the data structure in memory. We may refer to this as a pointer or *continuation*

reference. Values then follow this grammar:

Pos. Vals.	$V^+ ::=$	$i(d)$	$(\oplus)$	label $k$ with cont. ref. $d$
		$\langle \rangle$	$(\mathbf{1})$	final value
		$\langle c, d \rangle$	$(\otimes)$	address $c$ and cont. ref. $d$
		$\text{shift}(d_u)$	$(\downarrow)$	ephemeral reference to persistent cell $d_u$
Neg. Vals.	$V^- ::=$	$(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}$	$(\&)$	dispatch table based on label $\ell$ with cont. ref. $y$
		$\langle x, y \rangle \Rightarrow Q[x, y]$	$(\multimap)$	parameterized program continuation
		$\text{shift}(y_L) \Rightarrow Q[y_L]$	$(\uparrow)$	persistent program continuation
Values	$V ::=$	$V^+ \mid V^-$		

We now rewrite the previous transition rules using cells. First, cut allocates a fresh cell.

$$(\text{cutC}) \quad \text{proc}(c, x_m \leftarrow P[x_m] ; Q[x_m]) \longrightarrow \text{cell}(a_m, \_), \text{proc}(a_m, P[a_m]), \text{proc}(c, Q[a_m]) \quad (a_m \text{ fresh})$$

The object  $\text{cell}(a_m, \_)$  here is ephemeral independent of the mode  $m$  because it has been allocated but not yet initialized.

All the rules maintain the invariant that if there is an object  $\text{proc}(a, P)$  then we have a corresponding cell  $(a, \_)$ . We think of  $a$  as the *destination* for the value of  $P$ . By the very nature of the system, all computations have a destination.

Processes whose destination is of positive type will write to the destination cell. This is independent of whether the address is of linear or shared mode.

$$\begin{aligned} (\oplus W) \quad & \text{cell}(c_m, \_), \text{proc}(c_m, c.i(d)) \longrightarrow !_m \text{cell}(c_m, i(d)) \\ (\oplus D) \quad & !_m \text{cell}(c_m, i(d)), \text{proc}(e, \text{case } c_m (\ell(y) \Rightarrow Q_\ell[y])) \longrightarrow \text{proc}(e, Q_k[d]) \end{aligned}$$

When a cell is written to (here in the  $\oplus W$  rule) then it becomes persistent if the address is shared. The process reading from a cell  $c_m$  may block until the cell is initialized, which could be implemented via a lock or other low level shared memory primitives. Recall also that persistent objects  $!\phi$  will remain in the configuration unchanged so it can interact with other potential clients.

As an example of a negative type, consider the transition for type  $A \multimap B$ .

$$\begin{aligned} (\multimap W) \quad & \text{cell}(c_m, \_), \text{proc}(c_m, \langle z, y \rangle \leftarrow c ; Q[z, y]) \longrightarrow !_m \text{cell}(c_m, \langle z, y \rangle \Rightarrow Q[z, y]) \\ (\multimap D) \quad & !_m \text{cell}(c_m, \langle z, y \rangle \Rightarrow Q[z, y]), \text{proc}(d, c.\langle b, d \rangle) \longrightarrow \text{proc}(d, Q[b, d]) \end{aligned}$$

The process executing a receive under the message-passing semantics instead writes a continuation expression to memory. At a lower level of abstraction, we would expect the implementation to construct a closure and store a pointer to the closure in  $c_m$ . Again, writing to the cell means it takes on the structural properties of its address. Reading from the cell instantiates it with the argument channel and continuation address and executes the continuation.

Writing a closure to memory is a luxury afforded by shared memory: we should not expect in general to be able send a closure as a primitive operation in the message-passing setting.

Finally, we consider forwarding. The two different rules of the message-passing semantics in Fig. 5 are unified to

$$(\text{idC}) \quad !_m \text{cell}(d_m, V), \text{cell}(c_m, \_), \text{proc}(c_m, c_m \leftarrow d_m) \longrightarrow !_m \text{cell}(c_m, V)$$

Depending on  $m$  this either *moves* the value  $V$  (when  $m = L$ ) or *copies* the value  $V$  from cell  $c$  to  $d$  (when  $m = U$ ). In either case, values  $V$  should be expected to be small (pairs of labels and pointers, possibly a pointer to a closure), so this operation is meaningful from an (abstract) efficiency standpoint.

Since there are two forms of identity reduction in the sequent calculus, we would expect an alternative semantics to exist. Indeed, a configuration  $\text{cell}(c_m, \_), \text{proc}(c_m, c_m \leftarrow d_m)$  could transition to  $!\_m \text{cell}(c_m, \text{FWD}(d_m))$ . Then, references to  $c_m$  have to follow chains of forwarding pointers.

- (cutC)  $\text{proc}(c, x_m \leftarrow P[x_m] ; Q[x_m]) \longrightarrow \text{cell}(a_m, \_), \text{proc}(a_m, P[a_m]), \text{proc}(c, Q[a_m])$  ( $a_m$  fresh)
- (idC)  $!_m \text{cell}(d_m, V), \text{cell}(c_m, \_), \text{proc}(c_m, c_m \leftarrow d_m) \longrightarrow !_m \text{cell}(c_m, V)$
- (call)  $\text{proc}(c, x \leftarrow f \leftarrow \bar{d} ; Q[x]) \longrightarrow \text{proc}(a, x \leftarrow P[x, \bar{d}] ; Q[x])$   
 when  $z \leftarrow f \leftarrow \bar{y} = P[z, \bar{y}]$
- (write)  $\text{cell}(c_m, \_), \text{proc}(c_m, c @ V) \longrightarrow !_m \text{cell}(c_m, V)$
- ( $\oplus D$ )  $!_m \text{cell}(c_m, i(d)), \text{proc}(e, \text{case } c_m (\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L}) \longrightarrow \text{proc}(e, Q_i[d])$
- ( $1D$ )  $!_m \text{cell}(c_m, \langle \rangle), \text{proc}(e, \langle \rangle \leftarrow c ; Q) \longrightarrow \text{proc}(e, Q)$
- ( $\otimes D$ )  $!_m \text{cell}(c_m, \langle b, d \rangle), \text{proc}(e, \langle w, y \rangle \leftarrow c ; Q[w, y]) \longrightarrow \text{proc}(e, Q[b, d])$
- ( $\& D$ )  $!_m \text{cell}(c_m, (\ell(y) \Rightarrow P_\ell[y])), \text{proc}(c, c.i(d)) \longrightarrow \text{proc}(d, P_i[d])$
- ( $\multimap D$ )  $!_m \text{cell}(c_m, \langle w, y \rangle \Rightarrow P[w, y]), \text{proc}(d, c.\langle b, d \rangle) \longrightarrow \text{proc}(d, P[b, d])$
- ( $\downarrow D$ )  $\text{cell}(c_\downarrow, \text{shift}(d_\downarrow)), \text{proc}(e, \text{shift}(y_\downarrow) \leftarrow c_\downarrow ; Q[y_\downarrow]) \longrightarrow \text{proc}(e, Q[d_\downarrow])$
- ( $\uparrow D$ )  $!_\uparrow \text{cell}(c_\uparrow, \text{shift}(y_\uparrow) \Rightarrow P[y_\uparrow]), \text{proc}(d_\uparrow, c_\uparrow.\text{shift}(d_\uparrow)) \longrightarrow \text{proc}(d_\uparrow, P[d_\uparrow])$

Fig. 6. LNL<sup>†</sup> shared memory semanticsSee Fig. 7 for the definition of  $c @ V$ 

Besides tag checking (cell is empty, has a value, or has a forwarding pointer) these forwarding cells may accumulate when persistent. Moreover, the metatheory is less satisfying since the inversion properties in this formulation are weaker, so we prefer the version we have presented.

The complete set of rules can be found in Fig. 6. We have unified the rules that write using the notation introduced in Sec. 7.4.

### 7.1 Example: Bit Streams Revisited

Bit streams defined in Sec. 4.2 are an example of a *purely positive type*.

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

This means if we have any program  $f$  such that

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

and we execute a terminating program  $\text{proc}(c_0, c_0 \leftarrow f)$  the final configuration will consist of a linked list of bits, plus potentially some persistent (unreachable) cells. We take the same program *six* but change the comments to describe the shared memory behavior.

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

$$\begin{aligned} x \leftarrow \text{six} = & x_1 \leftarrow x.b0(x_1) ; & \% \text{ allocate } x_1 \text{ and write } b0(x_1) \text{ to } x \\ & x_2 \leftarrow x_1.b1(x_2) ; & \% \text{ allocate } x_2 \text{ and write } b1(x_2) \text{ to } x_1 \\ & x_3 \leftarrow x_2.b1(x_3) ; & \% \text{ allocate } x_3 \text{ and write } b1(x_3) \text{ to } x_2 \\ & x_4 \leftarrow x_3.\$(x_4) ; & \% \text{ allocate } x_4 \text{ and write } \$(x_4) \text{ to } x_3 \\ & x_4.\langle \rangle & \% \text{ write } \langle \rangle \text{ to } x_4 \end{aligned}$$

The write operations do not need to happen in the order described (we still have concurrency!), but when terminated we will have reached the configuration indicated below:

$$\begin{aligned} \text{proc}(c_0, c_0 \leftarrow \text{six}) \longrightarrow &^* \text{cell}(c_4, \langle \rangle), \\ & \text{cell}(c_3, \$(c_4)), \\ & \text{cell}(c_2, b1(c_3)), \\ & \text{cell}(c_1, b1(c_2)), \\ & \text{cell}(c_0, b0(c_1)) \end{aligned}$$

Surprisingly, we have the exact same program and outcome if we defined

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

except that all the cells would be persistent. If we had a transducer such as the *plus1* process, then the difference in modes would manifest itself in that ephemeral cells are deallocated when read while persistent cells, well, are persistent and would therefore have to be deallocated by a garbage collector.

## 7.2 Example Revisited: Circuits

The example of circuits over infinite bit streams from Sec. 5.1 can be transliterated into this new semantics, which forces addresses to be shared and therefore channels to be persistent.

$$\text{bits}_U^\infty = \oplus \{b0 : \text{bits}_U^\infty, b1 : \text{bits}_U^\infty\}$$

$$x : \text{bits}_U^\infty, y : \text{bits}_U^\infty \vdash \text{nor} :: (z : \text{bits}_U^\infty)$$

$$x : \text{bits}_U^\infty, y : \text{bits}_U^\infty \vdash \text{or} :: (z : \text{bits}_U^\infty)$$

$$z \leftarrow \text{or} \leftarrow x, y =$$

$$u \leftarrow \text{nor} \leftarrow x, y$$

$$z \leftarrow \text{nor} \leftarrow u, u$$

We do not repeat the definition of *nor*, but the *or* program now has the following effect on the configuration (which includes memory cells):

$$!\text{cell}(a, A(a')), !\text{cell}(b, B(b')), \text{proc}(c, c \leftarrow \text{or} \leftarrow a, b)$$

$$\longrightarrow^* !\text{cell}(d, D(d'))$$

$$\text{cell}(d', \_), \text{proc}(d', d' \leftarrow \text{nor} \leftarrow a', b'),$$

$$\text{cell}(c', \_), \text{proc}(c', c' \leftarrow \text{nor} \leftarrow d', d'),$$

$$!\text{cell}(c, C(c')) \quad (d, d', c' \text{ fresh})$$

## 7.3 Example Revisited: Map

Map provides an example of a negative linear type  $A \multimap B$  and a negative shared type  $\uparrow(A \multimap B)$ . The program is exactly the same as given in Sec. 5.2, so we do not repeat it here, only the type definition and process declaration.

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

$$f_U : \uparrow(A_L \multimap B_L), l : \text{list}_A \vdash \text{map} :: (k : \text{list}_B)$$

If we execute a program of type  $\cdot \vdash g :: (x : \text{list}_A)$  and it terminates, it will construct a memory representation of a list in the form

$$\text{cell}(c_0, \text{cons}(c_1)),$$

$$\text{cell}(c_1, \langle a_1, c_2 \rangle),$$

$$\text{cell}(c_2, \text{cons}(c_3)),$$

$$\text{cell}(c_3, \langle a_2, c_4 \rangle),$$

...

$$\text{cell}(c_{2n+1}, \text{nil}(c_{2n+1})),$$

$$\text{cell}(c_{2n+2}, \langle \rangle)$$

where  $a_1, \dots, a_n$  are the addresses of the list elements of type  $A$ . When composed with a process such as

$$\text{cell}(d_0, \_), \text{cell}(d_0, d_0 \leftarrow \text{copy} \leftarrow c_0)$$

then we transition to a final state containing an exact replica of the cells from the beginning in newly allocated cells, together with a persistent reference to the identity process waiting for a shift.

$$\begin{aligned}
c @ i(d) &= c.i(d) \\
c @ \langle \rangle &= c.\langle \rangle \\
c @ \langle b, d \rangle &= c.\langle b, d \rangle \\
c @ \text{shift}(d_u) &= c.\text{shift}(d_u) \\
c @ (\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L} &= \text{case } c(\ell(y) \Rightarrow Q_\ell[y])_{\ell \in L} \\
c @ \langle z, y \rangle \Rightarrow Q[z, y] &= \langle z, y \rangle \leftarrow c ; Q[z, y] \\
c @ \text{shift}(y_L) \Rightarrow Q[y_L] &= \text{shift}(y_L) \leftarrow c ; Q[y_L]
\end{aligned}$$

Fig. 7. Value composition and decomposition

$\text{cell}(d_0, \text{cons}(d_1)), \text{cell}(d_1, \langle a_1, d_2 \rangle), \dots, \text{cell}(d_{2n+2}, \langle \rangle),$   
 $!\text{cell}(f_u, \text{shift}(f') \Rightarrow \langle x, f'' \rangle \leftarrow f' ; f'' \leftarrow x)$

The *map* process repeatedly invoked this by providing freshly allocated (linear) destinations for  $f'$ .

#### 7.4 Metatheory Revisited

The development in Sec. 6 carries over with very few changes. We use cells instead of messages and services, while identity and composition remain unchanged. The typing rule for process object changes to

$$\frac{\Psi \geq m \quad \Psi_U, \Psi \vdash P :: (c_m : A_m)}{\Psi_U, \Psi', \Psi \vdash^M \text{cell}(c, \_), \text{proc}(c, P) :: (\Psi_U, \Psi', c_m : A_m)} \text{proc}$$

$$\frac{\Psi \geq m \quad \Psi_U, \Psi \vdash c @ V : (c_m : A_m)}{\Psi_U, \Psi', \Psi \vdash^M !_m \text{cell}(c_m, V) :: (\Psi_U, \Psi', c_m : A_m)} \text{cell}$$

The easiest way to type values is to turn them back into process expression which is described in Fig. 7.

We obtain session fidelity, which we now call type preservation since we are no longer in a message-passing setting.

**THEOREM 7.1 (TYPE PRESERVATION).** *If  $\Psi' \vdash^M C :: \Psi$  and  $C \longrightarrow \mathcal{D}$  then  $\Psi' \vdash^M \mathcal{D} :: (\Psi, \Psi_U)$  where  $\Psi_U$  may be empty or consist of a fresh address not in  $\Psi$  or  $\Psi'$ .*

**PROOF.** As in session fidelity (Thm 6.1). □

Similarly, we modify the typing of for the progress theorem.

$$\frac{}{\models^M (\cdot) :: (\cdot)} \text{id}$$

$$\frac{\models^M C :: (\Psi_U, \Psi', \Psi) \quad \Psi \geq m \quad \Psi_U, \Psi \vdash P :: (c_m : A_m)}{\models^M (C, \text{cell}(c, \_), \text{proc}(c, P)) :: (\Psi_U, \Psi', c_m : A_m)} \text{proc}$$

$$\frac{\models^M C :: (\Psi_U, \Psi', \Psi) \quad \Psi \geq m \quad \Psi_U, \Psi \vdash c_m @ V :: (c_m : A_m)}{\models^M (C, !_m \text{cell}(c_m, V)) :: (\Psi_U, \Psi', c_m : A_m)} \text{cell}$$

The two forms of typing are equivalent, as before, but we elide the straightforward statement and proof.

**THEOREM 7.2 (GLOBAL PROGRESS).** *If  $\models^M C :: \Psi$  then*

- (i) *either  $C \longrightarrow \mathcal{D}$  for some  $\mathcal{D}$ , or*
- (ii)  *$C$  consists entirely of memory cells  $!\_m \text{cell}(c_m, V)$*

PROOF. As for deadlock freedom (Thm. 6.3).  $\square$

We can also ask about the relationship between the message passing and shared memory semantics. Fortunately, we have carefully constructed them so that is very easy. We define  $C \sim \mathcal{D}$

$$\begin{aligned} \text{proc}(c, P) &\sim \text{cell}(c, \_), \text{proc}(c, P) \\ !_m \text{msg}(c_m, c @ V^+) &\sim !_m \text{cell}(c_m, V^+) \\ !_m \text{srv}(c_m, c @ V^-) &\sim !_m \text{cell}(c_m, V^-) \end{aligned}$$

and extend this compositionally to whole configuration. Then we have a strong bisimulation between the two formulations of the operational semantics.

THEOREM 7.3 (BISIMULATION). *Assume  $C \sim \mathcal{D}$ . Then*

- (i) *If  $C \longrightarrow C'$  then  $\mathcal{D} \longrightarrow \mathcal{D}'$  for a  $\mathcal{D}'$  with  $C' \sim \mathcal{D}'$*
- (ii) *If  $\mathcal{D} \longrightarrow \mathcal{D}'$  then  $C \longrightarrow C'$  for some  $C'$  with  $C' \sim \mathcal{D}'$*

PROOF. By cases on the possible reductions, using consistent renaming for freshly created channels/addresses.  $\square$

## 8 RELATED WORK

There have been a number of research threads relating linear logic to computation, too many to survey them all here.

One class of related work uses *linear typing* for a  $\lambda$ -calculus under  $\beta$ -reduction, or, more broadly, a functional language, to capture some intensional properties of terms or to optimize execution (see, for example, [Abramsky 1993; Bernardy et al. 2016; Igarashi and Kobayashi 2000; Mackie 1993]). *Interaction nets* proposed by Lafont [1990] are a native model of computation derived from linear logic that can interpret various other concurrency models, including for example Kahn process networks [Mackie 2016].

More closely related is work on linear types for the  $\pi$ -calculus and concurrent computation [Bellin and Scott 1994; Kobayashi et al. 1996]. This was later recognized to be related to *session types* that prescribe interactions between communicating processes [Honda 1993; Honda et al. 1998], eventually including buffered, asynchronous communication [Carbone et al. 2008; Gay and Vasconcelos 2010]. These lines were unified by giving Curry-Howard interpretations of intuitionistic [Caires and Pfenning 2010; Caires et al. 2016] and classical [Wadler 2012] linear logic in which linear propositions (including the exponential !A) correspond to session types, proofs correspond to programs, and cut reduction corresponds to synchronous communication. This is also closely related to the purely linear development by Cockett and Pastro [2009] who also provide a categorical semantics. It was later observed that we can also give a consistent asynchronous semantics to linear session-typed programs [DeYoung et al. 2012] and, conversely, provide a logical explanation for acknowledgments that implement synchronization [Pfenning and Griffith 2015]. These observations, however, depart from the proofs-as-programs interpretation in that the semantics is no longer directly based on cut reduction. The use of adjunctions to decompose !A in the context of session types is due to Pfenning and Griffith [2015], but the non-linear layer was populated only by  $\uparrow A$  with a more general interpretation given as a challenge. As far as we are aware, none of these works provide a system in which cut reduction corresponds to asynchronous communication, nor do they include multicast or a general non-linear intuitionistic layer. Toninho et al. [2013] integrate linear concurrent programs into an ambient functional language via a contextual monad. While this is both sound and pragmatic, it does not correspond to a uniform logical system or operational semantics as we have presented here.

On the side of parallelism in functional programming, our shared memory semantics appears to be closely related to *futures*, proposed by Halstead [1985] in the dynamically typed setting of

Multilisp. A construct (future  $M$ ) immediately spawns the parallel computation of  $M$  and returns a future that can be synchronized with when the value of  $M$  is needed. A future is modeled here as a shared memory address, accessible to the thread evaluating  $M$  as well as the thread that spawned the future. A significant difference is that sequential computation is the default in Multilisp, while concurrency is the default in LNL<sup>†</sup>. An interesting item of future work is to see if the proposal by Bernardy et al. [2016] to use a double-negation translation to enforce sequentiality might be used in our setting.

## 9 CONCLUSION

We have presented a new variant of a mixed linear/non-linear sequent calculus for Benton's LNL with a notion of cut reduction that corresponds asynchronous communication. We provide two natural operational interpretations that also include arbitrary equirecursive types and recursively defined processes. One semantics is via message passing, in which we can recognize multicast (one send with multiple recipients) in the behavior of shared positive types, and copy-on-recv in the behavior of shared negative types. We do not believe these important constructs were present in prior concurrent languages with logical underpinnings. The second interpretation is via shared memory that exploits the asymmetry of intuitionistic linear logic: right rules always write to memory and left rules always read. Both operational interpretations are in close correspondence and satisfy the expected type preservation and progress properties.

Among the most immediate items of future work, we would like to generalize the calculus to *adjoint logic* [Pruiksma et al. 2018a; Reed 2009] where a lattice of related modes with varying structural properties is available. The semantics might then be more naturally be formalized with subexponentials [Nigam and Miller 2009] or perhaps in focused adjoint logic [Pruiksma et al. 2018b]. We would also like to incorporate sharing in the sense of Balzer and Pfenning [2017]. We conjecture this would support writeable shared memory protected by critical regions with locks. Finally, we would like to consider establishing more control over parallelism using logical means as, for example, proposed by Bernardy et al. [2016].

## ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. CCF-1718267.

## REFERENCES

- Samson Abramsky. 1993. Computational Interpretations of Linear Logic. *Theoretical Computer Science* 111 (1993), 3–57.
- Stephanie Balzer and Frank Pfenning. 2015. Objects as Session-Typed Processes. In *Workshop on Programming based on Actors, Agents, and Decentralized Control (AGERE! 2015)*. ACM SIGPLAN. To appear.
- Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. In *International Conference on Functional Programming (ICFP)*. ACM, 37:1–37:29.
- G. Bellin and P. J. Scott. 1994. One the  $\pi$ -Calculus and Linear Logic. *Theoretical Computer Science* 135 (1994), 11–65.
- Nick Benton. 1994. A Mixed Linear and Non-Linear Logic: Proofs, Terms and Models. In *Selected Papers from the 8th International Workshop on Computer Science Logic (CLS'94)*, Leszek Pacholski and Jerzy Tiuryn (Eds.). Springer LNCS 933, Kazimierz, Poland, 121–135. An extended version appears as Technical Report UCAM-CL-TR-352, University of Cambridge.
- Jean-Philippe Bernardy, Víctor López Juan, and Josef Svenningsson. 2016. Composable Efficient Array Computations Using Linear Types. (Oct. 2016). Unpublished manuscript.
- Luis Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *Proceedings of the 21st International Conference on Concurrency Theory (CONCUR 2010)*. Springer LNCS 6269, Paris, France, 222–236.
- Luis Caires, Frank Pfenning, and Bernardo Toninho. 2016. Linear Logic Propositions as Session Types. *Mathematical Structures in Computer Science* 26, 3 (2016), 367–423. Special Issue on Behavioural Types.

- Marco Carbone, Kohei Honda, and Nobuko Yoshida. 2008. Multiparty Asynchronous Session Types. In *Proceedings of the 35th Symposium on Principles of Programming Languages (POPL '08)*, G.Necula and P.Wadler (Eds.). ACM, San Francisco, California, USA, 273–284.
- Iliano Cervesato, Frank Pfenning, David Walker, and Kevin Watkins. 2002. *A Concurrent Logical Framework II: Examples and Applications*. Technical Report CMU-CS-02-102. Department of Computer Science, Carnegie Mellon University. Revised May 2003.
- Iliano Cervesato and Andre Scedrov. 2009. Relating State-Based and Process-Based Concurrency through Linear Logic. *Information and Computation* 207, 10 (Oct. 2009), 1044–1077.
- J.R.B. Cockett and Craig Pastro. 2009. The Logic of Message-Passing. *Science of Computer Programming* 74 (2009), 498–533.
- Henry DeYoung, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2012. Cut Reduction in Linear Logic as Asynchronous Session-Typed Communication. In *Proceedings of the 21st Conference on Computer Science Logic (CSL 2012)*, P. Cégielski and A. Durand (Eds.). Leibniz International Proceedings in Informatics, Fontainebleau, France, 228–242.
- Simon J. Gay and Malcolm Hole. 2005. Subtyping for Session Types in the  $\pi$ -Calculus. *Acta Informatica* 42, 2–3 (2005), 191–225.
- Simon J. Gay and Vasco T. Vasconcelos. 2010. Linear Type Theory for Asynchronous Session Types. *Journal of Functional Programming* 20, 1 (Jan. 2010), 19–50.
- Robert H. Halstead. 1985. Multilisp: A Language for Parallel Symbolic Computation. *ACM Transactions on Programming Languages and Systems* 7, 4 (Oct. 1985), 501–539.
- Kohei Honda. 1993. Types for Dyadic Interaction. In *4th International Conference on Concurrency Theory (CONCUR'93)*. Springer LNCS 715, 509–523.
- Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *7th European Symposium on Programming Languages and Systems (ESOP'98)*. Springer LNCS 1381, 122–138.
- Atsushi Igarashi and Naoki Kobayashi. 2000. Garbage collection based on a linear type system. In *Preliminary Proceedings of the 3rd ACM SIGPLAN Workshop on Types in Compilation (TIC'00)*, Vol. 152.
- Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. 1996. Linearity and the Pi-Calculus. In *Proceedings of the 23rd Symposium on Principles of Programming Languages (POPL '96)*, H.-J. Boehm and G. Steele (Eds.). ACM, St. Petersburg Beach, Florida, USA, 358–371.
- Y. Lafont. 1990. Interaction Nets. In *17th Symposium on Principles of Programming Languages (POPL '90)*. ACM Press, San Francisco, California, 95–108.
- I. Mackie. 1993. Lilac — A Functional Programming Language Based on Linear Logic. *Journal of Functional Programming* 4, 4 (1993), 395–433.
- Ian Mackie. 2016. Compiling Process Networks to Interaction Nets. In *Computing with Terms and Graphs (TERMGRAPH'16) (EPTCS)*, Vol. 225. 5–14.
- Vivek Nigam and Dale Miller. 2009. Algorithmic Specifications in Linear Logic with Subexponentials. In *Proceedings of the 11th International Conference on Principles and Practice of Declarative Programming (PPDP)*. ACM, Coimbra, Portugal, 129–140.
- Frank Pfenning. 2004. Substructural Operational Semantics and Linear Destination-Passing Style. In *Proceedings of the 2nd Asian Symposium on Programming Languages and Systems (APLAS'04)*, W.-N. Chin (Ed.). Springer-Verlag LNCS 3302, Taipei, Taiwan, 196. Abstract of invited talk.
- Frank Pfenning and Dennis Griffith. 2015. Polarized Substructural Session Types. In *Proceedings of the 18th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2015)*, A. Pitts (Ed.). Springer LNCS 9034, London, England, 3–22. Invited talk.
- Frank Pfenning and Robert J. Simmons. 2009. Substructural Operational Semantics as Ordered Logic Programming. In *Proceedings of the 24th Annual Symposium on Logic in Computer Science (LICS 2009)*. IEEE Computer Society Press, Los Angeles, California, 101–110.
- Klaas Pruikma, William Chargin, Frank Pfenning, and Jason Reed. 2018a. Adjoint Logic. (April 2018). <http://www.cs.cmu.edu/~fp/papers/adjoint18b.pdf> Unpublished manuscript.
- Klaas Pruikma, William Chargin, Frank Pfenning, and Jason Reed. 2018b. Adjoint Logic and Its Concurrent Operational Interpretation. (Jan. 2018). <http://www.cs.cmu.edu/~fp/papers/adjoint18.pdf> Unpublished manuscript.
- Jason Reed. 2009. A Judgmental Deconstruction of Modal Logic. (May 2009). <http://www.cs.cmu.edu/~jcreed/papers/jdml2.pdf> Unpublished manuscript.
- Robert J. Simmons. 2012. *Substructural Logical Specifications*. Ph.D. Dissertation. Carnegie Mellon University. Available as Technical Report CMU-CS-12-142.
- Bernardo Toninho, Luís Caires, and Frank Pfenning. 2013. Higher-Order Processes, Functions, and Sessions: A Monadic Integration. In *Proceedings of the European Symposium on Programming (ESOP'13)*, M.Felleisen and P.Gardner (Eds.). Springer LNCS 7792, Rome, Italy, 350–369.



Philip Wadler. 2012. Propositions as Sessions. In *Proceedings of the 17th International Conference on Functional Programming (ICFP 2012)*. ACM Press, Copenhagen, Denmark, 273–286.