

Back to Futures

Klaas Pruiksmā and Frank Pfenning

Computer Science Department
Carnegie Mellon University
kpruiksm@andrew.cmu.edu and fp@cs.cmu.edu

Abstract. We briefly introduce the semi-axiomatic sequent calculus for linear logic whose natural computational interpretation is session-typed asynchronous communication. This natural asynchrony allows us to endow it with a shared-memory semantics that is weakly bisimilar to its more standard message-passing semantics. We then show how to further refine the concurrent shared memory semantics into a sequential one. Exploiting the expressive framework of adjoint logic, we show how to combine instances of message-passing, shared-memory, and sequential languages into a coherent whole. We exemplify this by providing rational reconstructions for SILL and futures, two approaches for introducing concurrency into functional programming languages. As a byproduct we obtain a first complete definition of typed linear futures.

Keywords: Session types · Linear logic · Futures

1 Introduction

The computational interpretation of constructive proofs depends not only on the logic, but on the fine structure of its presentation. For example, intuitionistic logic may be presented in *axiomatic form* (and we obtain *combinatory reduction* [13]), or via *natural deduction* (which corresponds to the λ -calculus [30]), or a sequent calculus with a stoup (and we discern explicit substitutions [25]).

More recently, a correspondence between linear logic in its sequent formulation and the session-typed synchronous π -calculus has been discovered [3, 10, 7, 54] and analyzed. Cut reduction here corresponds to *synchronous* communication. We can also give an asynchronous message-passing semantics [19] to the same syntax (in the sense that input is blocking but output is not), but this is no longer directly related to cut reduction in the sequent calculus [16]. In this paper we briefly introduce a new style of inference system that combines features from Hilbert’s axiomatic form [26] and Gentzen’s sequent calculus [20] which we call *semi-axiomatic sequent calculus* [15] (Section 2). Cut reduction in the semi-axiomatic sequent calculus corresponds to *asynchronous communication* somewhat like the *asynchronous π -calculus* [6, 28] except that message order must be preserved.

In this paper we show that the apparently small change from the ordinary to the semi-axiomatic sequent calculus has profound and far-reaching consequences.

The first of these is that if we stick to intuitionistic (rather than classical) linear logic, we can easily give a natural (although restricted to write-once memory) shared-memory semantics which is weakly bisimilar with the distributed message-passing semantics (Section 4). In this semantics, channels are reinterpreted as locations in shared memory. Naively, we might expect that sending a message corresponds to writing to memory, while receiving a message reads it. Instead all right rules in the sequent calculus write to memory and all left rules read from memory. This presents the opportunity for a simple *sequential* semantics consistent with the more typical interpretations of intuitionistic logic, but here the role of memory is made explicit (Section 5). The sequential semantics is a refinement of the shared memory semantics that can be seen as a particularly simple stack-based scheduler. This observation, together with the conceptual tools from *adjoint logic* [48, 35, 36, 47] allow us to reconstruct SILL [52, 51, 22], fork/join parallelism [11], and futures [24], three approaches that introduce concurrency into (functional) programming languages (see Sections 7 and 8 and appendix A.4). As a side effect, we obtain a rigorous treatment of *linear futures*, already anticipated and used in parallel complexity analysis by Belloch and Reid-Miller [5] without fully developing them.

For simplicity, we focus mainly on purely linear channels and processes. The ideas are robust enough to easily generalize to the case where weakening and contraction are permitted on the logical side, which leads to memory structures that may be referenced zero (weakening) or multiple (contraction) times (Section 9). Both concurrent and sequential semantics generalize, which means we have captured SILL and a typed version of futures in their usual generality: concurrency primitives layered over a fully expressive functional language.

The principal contributions of this paper are:

1. a shared-memory semantics for session-typed programs based on a computational interpretation of linear logic with a direct proof of bisimilarity to the usual message-passing semantics;
2. a sequential refinement of this semantics;
3. generalizations of these constructions from linear logic to adjoint logic, which supports conservative combinations of multiple modes of truth with varying structural properties among weakening and contraction;
4. a logical reconstruction of SILL which combines functional programming with session-typed message-passing concurrency;
5. a logical reconstruction of typed futures (including a rigorous definition of linear futures) within adjoint logic.

2 A Linear Sequent Calculus for Asynchronous Communication

The Curry-Howard interpretation of linear logic [21, 3, 10, 7, 8, 54] interprets linear propositions as session types [27, 29], proofs in the sequent calculus as processes, and cut reduction as communication. Communication is *synchronous* in

the sense that sender and receiver proceed to their continuations in one joint step. Asynchronous communication, in the sense that the sender may proceed before the message is received, can also be modeled operationally [19] (departing from the correspondence with linear logic) or logically [16]. The latter interprets a message as a process whose only job it is deliver a message. Nevertheless, to be faithful to the logical foundation, an implementation would have to support the default of synchronous communication.

One can ask if there is a different formulation of linear logic that *forces* asynchronous communication because cut reduction itself is asynchronous. An instantiation of this idea can be found in Solos [33, 23] which goes to an extreme in the sense that both send and receive of channels are asynchronous operations, although branching on labels received still has to block. Here we present a practical intermediate point where all send actions are asynchronous but receive actions block.

We call the underlying sequent calculus *semi-axiomatic*. It retains the invertible left or right rules for each connective and replaces the noninvertible ones with axioms, so half the rules will be axioms, while the other half will be the familiar sequent calculus rules. They are presented in Figure 1.

$$\begin{array}{c}
\frac{}{A \vdash A} \text{id} \quad \frac{\Gamma \vdash A \quad \Delta, A \vdash C}{\Gamma, \Delta \vdash C} \text{cut} \\
\\
\frac{}{A_i \vdash A_1 \oplus A_2} \oplus R_i^0 \quad \frac{\Gamma, A_1 \vdash C \quad \Gamma, A_2 \vdash C}{\Gamma, A_1 \oplus A_2 \vdash C} \oplus L \\
\\
\frac{}{A, B \vdash A \otimes B} \otimes R^0 \quad \frac{\Gamma, A, B \vdash C}{\Gamma, A \otimes B \vdash C} \otimes L \\
\\
\frac{}{\cdot \vdash \mathbf{1}} \mathbf{1}R^0 \quad \frac{\Gamma \vdash C}{\Gamma, \mathbf{1} \vdash C} \mathbf{1}L \\
\\
\frac{\Gamma \vdash A_1 \quad \Gamma \vdash A_2}{\Gamma \vdash A_1 \& A_2} \&R \quad \frac{\Gamma, A_i \vdash C}{\Gamma, A_1 \& A_2 \vdash C} \&L_i^0 \\
\\
\frac{\Gamma, A \vdash B}{\Gamma \vdash A \multimap B} \multimap R \quad \frac{}{A, A \multimap B \vdash B} \multimap L^0
\end{array}$$

Fig. 1. Semi-Axiomatic Sequent Calculus, linear fragment

In the presence of cut, it is easy to show that $\Gamma \vdash A$ in the ordinary sequent calculus if and only if $\Gamma \vdash A$ in the semi-axiomatic sequent calculus. In order to state and prove a result that is analogous to cut elimination, we syntactically

classify some proofs as *normal* and then show that for each proof of $\Gamma \vdash A$ there exists one that is normal. Normal forms allow some *analytic* cuts (which preserve the subformula property) and can also be translated compositionally to cut-free proofs in the ordinary sequent calculus. This is analogous to the different notions of normal form for natural deductions [45, 30] and combinatory terms [13], which is where the Curry-Howard isomorphism originated.

In this paper, our interest lies elsewhere, so we do not carry out this proof-theoretic analysis (see [15] for the results on the additive fragment), but we do extract the key steps of cut reduction from the proof of normalization. For this purpose, we label antecedents and succedent with distinct variables so we can better recognize the effects of cut reduction. For internal choice, the reduction selects a branch (\mathcal{E}_2 , in the example) and also substitutes for the antecedent.

$$\frac{\frac{\frac{}{y : A_2 \vdash x : A_1 \oplus A_2} \oplus R_2^0 \quad \frac{\frac{\mathcal{E}_1}{\Gamma, x : A_1 \vdash z : C} \quad \frac{\mathcal{E}_2}{\Gamma, x : A_2 \vdash z : C}}{\Gamma, x : A_1 \oplus A_2 \vdash z : C} \oplus L}{\Gamma, y : A_2 \vdash z : C} \text{cut}_{A_1 \oplus A_2} \quad \frac{[y/x]\mathcal{E}_2}{\Gamma, y : A_2 \vdash z : C}}{\Gamma, y : A_2 \vdash z : C} \longrightarrow$$

For linear implication, we substitute both for antecedent and succedent.

$$\frac{\frac{\frac{\mathcal{D}_0}{\Gamma, y : A \vdash x : B}}{\Gamma \vdash x : A \multimap B} \multimap R \quad \frac{}{u : A, x : A \multimap B \vdash w : B} \multimap L^0}{\Gamma, u : A \vdash w : B} \text{cut}_{A \multimap B} \longrightarrow \frac{[u/y, w/x]\mathcal{D}_0}{\Gamma, u : A \vdash w : B}$$

We see that in both of these cases the cut disappears entirely rather than being reduced to smaller cuts as in the ordinary sequent calculus. In the operational interpretation presented in the next section, axioms correspond to messages, cut reduction corresponds to message receipt, and certain cuts correspond to the (asynchronous) sending of messages.

3 Process Expressions and a Message-Passing Semantics

We now assign process expressions to deductions in the linear semi-axiomatic sequent calculus and formalize their operational semantics. This represents a special case of the semantics of adjoint logic [47] so we omit some details. The typing judgment for processes has the form

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

where P is a process expression providing along channel x and using channels x_i . We have the generic rules of cut and identity.

$$\frac{\Gamma \vdash P :: (x : A) \quad \Delta, x : A \vdash Q :: (z : C)}{\Gamma, \Delta \vdash (x \leftarrow P ; Q) :: (z : C)} \text{cut} \quad \frac{}{y : A \vdash (x \leftarrow y) :: (x : A)} \text{id}$$

We present the dynamic semantics in the form of *multiset rewriting rules* [9]. A *configuration* is a multiset of *semantic objects* $\text{proc}(c, P)$ where c is a channel and P is a process providing c . These semantic objects are *ephemeral* in the sense that if they appear on the left-hand side of a rewriting rule they are consumed and replaced by the right-hand side. Rules can be applied to any subconfiguration, leaving the remainder of the configuration unchanged. All the channels provided in a configuration must be pairwise distinct. By convention, we write the provider of a channel to the left of its client. We then have the following rule for cut, which corresponds to spawning a new process.

$$\text{proc}(c, x \leftarrow P ; Q) \mapsto \text{proc}(a, [a/x]P), \text{proc}(c, [a/x]Q) \quad (a \text{ fresh}) \textit{ cut/spawn}$$

In the reduction rules we use c, d, e for *channels*, which are runtime artifacts, and we continue to use x, y, z for *variables* in programs that stand for channels that are created at runtime. In an object $\text{proc}(c, P)$ the expression P will refer to channels, but will not have any free variables.

In some sense, the most natural semantics for identity $d \leftarrow c$ is a global renaming of c to d or d to c . This is difficult to implement, however, so we instead look for a local version of the semantics. As can be seen from the progress and preservation theorems, it is sufficient for $\text{proc}(d, d \leftarrow c)$ to communicate with the provider of c , and that only when the provider of $\text{proc}(c, P)$ itself tries to communicate along c . We express this as saying that P is *poised on* c . This is also the solution adopted in Concurrent C0 [55].

$$\text{proc}(c, P), \text{proc}(d, d \leftarrow c) \mapsto \text{proc}(d, [d/c]P) \quad (P \text{ poised on } c) \textit{ identity/forward}$$

We introduce additional process expressions in the remainder of this section. A summary of all the rules can be found in Figures 4 to 6. In the purely linear case presented so far, the only mode is $m = L$, so the mode subscripts can be ignored and we have no shifts. Also, Γ_W and Γ_C are always empty. For a summary restricted to the linear fragment, please refer to Appendix A.2.

Internal and External Choice. In order to support programming more naturally, we generalize internal choice and write $\oplus\{\ell : A_\ell\}_{\ell \in L}$ where L is a non-empty finite set of labels. The binary choice can be recovered as $A_1 \oplus A_2 \triangleq \oplus\{\pi_1 : A_1, \pi_2 : A_2\}$. We then obtain the following rules:

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

$$\frac{\Gamma, y : A_\ell \vdash Q_\ell :: (z : C) \quad (\text{for all } \ell \in L)}{\Gamma, x : \oplus\{\ell : A_\ell\}_{\ell \in L} \vdash \mathbf{case} \ x \ (\ell(y) \Rightarrow Q_\ell)_{\ell \in L} :: (z : C)} \oplus L$$

The right rule corresponds to a message carrying the label i and a continuation channel y . The left rule provides a branch for every possible label $\ell \in L$ and also a bound name y for the continuation channel. Interaction between these objects is captured in the following computation rule.

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

This rule represents receipt of the message $i(d)$ along channel c . There is no corresponding rule to *send* such a message, since sending is achieved asynchronously by spawning a process (using cut).

External choice $\&\{\ell : A_\ell\}_{\ell \in L}$ behaves dually. The nullary *left* rule represents a message $x.i(y)$ with label i and continuation channel y , while the *right* rule receives and branches on such a message. These rules can be found in Figure 5 or in Appendix A.2. Computationally, we have

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

The first process above continues, but now providing the continuation channel d that is received with the message.

Multiplicative Unit. Because we use it in our examples, we briefly present the multiplicative unit $\mathbf{1}$. The right rule corresponds to a message $\langle \rangle$ that also closes the channel after it is received, so there is no continuation. The left rule just waits for this message and then proceeds.

$$\frac{}{\cdot \vdash x.\langle \rangle :: (x : \mathbf{1})} \mathbf{1}R^0 \quad \frac{\Gamma \vdash Q :: (z : C)}{\Gamma, x : \mathbf{1} \vdash \text{case } x(\langle \rangle \Rightarrow Q) :: (z : C)} \mathbf{1}L$$

Even though the usual $\mathbf{1}R$ rule already has no premises, we call it here $\mathbf{1}R^0$ to emphasize the principle that in the semi-axiomatic sequent calculus, all noninvertible rules become axioms. Operationally, the unit is straightforward.

$$\text{proc}(c, c.\langle \rangle), \text{proc}(e, \text{case } c(\langle \rangle \Rightarrow Q)) \mapsto \text{proc}(e, Q)$$

Linear Implication and Multiplicative Conjunction. As is true generally for session types, the provider of a linear implication $c : A \multimap B$ receives a channel $e : A$. But because communication is asynchronous, it also must receive a continuation channel $d : B$, so the message has the form $c.\langle e, d \rangle$. The provider has a corresponding $\text{case } c(\langle w, y \rangle \Rightarrow P)$. The actions for the multiplicative conjunction $c : A \otimes B$ are exactly the same, just reversing the role of provider and client. The typing rules are in Figure 6 and the computation rules are included in the ones presented later in this section and in Appendix A.1.

Recursive Types and Processes. So far, we have strictly adhered to the logical foundation provided by the Curry-Howard interpretation of linear logic in its semi-axiomatic formulation. In a programming language, we also need recursive types and recursively defined processes. As is customary in session types, we use *equirecursive types*, collected in a signature Σ in which we also collect recursive process definitions and their types. For each type definition $t = A$, the type A must be *contractive* so we can treat types *equirecursively* with a straightforward coinductive definition and efficient algorithm for type equality [18].

A definition of the process p has the form $x \leftarrow p \leftarrow y_1, \dots, y_n = P$, where the type of p is given as $y_1 : B_1, \dots, y_n : B_n \vdash p :: (x : A)$.

Signature $\Sigma ::= \cdot \mid \Sigma, t = A \mid \Sigma, \overline{y_m} : \overline{B_m} \vdash p :: (x_k : A_k) \mid \Sigma, x_k \leftarrow p \leftarrow \overline{y_m} = P$

For valid signatures we require that each declaration $\overline{y_m : B_m} \vdash p :: (x_k : A_k)$ has a corresponding definition $x_k \leftarrow p \leftarrow \overline{y_m} = P$ with $\overline{y_m : B_m} \vdash P :: (x_k : A_k)$. This means that all type and process definitions can be mutually recursive.

A call is then typed by

$$\frac{y_1 : B_1, \dots, y_n : B_n \vdash p :: (x : A) \in \Sigma}{w_1 : B_1, \dots, w_n : B_n \vdash z \leftarrow p \leftarrow w_1, \dots, w_n :: (z : A)} \text{ call}$$

Operationally, we simply unfold the definition so this rule does not require communication.

$$\text{proc}(c, c \leftarrow p \leftarrow \overline{d}) \mapsto \text{proc}(c, [c/x, \overline{d/y}]P) \quad \text{for } x \leftarrow p \leftarrow \overline{y} = P \in \Sigma \quad (\text{call})$$

In the remainder of this paper we assume that we have a fixed valid signature Σ , so we annotate neither the typing judgment nor the computation rules with an explicit signature.

Example: Binary Numbers. As a first simple example we consider binary numbers, defined as a type *bin*.

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

For example, the number $6 = (110)_2$ would be represented by a sequence of labels $\mathbf{e}, \mathbf{b1}, \mathbf{b1}, \mathbf{b0}$. The first bit to be received would be $\mathbf{b0}$. It arrives along a channel, say c_0 , and carries with it the continuation channel c_1 for the remaining bits. Writing out the whole sequence as a configuration we have

$$\text{proc}(c_4, c_4.\langle \rangle), \text{proc}(c_3, c_3.\mathbf{e}(c_4)), \text{proc}(c_2, c_2.\mathbf{b1}(c_3)), \\ \text{proc}(c_1, c_1.\mathbf{b1}(c_2)), \text{proc}(c_0, c_0.\mathbf{b0}(c_1))$$

Example: Computing with Binary Numbers. We implement a process *succ* which receives the bits of a binary number n along a channel y and produces the bits for the binary number $n + 1$ along x . For a spawn/cut in the examples, we consider a line break like a semicolon. Also, in the examples it is helpful to have a reverse cut, where the order of the premises is reversed.

$$\frac{\Delta, x : A \vdash Q :: (z : C) \quad \Gamma \vdash P :: (x : A)}{\Gamma, \Delta \vdash (x \leftarrow Q ; P) :: (z : C)} \text{ cut}^R$$

Operationally, *cut* and cut^R behave identically. Using the channel names, it is easy to disambiguate which syntactic form of cut is used, so we will point out uses of reverse cut only in this first example. As a general convention in the example processes, we write the continuation of a channel x as x' . The code for the successor process is in Figure 2.

To implement *plus2* we can just compose *succ* with itself. In the message-passing semantics, the two successor processes form a concurrently executing pipeline. This definition is also in Figure 2. We have taken a small syntactic liberty in that the second line of *plus2* should technically be $y \leftarrow (y \leftarrow \text{succ} \leftarrow z)$. We will continue to use the same abbreviation in the remaining examples.

```

bin =  $\oplus\{\mathbf{b0} : bin, \mathbf{b1} : bin, \mathbf{e} : \mathbf{1}\}$ 

y : bin  $\vdash$  succ :: (x : bin)
x  $\leftarrow$  succ  $\leftarrow$  y =
  case y (  $\mathbf{b0}(y') \Rightarrow x' \leftarrow x.\mathbf{b1}(x')$            % send b1 along x (cutR)
          x'  $\leftarrow$  y'                               % forward the remaining bits
        |  $\mathbf{b1}(y') \Rightarrow x' \leftarrow x.\mathbf{b0}(x')$            % send b0 along x (cutR)
          x'  $\leftarrow$  succ  $\leftarrow$  y'                   % carry, as a recursive call
        |  $\mathbf{e}(y') \Rightarrow x' \leftarrow x.\mathbf{b1}(x')$            % send b1 along x (cutR)
          x''  $\leftarrow$  x'. $\mathbf{e}(x'')$                        % followed by e (cutR)
          x''  $\leftarrow$  y' )                               % terminate, by forward at type 1

z : bin  $\vdash$  plus2 :: (x : bin)
z  $\leftarrow$  plus2  $\leftarrow$  x =
  y  $\leftarrow$  succ  $\leftarrow$  z           % spawn first successor in pipeline
  x  $\leftarrow$  succ  $\leftarrow$  y           % continue as second successor

```

Fig. 2. Successor and plus2 processes on binary numbers

Example: A Binary Counter. As a second example, now using external choice, we implement a binary counter for which the client has two choices: it can either increment the counter (message `inc`) or it can request the value in the form of a binary number from the preceding example (message `val`).

$$ctr = \&\{\mathbf{inc} : ctr, \mathbf{val} : bin\}$$

We implement the counter as a chain of processes, each carrying one bit of information, and a final process representing the end of the chain, see Figure 3. An increment message travels through the chain as far as necessary to account for all carries, and the value message travels all the way to the end, converting each process *bit0*, *bit1*, and *end* to its corresponding message `b0`, `b1`, and `e`.

Summary. The semantics presented in this section is a *message-passing semantics* and therefore suitable for distributed computation. It is easy to identify the messages and the threads of control that define the state of each process. A *message* has the form $c.V$ where c is a channel and V is a *value* carried by the message. The values are all small and either of the form $i(d)$ (where i is a label and d a continuation channel), $\langle e, d \rangle$ where e is a channel and d the continuation channel, and $\langle \rangle$ representing the final message on a channel with no continuation. They are uniformly received by a process of the form **case** $c K$ where K is a continuation expression. This is either $(\ell(y) \Rightarrow P_\ell)_{\ell \in L}$ (for internal or external choice), or $\langle \langle w, y \rangle \Rightarrow P \rangle$ (for linear implication or multiplicative conjunction), or $\langle \langle \rangle \Rightarrow P \rangle$ (for multiplicative unit).

A central operation (in both this message-passing semantics and the shared memory semantics of section 4) is to pass a value V to a continuation K , written

```

ctr = &{inc : ctr, val : bin}

y : ctr ⊢ bit0 : ctr
y : ctr ⊢ bit1 : ctr
· ⊢ end : ctr

x ← bit0 ← y =
  case x ( inc(x') ⇒ x' ← bit1 ← y           % continue as bit1
          | val(x') ⇒ x'' ← x'.b0(x'')         % respond with b0 along x'
              y' ← y.val(y')                   % convert remaining bits to binary
              x'' ← y' )                       % and forward

x ← bit1 ← y =
  case x ( inc(x') ⇒ y' ← y.inc(y')           % carry bit as message inc along y
          | val(x') ⇒ x'' ← x'.b1(x'')         % respond with b1 along x'
              y' ← y.val(y')                   % convert remaining bits to binary
              x' ← y' )                       % and forward

x ← end =
  case x ( inc(x') ⇒ y ← end                   % spawn new end process
          | val(x') ⇒ x'' ← x'.e(x'')         % respond with e as empty sequence
              x''.< > )                       % close channel x'' by sending < >

```

Fig. 3. A binary counter

as $V \circ K$ and defined by

$$\begin{aligned}
i(d) \circ (\ell(y) \Rightarrow P_\ell)_{\ell \in L} &= [d/y]P_i && (\oplus, \&) \\
\langle e, c \rangle \circ (\langle w, y \rangle \Rightarrow P) &= [e/w, c/y]P && (\otimes, \multimap) \\
\langle \rangle \circ (\langle \rangle \Rightarrow P) &= P && (\mathbf{1})
\end{aligned}$$

With this notation we can summarize the transition rules as follows:

$$\begin{aligned}
\text{proc}(c, x \leftarrow P ; Q) &\mapsto \text{proc}(a, [a/x]P), \text{proc}(c, [a/x]Q) && (a \text{ fresh}) \quad \text{cut/spawn} \\
\text{proc}(c, P), \text{proc}(d, d \leftarrow c) &\mapsto \text{proc}(d, [d/c]P) && (P \text{ poised on } c) \quad \text{id/forward} \\
\text{proc}(c, c \leftarrow p \leftarrow \bar{d}) &\mapsto \text{proc}(c, [c/x, \bar{d}/y]P) && \text{for } x \leftarrow p \leftarrow \bar{y} = P \in \Sigma \quad \text{call} \\
\text{proc}(c, c.V), \text{proc}(e, \text{case } c K) &\mapsto \text{proc}(e, V \circ K) && (\oplus, \otimes, \mathbf{1}) \\
\text{proc}(c, \text{case } c K), \text{proc}(d, V) &\mapsto \text{proc}(d, V \circ K) && (\&, \multimap)
\end{aligned}$$

Process Configurations. A process configuration \mathcal{C} is either a single process $\text{proc}(c, P)$ providing along c , empty (\cdot) , or the join $\mathcal{C}_1, \mathcal{C}_2$ of two configurations. We think of the join as an associative and commutative operator and impose the condition that all channels c provided in a configuration are distinct. Configurations use some channels and provide others, so we type them with the judgment

$\Gamma \vdash \mathcal{C} :: \Delta$ defined by the following rules.

$$\frac{\Gamma \vdash P :: (c : A)}{\Delta, \Gamma \vdash \mathbf{proc}(c, P) :: (\Delta, c : A)} \quad \frac{}{\Delta \vdash (\cdot) :: \Delta} \quad \frac{\Gamma \vdash \mathcal{C}_1 :: \Delta_1 \quad \Delta_1 \vdash \mathcal{C}_2 :: \Delta_2}{\Gamma \vdash \mathcal{C}_1, \mathcal{C}_2 :: \Delta_2}$$

In the first rule, Δ contains channels not referenced by P which are therefore still available to further clients to the right of $\mathbf{proc}(c, P)$ and we can read the rule for the empty configuration similarly. We then have our first theorems: preservation and progress. The proofs follow standard patterns in the literature [47] and are therefore omitted here. Their origins lie in the correctness of the presented cut reductions in the semi-axiomatic sequent calculus.

Theorem 1 (Type Preservation [47]). *If $\Gamma \vdash \mathcal{C} :: \Delta$ and $\mathcal{C} \mapsto \mathcal{C}'$ then $\Gamma \vdash \mathcal{C}' :: \Delta$.*

A process is *poised on channel c* if it is trying to communicate along c , that is, a process is poised on c if it has the form $\mathbf{proc}(_, c.V)$ (send on c), $\mathbf{proc}(_, \mathbf{case } c K)$ (receive on c), or $\mathbf{proc}(d, d \leftarrow c)$ (forward between c and d).

Theorem 2 (Progress [47]). *If $\cdot \vdash \mathcal{C} :: \Delta$ then either*

- (i) $\mathcal{C} \mapsto \mathcal{C}'$ for some \mathcal{C}' , or
- (ii) every $\mathbf{proc}(c, P) \in \mathcal{C}$ is poised on c .

4 A Shared Memory Semantics

The semantics presented in the previous section is a *message-passing semantics* and therefore suitable for distributed computation. It is easy to identify the messages and the threads of control that define the state of each process. A message has the form $c.V$ and is encapsulated in processes of the form $\mathbf{proc}(_, c.V)$. Messages are small: no process expressions or complicated data structures need to be transmitted. As for threads of control, *cut/spawn* spawns a new process $[a/x]P$ and continues as $[a/x]Q$, *id/forward* terminates, and in the rules for logical connectives the process receiving the message continues while the message terminates. A process $\mathbf{proc}(_, \mathbf{case } c K)$ represents a thread of control with continuation expression K , which proceeds as $V \circ K$ when a value V is received along channel c .

We can now implement this message-passing semantics using a distributed model of computation, or we could implement it using shared memory, depending on our application or programming language. In fact, existing implementations of session types (see, for example, [31, 49, 55, 32]) leverage libraries that may support one or the other or even both. Because of the varying levels of abstractions and the complexity of the linear typing invariant, proving the correctness of such an implementation is difficult and we are not aware of any successful efforts in that direction.

Types	$A_m, B_m ::= \oplus_m \{\ell : A_m^\ell\}_{\ell \in L}$	internal choice
	$A_m \otimes_m B_m$	multiplicative conjunction
	$\mathbf{1}_m$	multiplicative unit
	$\&_m \{\ell : A_m^\ell\}_{\ell \in L}$	external choice
	$A_m \multimap_m B_m$	linear implication
	$\downarrow_m^r A_r$	down shift, $r \geq m$
	$\uparrow_k^m A_k$	up shift, $m \geq k$
	t	type variables
Processes	$P, Q ::= x_m \leftarrow P ; Q$	cut/spawn
	$x_m \leftarrow y_m$	id/forward
	$x_m.V$	send value V
	case $x_m K$	receive and decompose a value
	$x_k \leftarrow p \leftarrow \overline{y_m}$	call
Values	$V ::= i(y_m)$	label i with cont. y_m ($\oplus, \&$)
	$\langle w_m, y_m \rangle$	channel w_m with cont. y_m (\otimes, \multimap)
	$\langle \rangle$	closing message ($\mathbf{1}$)
	shift (y_k)	shift with cont. y_k (\downarrow, \uparrow)
Continuations K	$::= (\ell(y_m) \Rightarrow P)_\ell \in L$	branches for labels $\ell \in L$ with conts. y_m ($\oplus, \&$)
	$(\langle w_m, y_m \rangle \Rightarrow P)$	receiving channel w_m and cont. y_m (\otimes, \multimap)
	$(\langle \rangle \Rightarrow P)$	waiting on closing message ($\mathbf{1}$)
	(shift (y_k) $\Rightarrow P)$	receiving cont. channel y_k (\downarrow, \uparrow)

Fig. 4. ADJ0, types and process expressions. For the purely linear fragment, $m = \mathbf{L}$ and there are no shifts.

Our approach to understanding the implementation of session types is to give a high-level definition of a shared memory semantics. We use the style of *substructural operational semantics* (SSOS) [42] and leverage the technique of *destination-passing style* [44], which also has shown promise for a high-performance implementation [50]. A significant advantage of SSOS is the modularity it affords, which allows us to subsequently integrate the semantics with others (see Section 6).

A key restriction behind the message-passing semantics is that messages are *small*. In particular, they should carry no code in the form of closures—constructing them and controlling their size is a well-known problem in the practice of distributed computing in high-level languages (see, for example, [38]). In the shared memory setting no such restriction is necessary since code can reside in shared memory without difficulty. On the other hand, we would like to explicitly allocate shared memory and ensure that concurrent access is safe.

Here are the central ideas of the shared-memory semantics:

1. Channels are reinterpreted as addresses in shared memory.
2. Cut/spawn is the only way to allocate a new cell in shared memory.
3. Identity/forward will move data between cells.

$$\begin{array}{c}
\frac{}{\Gamma_W, y : A_m \vdash x \leftarrow y :: (x : A_m)} \text{id} \\
\frac{(\Gamma_C, \Delta \geq m \geq r) \quad \Gamma_C, \Delta \vdash P :: (x : A_m) \quad \Gamma_C, \Delta', x : A_m \vdash Q :: (z : C_r)}{\Gamma_C, \Delta, \Delta' \vdash (x \leftarrow P ; Q) :: (z : C_r)} \text{cut} \\
\frac{(i \in L)}{\Gamma_W, y : A_m^i \vdash x.i(y) :: (x : \oplus_m \{\ell : A_m^\ell\}_{\ell \in L})} \oplus R^0 \\
\frac{\Gamma, (x : \oplus_m \{\ell : A_m^\ell\}_{\ell \in L})^\alpha, y : A_m^\ell \vdash Q_\ell :: (z : C_r) \quad (\text{for all } \ell \in L)}{\Gamma, x : \oplus_m \{\ell : A_m^\ell\}_{\ell \in L} \vdash \mathbf{case} x (\ell(y) \Rightarrow Q_\ell)_{\ell \in L} :: (z : C_r)} \oplus L_\alpha \\
\frac{\Gamma \vdash P_\ell :: (y : A_m^\ell) \quad (\text{for all } \ell \in L)}{\Gamma \vdash \mathbf{case} x (\ell(y) \Rightarrow P_\ell)_{\ell \in L} :: (x : \&_m \{\ell : A_m^\ell\}_{\ell \in L})} \& R \\
\frac{(i \in L)}{\Gamma_W, x : \&_m \{\ell : A_m^\ell\}_{\ell \in L} \vdash x.i(y) :: (y : A_m^i)} \& L^0 \\
\frac{}{\cdot \vdash x.\langle \rangle :: (x : \mathbf{1}_m)} \mathbf{1}R^0 \quad \frac{\Gamma, (x : \mathbf{1}_m)^\alpha \vdash P :: (z : C_r)}{\Gamma, x : \mathbf{1}_m \vdash \mathbf{case} x (\langle \rangle \Rightarrow P) :: (z : C_r)} \mathbf{1}L_\alpha \\
\frac{\overline{y_m : B_m} \vdash p :: (x_k : A_k) \in \Sigma}{\overline{w_m : B_m} \vdash z \leftarrow p \leftarrow \overline{w} :: (z_k : A_k)} \text{call}
\end{array}$$

Fig. 5. ADJ0 ($\oplus, \&, \mathbf{1}$). $\alpha \in \{0, 1\}$ with $\alpha = 1$ permitted only if $C \in \sigma(m)$.

$$\begin{array}{c}
\frac{\Gamma, w : A_m \vdash P :: (y : B_m)}{\Gamma \vdash \mathbf{case} x (\langle w, y \rangle \Rightarrow P) :: (x : A_m \multimap_m B_m)} \multimap R \\
\frac{}{\Gamma_W, w : A_m, x : A_m \multimap_m B_m \vdash x.\langle w, y \rangle :: (y : B_m)} \multimap L^0 \\
\frac{}{\Gamma_W, w : A_m, y : B_m \vdash x.\langle w, y \rangle :: (x : A_m \otimes_m B_m)} \otimes R^0 \\
\frac{\Gamma, (x : A_m \otimes_m B_m)^\alpha, w : A_m, y : B_m \vdash P :: (z : C_r)}{\Gamma, x : A_m \otimes_m B_m \vdash \mathbf{case} x (\langle w, y \rangle \Rightarrow P) :: (z : C_r)} \otimes L_\alpha
\end{array}$$

Fig. 6. ADJ0 (\multimap, \otimes). $\alpha \in \{0, 1\}$ with $\alpha = 1$ permitted only if $C \in \sigma(m)$.

4. A process $\text{proc}(c, P)$ that provides c will write to location c and then terminate.
5. A process $\text{proc}(d, Q)$ that is poised on a channel $c \neq d$ will read from location c (once it is available) and then continue.

In addition, memory cells that have been read can be deallocated due to linear typing. The counterintuitive part of this interpretation is that a process providing $c : A \& B$ does not read a value from shared memory. Instead, it writes a continuation to memory and terminates. Conversely, a client of such a channel does not write a value to shared memory. Instead, it continues by jumping to the continuation. A memory cell may therefore contain either a value V or a continuation K , a class of expressions we denote by W .

$$\text{Cell contents} \quad W ::= V \mid K$$

We formalize these principles using the following semantic objects:

1. $\text{thread}(c, P)$: thread P with destination c
2. $\text{cell}(c, _)$: cell c that has been allocated, but not yet written
3. $\text{cell}(c, W)$: cell c containing W

We maintain the invariant that in a configuration either $\text{thread}(c, P)$ appears together with $\text{cell}(c, _)$, or we have just $\text{cell}(c, W)$.

At a lower level of abstraction the continuation would presumably be a closure, pairing a code pointer with an environment assigning addresses to its free channel variables. Unlike messages, values do not carry their destination so we type cell contents with the straightforward judgment $\Gamma \vdash W : A$ whose rules we elide. The typing rules for configurations can be found in Section 9, where the context Γ_C (for antecedents allowing contraction) are always empty for now. It is convenient to type a thread together with its pre-allocated destination. The rules for the operational semantics just formalize these intuitions.

$$\begin{array}{ll}
\text{thread}(c, x \leftarrow P ; Q) \mapsto \text{thread}(a, [a/x]P), \text{cell}(a, _), \text{thread}(c, [a/x]Q) & (a \text{ fresh}) \text{ cut/spawn} \\
\text{cell}(c, W), \text{thread}(d, d \leftarrow c), \text{cell}(d, _) \mapsto \text{cell}(d, W) & \text{id/move} \\
\text{thread}(c, c \leftarrow p \leftarrow \bar{d}) \mapsto \text{thread}(c, [c/x, \bar{d}/y]P) \quad \text{for } x \leftarrow p \leftarrow \bar{y} = P \in \Sigma & \text{call} \\
\text{thread}(c, c.V), \text{cell}(c, _) \mapsto \text{cell}(c, V) & (\oplus R^0, \otimes R^0, \mathbf{1}R^0) \\
\text{cell}(c, V), \text{thread}(e, \text{case } c K) \mapsto \text{thread}(e, V \circ K) & (\oplus L, \otimes L, \mathbf{1}L) \\
\text{thread}(c, \text{case } c K), \text{cell}(c, _) \mapsto \text{cell}(c, K) & (\multimap R, \& R) \\
\text{cell}(c, K), \text{thread}(d, c.V) \mapsto \text{thread}(d, V \circ K) & (\multimap L^0, \& L^0)
\end{array}$$

Note in particular how the last rule reverses the earlier role of messages and processes. The thread computing with destination d continues execution by jumping to the stored continuation after substitution.

This semantics satisfies the expected variants of preservation and progress. Preservation does not change at all.

Theorem 3 (Type Preservation). *If $\Gamma \vdash C :: \Delta$ and $C \mapsto C'$ then $\Gamma \vdash C' :: \Delta$.*

Progress changes in that a configuration that cannot take a step must have filled in all of its destination cells.

Theorem 4 (Progress). *If $\cdot \vdash \mathcal{C} :: \Delta$ then either*

- (i) $\mathcal{C} \mapsto \mathcal{C}'$ for some \mathcal{C}' , or
- (ii) for every channel $c : A \in \Delta$ there is an object $\text{cell}(c, W) \in \mathcal{C}$.

We can now establish a weak bisimulation with the earlier semantics. We define a relation R between configurations in the message-passing and shared-memory semantics.

$$\begin{array}{ll}
\text{proc}(c, x \leftarrow P ; Q) & R \text{ thread}(c, x \leftarrow P ; Q), \text{cell}(c, _) \\
\text{proc}(c, c \leftarrow d) & R \text{ thread}(c, c \leftarrow d), \text{cell}(c, _) \\
\text{proc}(c, c \leftarrow p \leftarrow \bar{d}) & R \text{ thread}(c, c \leftarrow p \leftarrow \bar{d}), \text{cell}(c, _) \\
\text{proc}(c, c.V) & R \text{ thread}(c, c.V), \text{cell}(c, _) \\
\text{proc}(c, c.V) & R \text{ cell}(c, V) \\
\text{proc}(e, \text{case } c K) & R \text{ thread}(e, \text{case } c K), \text{cell}(e, _) \\
\text{proc}(d, c.V) & R \text{ thread}(d, c.V), \text{cell}(d, _) \\
\text{proc}(c, \text{case } c K) & R \text{ thread}(c, \text{case } c K), \text{cell}(c, _) \\
\text{proc}(c, \text{case } c K) & R \text{ cell}(c, K)
\end{array}$$

We extend this relation to whole configurations compositionally, recalling that the join operation for configurations is associative and commutative with unit (\cdot) .

$$\frac{}{(\cdot) R (\cdot)} \quad \frac{\mathcal{C}_1 R \mathcal{D}_1 \quad \mathcal{C}_2 R \mathcal{D}_2}{(\mathcal{C}_1, \mathcal{C}_2) R (\mathcal{D}_1, \mathcal{D}_2)}$$

It is easy to see that R is a weak bisimulation. In the case of cut, we have to take advantage of the stipulation that new names may be chosen arbitrarily as long as they are fresh.

Theorem 5 (Weak Bisimulation). *The relation R is a weak bisimulation between the message-passing and the shared-memory semantics.*

The threads of control in the two semantic specifications are different, however. As we can see from the bisimulation itself, some messages $\text{proc}(_, c.V)$ and $\text{proc}(_, \text{case } c K)$ corresponds to threads and others to memory cells. We return to the examples to illustrate the differences.

Example Revisited: Computing with Binary Numbers, Figure 2. Recall that a *succ* process was a transducer from sequences of bits to sequences bits, all of those being messages. Under the shared memory semantics, *succ* is still the only process, and it reads the elements of the input sequence from memory and writes the elements of the output sequence to memory. In the *plus2* example, under both interpretations there are two *succ* processes in a pipeline operating concurrently, either on messages or on memory content. We see from this that the two interpretations are quite close for *positive* types $(\oplus, \otimes, \mathbf{1})$.

Example Revisited: A Binary Counter, Figure 3. In the message passing semantics for the binary counter we have a sequence of processes, one for each bit and one to mark the end. An increment message *inc* is passed through this sequence until there are no further carries to be processed. In the shared memory semantics each bit is a continuation expression stored in memory. An increment is actually a process that executes such a continuation, writes a new continuation to memory and then reads and executes the next continuation (if necessitated by a carry) or terminates. When receiving a value message, bits *b0* and *b1* are written to memory as each continuation *bit0* and *bit1* is reached and executed.

This alternative operational interpretation is greatly aided by our insistence on an intuitionistic sequent calculus. It is precisely the asymmetry between multiple antecedents and a single conclusion that supports the shared memory interpretation, together with the idea to take the semi-axiomatic sequent calculus as our formal basis.

5 A Sequential Semantics

Once we have introduced threads communicating through shared memory in the particular way we have, it is only a small step to a sequential semantics. The key idea is that for cut/spawn $x \leftarrow P ; Q$ we execute P with destination x and wait until P has terminated (and therefore deposited a value in the location denoted by x). We then proceed with the execution of Q . This would not have been possible in the message-passing semantics because processes pause and resume as messages are exchanged between them. We choose here a call-by-value sequential semantics because *futures* (see Section 8) have first been described in this context. Of course, in the language so far we can code sequential computation by performing additional synchronization before proceeding with the continuation of a cut. As with the coding of asynchronous communication under the synchronous semantics, this would present at best a tortuous path towards an understanding how to combine sequential and concurrent modes of computation since a type by itself would remain agnostic to its intended interpretation.

The idea of the sequential semantics is to split the $\text{thread}(c, P)$ objects into two: $\text{eval}(c, P)$ which evaluates P with destination c and $\text{cont}(c, d, Q)$ which waits for the evaluation of c to complete before continuing with the evaluation of Q with destination d . In addition, we have a special case of the $\text{cell}(c, W)$ object, $\text{retn}(c, W)$, which means that W is returned to destination c , but has not yet been committed to storage. The configuration then always has one of the following two forms, adhering to our convention to show a provider to the left of its client.

$$\begin{aligned} & \text{cell}(d_1, W_1), \dots, \text{eval}(c_0, P_0), \text{cont}(c_0, c_1, P_1), \text{cont}(c_1, c_2, P_2), \dots \\ & \text{cell}(d_1, W_1), \dots, \text{retn}(c_0, W), \text{cont}(c_0, c_1, P_1), \text{cont}(c_1, c_2, P_2), \dots \end{aligned}$$

The continuation objects are threaded in the form of a stack, while the cells represent a global store as before. However, we no longer have uninitialized cells since a cell will always be written by its provider before it is accessed by the client. When typing configurations we have to express the ordered threading

of continuations. We accomplish this by generalizing the typing judgment to $\Gamma \mid \iota \vdash \mathcal{C} :: \Delta \mid \kappa$ where ι, κ are either empty (\cdot) or a single channel c denoting the top of the continuation stack.

$$\frac{\Gamma \vdash P :: (c : A)}{\Delta, \Gamma \mid \cdot \vdash \text{eval}(c, P) :: (\Delta, c : A) \mid c} \quad \frac{\Gamma \vdash P :: (c : A)}{\Delta, \Gamma \mid a \vdash \text{cont}(a, c, P) :: (\Delta, c : A) \mid c}$$

$$\frac{\Gamma \vdash W : A}{\Delta, \Gamma \mid \cdot \vdash \text{retn}(c, W) :: (\Delta, c : A) \mid c} \quad \frac{\Gamma \vdash W : A}{\Delta, \Gamma \mid \cdot \vdash \text{cell}(c, W) :: (\Delta, c : A) \mid \cdot}$$

$$\frac{\Delta \mid \kappa \vdash (\cdot) :: \Delta \mid \kappa}{\Gamma \mid \iota \vdash \mathcal{C}_1 :: \Delta_1 \mid \kappa_1 \quad \Delta_1 \mid \kappa_1 \vdash \mathcal{C}_2 :: \Delta_2 \mid \kappa_2} \quad \Gamma \mid \iota \vdash \mathcal{C}_1, \mathcal{C}_2 :: \Delta_2 \mid \kappa_2$$

The dynamics consists of the following rules.

$$\begin{array}{ll} \text{eval}(c, x \leftarrow P ; Q) \mapsto \text{eval}(a, [a/x]P), \text{cont}(a, c, [a/x]Q) & (a \text{ fresh}) \quad \text{cut/spawn} \\ \text{retn}(a, W), \text{cont}(a, c, Q) \mapsto \text{cell}(a, W), \text{eval}(c, Q) & \text{return/write} \\ \\ \text{cell}(c, W), \text{eval}(d, d \leftarrow c) \mapsto \text{retn}(d, W) & \text{id/move} \\ \text{eval}(c, c \leftarrow p \leftarrow \bar{d}) \mapsto \text{eval}(c, [c/x, \bar{d}/y]P) & \text{for } x \leftarrow p \leftarrow \bar{y} = P \in \Sigma \text{ call} \\ \text{eval}(c, c.V) \mapsto \text{retn}(c, V) & (\oplus R^0, \otimes R^0, \mathbf{1}R^0) \\ \text{cell}(c, V), \text{eval}(e, \text{case } c K) \mapsto \text{eval}(e, V \circ K) & (\oplus L, \otimes L, \mathbf{1}L) \\ \text{eval}(c, \text{case } c K) \mapsto \text{retn}(c, K) & (\multimap R, \& R) \\ \text{cell}(c, K), \text{eval}(d, V) \mapsto \text{eval}(d, V \circ K) & (\multimap L^0, \& L^0) \end{array}$$

We have new but uninteresting versions of preservation and progress.

Theorem 6 (Type Preservation). *If $\Gamma \mid \iota \vdash \mathcal{C} :: \Delta \mid \kappa$ and $\mathcal{C} \mapsto \mathcal{C}'$ then $\Gamma \mid \iota \vdash \mathcal{C}' :: \Delta \mid \kappa$.*

Progress changes in that a configuration that cannot take a step must have filled in all of its destinations either by writing values to memory or returning a value to the continuation on top of the stack.

Theorem 7 (Progress). *If $\cdot \vdash \mathcal{C} :: \Delta \mid \kappa$ then either*

- (i) $\mathcal{C} \mapsto \mathcal{C}'$ for some \mathcal{C}' , or
- (ii) we have
 - (a) for every channel $c : A \in \Delta$ with $c \neq \kappa$ there is a $\text{cell}(c, W) \in \mathcal{C}$, and
 - (b) if $\kappa = c_0$ then $\text{retn}(c_0, W_0) \in \mathcal{C}$ for some W_0 .

The sequential semantics is no longer bisimilar to the concurrent one, but it represents a particular schedule. A simple way of stating this is that the concurrent semantics simulates the sequential one. More might be said by developing a notation of observational equivalence [41], but this is beyond the scope of this paper.

$$\begin{array}{ll} \text{thread}(c, P), \text{cell}(c, _) S \text{eval}(c, P) \\ \text{thread}(c, P), \text{cell}(c, _) S \text{cont}(d, c, P) \\ \text{cell}(c, W) & S \text{cell}(c, W) \\ \text{cell}(c, W) & S \text{retn}(c, W) \end{array}$$

Theorem 8 (Weak Simulation). *The relation S^{-1} is a weak simulation.*

Example Revisited: Computing with Binary Numbers, Figure 2. Under the sequential semantics, a single *succ* process will consume the given input sequence until it encounters the first *b0* (which requires no carry), building up continuation objects that then construct the output. When the two are composed in *plus2*, the first *succ* must finish its computation and write all results to memory before the second *succ* starts.

Example Revisited: A Binary Counter, Figure 3. In the binary counter, an increment represents the only thread of control. It invokes the sequence of continuations in memory in turn until it encounters a *bit0* process, at which point it returns and writes fresh continuations. When multiple increments are executed, each has to finish completely before the next one starts. A value message instead traverses the entire sequence, invoking the stored continuations, and then writes corresponding bits to memory as it returns from each.

6 Adjoint Logic: Combining Languages

At this point we have developed several different versions of an operational semantics for the same source language and explored their relationship. They have different use cases: the message-passing semantics is suitable for distributed computation, the shared-memory semantics for multi-threaded computation, and the sequential semantics for functional programming (linear, for now). One obvious generalization of this is to allow the structural rules of weakening and contraction on the logical side and, correspondingly, persistent cells on the computational side. Another question is how we can combine these languages in a coherent manner, since an application may demand, say, mostly sequential programming with some concurrency.

In this section we review *adjoint logic* [48, 35, 36, 47], a general framework for combining logics with different structural properties. In the next sections we explore particular combinations to reconstruct the essence of SILL [52, 51, 22] and futures [24] as instances of the general schema.

In adjoint logic, propositions are stratified into distinct layers, each identified by a *mode*. For each mode m there is a set $\sigma(m) \subseteq \{W, C\}$ of structural properties satisfied by antecedents of mode m in a sequent. Here, W stands for weakening and C for contraction. For simplicity, we always assume exchange is possible. In addition, any instance of adjoint logic specifies a preorder $m \geq k$ between modes, expressing that the proof of a proposition A_k of mode k may depend on assumptions B_m . In order for cut elimination to hold, this ordering must be compatible with the structural properties: if $m \geq k$ then $\sigma(m) \supseteq \sigma(k)$. Sequents then have the form $\Gamma \vdash A_k$ where, critically, each antecedent B_m in Γ satisfies $m \geq k$. We express this concisely as $\Gamma \geq k$.

A prototypical and inspirational example is Benton's LNL [4] combining linear and nonlinear intuitionistic logic. There are two modes, a linear one L with

$\sigma(\mathbf{L}) = \{\}$ and an unrestricted one \mathbf{U} with $\sigma(\mathbf{U}) = \{W, C\}$. Critically, $\mathbf{U} > \mathbf{L}$, so the proof of a linear proposition can depend on unrestricted assumptions, while a proof of an unrestricted proposition can only depend on other unrestricted propositions but not linear ones.

We can go back and forth between the layers using *shifts* $\uparrow_k^m A_k$ (up from k to m requiring $m \geq k$) and $\downarrow_m^r A_r$ (down from r to m requiring $r \geq m$). A given pair \uparrow_k^m and \downarrow_m^r forms an adjunction, justifying the name adjoint logic.

Propositions $A_m, B_m ::= \oplus_m \{\ell : A_m^\ell\}_{\ell \in L} \mid A_m \multimap_m B_m \mid \dots \mid \uparrow_k^m A_k \mid \downarrow_m^r A_r$

We also sometimes restrict the available propositions of a given mode. For example, we can recover a formulation of intuitionistic linear logic [21] with the two modes \mathbf{L} and \mathbf{U} from LNL with $\mathbf{U} > \mathbf{L}$, but stipulating that $A_u ::= \uparrow_{\mathbf{L}}^u A_{\mathbf{L}}$. Then $!A_{\mathbf{L}} = \downarrow_{\mathbf{L}}^u \uparrow_{\mathbf{L}}^u A_{\mathbf{L}}$.

The formulation of adjoint logic for a fixed set of modes denoting linear, affine, and unrestricted propositions has previously been proposed [43]. More recently, this has been generalized to a uniform message-passing semantics for all of adjoint logic [47] which gives rise to new communication patterns. One example is *multicast*, where a single message is sent to multiple recipients.

We skip a presentation of the pure sequent calculus [46] and just show its semi-axiomatic form in Figures 5, 6, and 7 for the new shift rules. For now, we require all modes to be linear, but remarkably few changes need to be made to handle the general case (see Section 9). When reading these rules it is important to keep in mind the general presupposition that for a sequent $\Gamma \vdash A_k$ we have $\Gamma \geq k$. All other sequents are meaningless and hence disallowed a priori. We therefore always assume the conclusion of the rule satisfies the presupposition and have to add sufficient conditions that it will hold for all premises. In this asynchronous formulation a condition is necessary only for cut: the cut proposition A_m must satisfy $m \geq k$ so that the second premise is well-formed. Furthermore, since m may be strictly greater than k , we must verify that $\Gamma \geq m$, that is, the proof of A_m may depend on all of its antecedents.

$$\frac{}{\Gamma_W, y : A_m \vdash x.\mathbf{shift}(y) :: (x : \downarrow_k^m A_m)} \downarrow R^0 \quad \frac{\Gamma, x : A_m \vdash Q :: (z : C_r)}{\Gamma, x : \downarrow_k^m A_m \vdash \mathbf{case} x (\mathbf{shift}(y) \Rightarrow Q) :: (z :: C_r)} \downarrow L$$

$$\frac{\Gamma \vdash P :: (y : A_k)}{\Gamma \vdash \mathbf{case} x (\mathbf{shift}(y) \Rightarrow P) :: (x : \uparrow_k^m A_k)} \uparrow R \quad \frac{}{\Gamma_W, x : \uparrow_k^m A_k \vdash x.\mathbf{shift}(y) :: (y : A_k)} \uparrow L^0$$

Fig. 7. ADJ0 ($\uparrow\downarrow$).

6.1 A Shared Memory Semantics

We do not provide here a uniform message-passing semantics suitable for a distributed implementation, which has been provided in prior work [47]. Instead, we explore in the remainder of this paper how to combine *different* operational rules at different modes. We see that only the four shift rules and cut deal with different modes. We have to reexamine cut in each case, but the shift rules remain the same for a shared memory semantics because their actions are prescribed by their polarity. Specifically, $\downarrow A$ is positive and therefore sends, while $\uparrow A$ is negative and therefore receives. We annotate each channel with its mode because, unlike in the typing rules, this information is not readily discernible. Here, we imagine a uniform shared-memory semantics for all modes. The language of types with shifts we have already described; processes change only to the extent that the values and continuations in them change. We have a new value V for shifts $\mathbf{shift}(y_k)$ and a new continuation K which receives a shift ($\mathbf{shift}(y_k) \Rightarrow P$) where $\mathbf{shift}(d_k) \circ (\mathbf{shift}(y_k) \Rightarrow P) = [d_k/y_k]P$. These shifts can go in both directions and in this way cover the four different rules.

6.2 Dynamic Interactions Between Layers

In the sequential semantics, we have a different set of semantic objects, but they behave in the same manner as above. Assume the mode k has a sequential semantics and m and r are generic so we represent them as threads.

$$\begin{array}{ll}
\mathbf{eval}(c_k, c_k.\mathbf{shift}(d_m)) \mapsto \mathbf{retn}(c_k, \mathbf{shift}(d_m)) & (\downarrow R^0) \\
\mathbf{cell}(c_k, \mathbf{shift}(d_m)), \mathbf{thread}(e_r, \mathbf{case } c_k (\mathbf{shift}(y_m) \Rightarrow Q)) \mapsto \mathbf{thread}(e_r, [d_m/y_m]Q) & (\downarrow L) \\
\mathbf{thread}(c_m, \mathbf{case } c_m (\mathbf{shift}(y_k) \Rightarrow P)), \mathbf{cell}(c_m, _) \mapsto \mathbf{cell}(c_m, \mathbf{shift}(y_k) \Rightarrow P) & (\uparrow R) \\
\mathbf{cell}(c_m, \mathbf{shift}(y_k) \Rightarrow P), \mathbf{eval}(d_k, c_m.\mathbf{shift}(d_k)) \mapsto \mathbf{eval}(d_k, [d_k/y_k]P) & (\uparrow L^0)
\end{array}$$

The last line (for $\uparrow L^0$) is noteworthy because evaluation, usually entirely sequential, might have to synchronize on the cell c_m if m is a concurrent mode and therefore allocates $\mathbf{cell}(c_m, _)$ in an uninitialized state.

Further interactions take place in a cut. If we assume k is a sequential mode and m is a concurrent mode, then we have the following two cross-mode cuts.

$$\begin{array}{l}
\mathbf{thread}(c_m, x_k \leftarrow P ; Q) \mapsto \mathbf{eval}(a_k, [a/x]P), \mathbf{cont}(a_k, c_m, [a/x]Q) \quad (a_k \text{ fresh}) \\
\mathbf{eval}(c_k, x_m \leftarrow P ; Q) \mapsto \mathbf{thread}(a_m, [a/x]P), \mathbf{cell}(a_m, _), \mathbf{eval}(c_k, [a/x]Q) \quad (a_m \text{ fresh})
\end{array}$$

The invariants we are maintaining are that a sequential cell is always already available before it is referenced (first rule) and that a thread always has an uninitialized destination (second rule). As we can see, the mode of the fresh channel or location determines the behavior of the cut. Note also that we must have $k \geq m$ for the first rule, and $m \geq k$ in the second rule, or the given configurations would not be well-typed.

7 Reconstructing the Essence of SILL

SILL [52, 51, 22] embeds the concurrency primitives of linear logic in a functional language. Briefly, and omitting the presence of *shared channels*, we have

$$\begin{array}{l} \text{Functional types } \tau, \sigma ::= \tau \rightarrow \sigma \mid \tau \times \sigma \mid \mathbf{1} \mid \tau + \sigma \mid \cdots \mid \{S \leftarrow S_1, \dots, S_n\} \\ \text{Session types } S, T ::= S \multimap T \mid S \& T \mid S \otimes T \mid \mathbf{1} \mid S \oplus T \mid \tau \wedge S \mid \tau \supset S \end{array}$$

where the functional layer is pure and sequential. One key construct is the functional type $\{S \leftarrow S_1, \dots, S_n\}$ whose values are process expressions $x \leftarrow \{P\} \leftarrow y_1, \dots, y_n$ such that $y_1 : S_1, \dots, y_n : S_n \vdash P :: (x : S)$. We simplify our task somewhat by discharging the antecedents, so that $\{S \leftarrow S_1, \dots, S_n\}$ becomes $\{S_1 \multimap \cdots \multimap S_n \multimap S\}$, which requires only a simple local transformation of programs. Note that values of this type are closed process expressions that can be constructed and passed around in the functional layer. At the top level we have a concurrent program that can call into the functional layer to construct process expressions that are then concurrently executed using a *bind* construct that corresponds to a cut.

The other two significant constructs are $\tau \wedge A$ and $\tau \supset A$ that send and receive a (functional) value of type τ and continue with A , respectively. Because no restrictions are placed on τ , this general formulation of SILL is more directly suitable for a shared memory semantics than a distributed message-passing semantics, even if it is formally agnostic.

We model this language with two modes, S^* and C such that $S^* > C$ and $\sigma(S^*) = \{W, C\}$ and $\sigma(C) = \{\}$. For now, however, we use the mode S with $\sigma(S) = \{\}$; see Section 9.1 for the general case. We assign the sequential semantics to S and the concurrent (shared-memory) semantics to C . We define a translation $\llbracket - \rrbracket$ which is overloaded for types, functional terms, and process expressions. We show only some exemplars.

$$\begin{array}{l} \llbracket \tau \rightarrow \sigma \rrbracket = \llbracket \tau \rrbracket \multimap_s \llbracket \sigma \rrbracket \\ \llbracket \{S\} \rrbracket = \uparrow_c^s \llbracket S \rrbracket \\ \llbracket S \multimap T \rrbracket = \llbracket S \rrbracket \multimap_c \llbracket T \rrbracket \\ \llbracket \tau \supset S \rrbracket = (\downarrow_c^s \llbracket \tau \rrbracket) \multimap_c \llbracket S \rrbracket \end{array}$$

Functional terms M in SILL are typed with $\Psi \vdash M : \tau$, where Ψ is a context of purely functional variables. We translate M with respect to a destination x such that $\llbracket \Psi \rrbracket_s \vdash \llbracket M \rrbracket_x :: (x_s : \llbracket \tau \rrbracket)$, where $\llbracket \Psi \rrbracket_s$ translates every $y : \sigma$ to $y_s : \llbracket \sigma \rrbracket$. Processes are already in close correspondence and are translated compositionally, where $\Psi ; \Gamma \vdash P :: (x : S)$ entails that $\llbracket \Psi \rrbracket_s, \llbracket \Gamma \rrbracket_c \vdash \llbracket P \rrbracket :: (x_c : \llbracket S \rrbracket)$. First, key cases of the translation of functional terms M .

$$\begin{array}{l} \llbracket w \rrbracket_x = x \leftarrow w \\ \llbracket \lambda w. M \rrbracket_x = \mathbf{case} \ x \langle w, y \rangle \Rightarrow \llbracket M \rrbracket_y \\ \llbracket M N \rrbracket_y = x \leftarrow \llbracket M \rrbracket_x ; w \leftarrow \llbracket N \rrbracket_w ; x.\langle w, y \rangle \end{array}$$

Note the strong analogy with translations to continuation-passing style. For example, the last line could be read as

$$\llbracket M N \rrbracket_y = \llbracket M \rrbracket(\lambda x. \llbracket N \rrbracket(\lambda w. x w y))$$

except that continuations are not themselves λ -terms but named by destinations. For process expressions, the translation is entirely straightforward except for the new constructs such as $\tau \supset S$. Here, we let v stand for the address of a value v of type τ whose translation has type $\downarrow_c^s \llbracket \tau \rrbracket$.

$$\begin{aligned} \llbracket x.\langle v, y \rangle \rrbracket &= w \leftarrow w.\mathbf{shift}(v) ; x.\langle w, y \rangle \\ \llbracket \mathbf{case} x (\langle v, y \rangle \Rightarrow P) \rrbracket &= \mathbf{case} x (\langle w, y \rangle \Rightarrow \mathbf{case} w (\mathbf{shift}(v) \Rightarrow \llbracket P \rrbracket)) \end{aligned}$$

The most interesting part are the functional constructs embedding process expressions. We define:

$$\begin{aligned} \llbracket \{S\} \rrbracket &= \uparrow_c^s \llbracket S \rrbracket \\ \llbracket x \leftarrow \{P\} \rrbracket_y &= \mathbf{case} y_s (\mathbf{shift}(x_c) \Rightarrow \llbracket P \rrbracket) \\ \llbracket x \leftarrow M ; Q \rrbracket &= x_c \leftarrow (y_s \leftarrow \llbracket M \rrbracket_y ; y_s.\mathbf{shift}(x_c)) ; Q \end{aligned}$$

We can see that the operational behavior is correct: the outer cut introducing x_c spawns a new process that first computes $\llbracket M \rrbracket$ sequentially to a value (which represents a suspended process). It then provides access to it along x_c which is used in Q . Together, this is exactly the semantics of $\{ \}E$ in SILL.

The biggest difference to SILL is that the functional layer here is linear. A generalization to a nonlinear one is unproblematic and sketched in Section 9. The choice that $S > C$ makes it straightforward to allow weakening and contraction in the sequential layer. There are a few other features such as polymorphism that are beyond the scope of the present paper.

8 Reconstructing Futures

Futures [24] present the opposite phenomenon from SILL: within functional computation we create a *future* that immediately returns a *promise* and spawns a concurrent computation. *Touching* a promise blocks until its value has been computed. Futures have been a popular mechanism for parallel execution in both statically and dynamically typed languages, and they are also used to encapsulate communication primitives. The simpler mechanism of fork/join parallelism [11] may be a helpful warmup exercise and can be found in Appendix A.4.

For the moment we are interested in *linear futures*, that is, futures that must be used exactly once. In order to preserve their invariants, the sequential language must also be linear. Uses for linear futures (without a full formalization) in the efficient expression of certain parallel algorithms have already been explored in prior work [5].

We present futures in the form of an extension to the purely sequential language and show how it can be elaborated into the adjoint framework in a compositional manner. The first change we need to make is to the preorder between modes: we now postulate $S \geq C$ and also $C \geq S$ because we have mutual dependency between concurrent and sequential computation. We then have the following new type and expressions in the otherwise purely sequential language:

$$\begin{aligned} \text{Types} \quad A &::= \dots \mid \mathbf{fut} A \\ \text{Expressions} \quad P &::= \dots \mid (z \leftarrow \mathbf{future} P ; x.\langle z \rangle) \mid \mathbf{touch} x (\langle y \rangle \Rightarrow P) \end{aligned}$$

We type them as

$$\frac{\Gamma \vdash P :: (z : A)}{\Gamma \vdash (z \leftarrow \mathbf{future} P ; x.\langle z \rangle) :: (x : \mathbf{fut} A)} \mathbf{fut}R$$

$$\frac{\Gamma, z : A \vdash Q :: (w : C)}{\Gamma, x : \mathbf{fut} A \vdash \mathbf{touch} x (\langle z \rangle \Rightarrow Q) :: (w : C)} \mathbf{fut}L$$

The intuition is that $(z \leftarrow \mathbf{future} P ; x.\langle z \rangle)$ will spawn a new thread to evaluate P with destination z , and immediately send the promise of z (written as $\langle z \rangle$) along x . It behaves almost exactly like a cut, except that with the sequential semantics of Section 5 the promise would not be sent until the evaluation of P had terminated. Here, we want it to be evaluated concurrently. The construct $\mathbf{touch} x (\langle z \rangle \Rightarrow Q)$ waits not only for a promise along x (which would be a standard **case** construct), but also for this promise to have been fulfilled by a value. More formally:

$$\begin{aligned} \mathbf{eval}(c, z \leftarrow \mathbf{future} P ; c.\langle z \rangle) &\mapsto \mathbf{eval}(e, [e/z]P), \mathbf{retn}(c, \langle e \rangle) \quad (e \text{ fresh}) \\ \mathbf{eval}(w, \mathbf{touch} c (\langle z \rangle \Rightarrow Q)), \mathbf{cell}(c, \langle e \rangle), \mathbf{retn}(e, W) &\mapsto \mathbf{eval}(w, [e/z]Q), \mathbf{cell}(e, W) \end{aligned}$$

The first clause violates our previous invariant that there is exactly one object **eval** or **retn** in a configuration. We should expect this since with futures we have multiple threads of sequential control, each one represented by an **eval** or **retn**. In the second clause we see that we explicitly synchronize by waiting until the cell e is filled with a value. It would be type-correct to substitute e for z without testing this condition, but then we would not have specified futures.

These constructs can already be expressed compositionally within the framework of adjoint logic. We define:

$$\begin{aligned} \mathbf{fut} A &\rightsquigarrow \downarrow_s^c \uparrow_s^c A_s \\ z \leftarrow \mathbf{future} P ; x.\langle z \rangle &\rightsquigarrow y_c \leftarrow (z_s \leftarrow P ; \mathbf{case} y_c (\mathbf{shift}(u_s) \Rightarrow u_s \leftarrow z_s)) ; x_s.\mathbf{shift}(y_c) \\ \mathbf{touch} x (\langle z \rangle \Rightarrow Q) &\rightsquigarrow \mathbf{case} x_s (\mathbf{shift}(y_c) \Rightarrow (z_s \leftarrow y_c.\mathbf{shift}(z_s) ; Q)) \end{aligned}$$

The typing of the expansions and the arguments for their operational correctness can be found in Appendices A.5 and A.6.

9 Weakening and Contraction

For simplicity of presentation, so far all of our languages and semantics have been entirely linear, even though adjoint logic permits the use of nonlinear modes. Surprisingly few changes are required to add weakening and contraction and thereby obtain a much more expressive language. For example, we will have a fully expressive functional language as a component of SILL as well as ordinary (nonlinear) futures embedded in a (nonlinear) functional language.

Because our primary interest in this paper is a shared-memory semantics, we use a formulation where the structural rules of weakening and contraction are

implicit. For a message-passing semantics and rules with explicit weakening and contraction, see [47].

The only change is the careful management of antecedents. We write Γ_W for antecedents all of whose modes admit weakening (used in the zero-premise rules), and Γ_C for antecedents all admitting contraction (used in cut). In some left rules, if the mode of a channel $x_m : A_m$ admits contraction, we may choose to retain it in the premise. The notation $(x_m : A_m)^\alpha$ expresses this, where $\alpha = 1$ means the assumption is present, and $\alpha = 0$ means that it is not. The typing rules can be found in Figures 5 to 7, now used in their full generality.

We generalize our semantic objects, allowing cells to be persistent if their mode satisfies contraction. A persistent object $!\phi$ on the left-hand side of a rewriting rule is not consumed, but remains in the configuration [9]. We do not explicitly account for weakening, because ephemeral or persistent objects may simply be left over at the end of the computation. We write $!_m\phi$ for a semantic object which is persistent ($!\phi$) if $C \in \sigma(m)$ and otherwise ephemeral (written as just ϕ as we have so far).

1. $\text{thread}(c_m, P)$: thread P with destination c ; always ephemeral
2. $\text{cell}(c_m, _)$: cell c that has been allocated, but not yet written; always ephemeral
3. $!_m\text{cell}(c_m, W)$: cell c containing W , either ephemeral or persistent

$$\frac{\Gamma_C, \Gamma \vdash P :: (c_m : A_m)}{\Gamma_C, \Delta, \Gamma \vdash \text{thread}(c_m, P), \text{cell}(c_m, _) :: (\Gamma_C, \Delta, c_m : A_m)} \quad \frac{}{\Delta \vdash (\cdot) :: \Delta}$$

$$\frac{\Gamma_C, \Gamma \vdash W : A}{\Gamma_C, \Delta, \Gamma \vdash \text{cell}(c_m, W) :: (\Gamma_C, \Delta, c_m : A_m)} \quad \frac{\Gamma \vdash \mathcal{C}_1 :: \Delta_1 \quad \Delta_1 \vdash \mathcal{C}_2 :: \Delta_2}{\Gamma \vdash \mathcal{C}_1, \mathcal{C}_2 :: \Delta_2}$$

The operational semantics is just as before, the only change being additional $!_m$ prefixes for some cells.

$$\begin{aligned} \text{thread}(c_k, x_m \leftarrow P ; Q) &\mapsto \text{thread}(a_m, [a_m/x_m]P), \text{cell}(a_m, _), \text{thread}(c_k, [a_m/x_m]Q) \\ &\quad (a_m \text{ fresh}) \text{ cut/spawn} \\ !_m\text{cell}(c_m, W), \text{thread}(d_m, d_m \leftarrow c_m), \text{cell}(d_m, _) &\mapsto !_m\text{cell}(d_m, W) \text{ id/move} \\ \text{thread}(c_k, c_k \leftarrow p \leftarrow \overline{d_m}) &\mapsto \text{thread}(c_k, [c/x, \overline{d/y}]P) \text{ for } x \leftarrow p \leftarrow \overline{y} = P \in \Sigma \text{ call} \\ \text{thread}(c_m, c_m.V), \text{cell}(c_m, _) &\mapsto !_m\text{cell}(c_m, V) \quad (\oplus R^0, \otimes R^0, \mathbf{1}R^0) \\ !_m\text{cell}(c_m, V), \text{thread}(e_k, \text{case } c_m K) &\mapsto \text{thread}(e_k, V \circ K) \quad (\oplus L, \otimes L, \mathbf{1}L) \\ \text{thread}(c_m, \text{case } c_m K), \text{cell}(c_m, _) &\mapsto !_m\text{cell}(c_m, K) \quad (\multimap R, \&R) \\ !_m\text{cell}(c_m, K), \text{thread}(d_m, c.V) &\mapsto \text{thread}(d_m, V \circ K) \quad (\multimap L^0) \end{aligned}$$

In the revised statement of type preservation, we need to allow additional channels to be provided by a configuration. This could arise in two situations: channels that admit weakening no longer have any clients, or a (fresh) channel admitting contraction should be visible in the interface.

Theorem 9 (Type Preservation). *If $\Gamma \vdash \mathcal{C} :: \Delta$ and $\mathcal{C} \mapsto \mathcal{C}'$ then $\Gamma \vdash \mathcal{C}' :: \Delta, \Delta'$, where all channels in Δ' admit either weakening or contraction.*

The progress theorem does not have to change, but the proof has to accommodate some additional cases.

Theorem 10 (Progress). *If $\cdot \vdash \mathcal{C} :: \Delta$ then either*

- (i) $\mathcal{C} \mapsto \mathcal{C}'$ for some \mathcal{C}' , or
- (ii) for every channel $c_m : A \in \Delta$ there is an object $\text{cell}(c_m, W) \in \mathcal{C}$.

The sequential semantics changes in a completely analogous way, using the semantic objects

1. $\text{eval}(c_m, P)$: evaluate P with destination c_m ; always ephemeral
2. $\text{cont}(c_m, d_k, P)$: cont. waiting on c_m with destination d_k ; always ephemeral
3. $\text{retn}(c_m, W)$: return W to c_m , always ephemeral
4. $!_m \text{cell}(c_m, W)$: cell c_m holding value W ; may be ephemeral or persistent

The results stated in the previous sections continue to hold under these extensions with some small modifications to accommodate persistence.

9.1 SILL Revisited

As already indicated in Section 7, we obtain SILL with the following preorder of modes, where S^* is sequential and subject to weakening and contraction, and C is concurrent and linear, and C^* is concurrent with shared channels but limited to $A_{C^*} ::= \uparrow_C^{C^*} A_C$. The C^* and C layers together constitute intuitionistic linear logic (with the possibility of embedded functional computation in the form of $\downarrow_C^{S^*} A_{S^*}$).

$$S^* > C^* > C \quad \text{with } \sigma(S^*) = \sigma(C^*) = \{W, C\} \text{ and } \sigma(C) = \{\}$$

The distinguishing characteristic of SILL is that we have two layers populated with a full set of connectives: the functional (S^*) and the linear concurrent (C) ones. The functional layer, since it is subject to weakening and contraction as usual, cannot refer to linear channels but it can compose concurrent computations. The linear layer admits a message-passing as well as a shared memory semantics, although without any restrictions a distributed semantics would have to serialize closures.

9.2 Futures Revisited

In the preorder for futures, the top two modes represent (a statically typed version of) traditional futures, where both the underlying functional language with sequential semantics (S^*) and the futures (C^*) are subject to weakening and contraction. The lower two modes represent linear futures C embedded in a linear functional language S , as presented in Section 8. Here, the two concurrent layers (whether subject to weakening and contraction or not) are occupied only by shifts $A_C ::= \uparrow_S^C A_S$ and $A_{C^*} ::= \uparrow_{S^*}^{C^*} A_{S^*}$.

$$\begin{array}{ccc}
S^* & \rightleftarrows & C^* \\
\downarrow & \swarrow & \downarrow \\
S & \rightleftarrows & C
\end{array}
\quad \text{with } \sigma(S^*) = \sigma(C^*) = \{W, C\} \\
\text{and } \sigma(S) = \sigma(C) = \{ \}$$

The distinguishing characteristic of futures is that the concurrent layer is occupied only by a single construct, namely the “future”. In the distributed setting this may be somewhat problematic since serializing closures may be necessary even for relatively simple examples, but it fits perfectly into the shared memory semantics.

Linearity is then an orthogonal consideration, allowing the eager deallocation of futures once they have been touched. However, this is only safely possible if a the surrounding functional language has a clearly delineated linear fragment.

10 Conclusion

We have presented a shared-memory semantics for a rich session-typed language based on a semi-axiomatic sequent calculus for linear logic. We then generalized this to the setting of adjoint logic (augmented with recursion), where weakening and contraction are also allowed, to reconstruct both SILL and futures from first principles. This opens up the possibility of further combinations of semantics, such as message-passing with shared memory, or futures together with SILL, and so on.

There are several interesting items of future work. One is the thorough analysis of the semi-axiomatic sequent calculus from the proof-theoretic perspective, which is ongoing work [15]. Another is a more modular and general understanding of how to combine different operational interpretations of layers in adjoint logic. The semantic stage for this has been set with work on polarization, effects and co-effects including linearity (see, for example, [12]) and call-by-push-value and similarly encompassing frameworks ([34, 39]).

In other ongoing research we are trying to recover the correspondence to logic by restricting recursive types to least and greatest fixed points (which can also be viewed as inductive and coinductive types) [14]. This has to go along with a suitable restriction on recursively defined processes. Some interesting work along these lines has already been carried out starting on the logical side [17, 1, 40]. As far as we are aware, only Toninho et al. [53] and Lindley and Morris [37] explicitly address the question in the context of session types.

Finally, the adjoint framework permits us to formulate a type system with shared resources [2] that can be acquired, interacted with, and then released. This represents a departure from a pure Curry-Howard isomorphism in that computation proceeds not only by proof reduction, but also by proof construction (acquire) and proof destruction (release). For example, it is straightforward to model *mutable* shared memory where memory access is protected by mutexes. We conjecture that this form of sharing dovetails nicely with the approach developed here because of the same logical roots.

Acknowledgments. We would like to thank Henry DeYoung for discussions regarding the subject of this paper and comments on an earlier draft. This material is based upon work supported by the National Science Foundation under Grant CCF-1718267.

References

1. Baelde, D.: Least and greatest fixed points in linear logic. *ACM Transactions on Computational Logic* **13**(1), 2:1–2:44 (2012)
2. Balzer, S., Pfenning, F.: Manifest sharing with session types. In: *International Conference on Functional Programming (ICFP)*. pp. 37:1–37:29. ACM (Sep 2017), extended version available as Technical Report CMU-CS-17-106R, June 2017.
3. Bellin, G., Scott, P.J.: One the π -calculus and linear logic. *Theoretical Computer Science* **135**, 11–65 (1994)
4. Benton, N.: A mixed linear and non-linear logic: Proofs, terms and models. In: Pacholski, L., Tiuryn, J. (eds.) *Selected Papers from the 8th International Workshop on Computer Science Logic (CLS'94)*. pp. 121–135. Springer LNCS 933, Kazimierz, Poland (Sep 1994), an extended version appears as Technical Report UCAM-CL-TR-352, University of Cambridge
5. Bletloch, G.E., Reid-Miller, M.: Pipeling with futures. *Theory of Computing Systems* **32**, 213–239 (1999)
6. Boudol, G.: Asynchrony and the π -calculus. *Rapport de Recherche 1702*, INRIA, Sophia-Antipolis (1992)
7. Caires, L., Pfenning, F.: Session types as intuitionistic linear propositions. In: *Proceedings of the 21st International Conference on Concurrency Theory (CONCUR 2010)*. pp. 222–236. Springer LNCS 6269, Paris, France (Aug 2010)
8. Caires, L., Pfenning, F., Toninho, B.: Linear logic propositions as session types. *Mathematical Structures in Computer Science* **26**(3), 367–423 (2016), special Issue on Behavioural Types.
9. Cervesato, I., Scedrov, A.: Relating state-based and process-based concurrency through linear logic. *Information and Computation* **207**(10), 1044–1077 (Oct 2009)
10. Cockett, J.R.B., Pastro, C.A.: The logic of message-passing. *Science of Computer Programming* **74**(8), 498–533 (2009)
11. Conway, M.E.: A multiprocessor system design. In: *Proceedings of the Fall Joint Computer Conference (AFIPS'63)*. pp. 139–146. ACM, Las Vegas, Nevada, USA (Nov 1963)
12. Curien, P.L., Fiore, M.P., Munch-Maccagnoni, G.: A theory of effects and resources: Adjunction models and polarised calculi. In: Bodík, R., Majumdar, R. (eds.) *Proceedings of the 43rd Symposium on Principles of Programming Languages (POPL 2016)*. pp. 44–56. ACM, St. Petersburg, Florida, USA (Jan 2016)
13. Curry, H.B.: Functionality in combinatory logic. *Proceedings of the National Academy of Sciences, U.S.A.* **20**, 584–590 (1934)
14. Derakhshan, F., Pfenning, F.: Circular proofs as session-typed processes: A local validity condition. *arXiv 1908.01909* (2019)
15. DeYoung, H.: *Session-Typed Ordered Logical Specifications*. Ph.D. thesis, Computer Science Department, Carnegie Mellon University (2019)
16. 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*. pp. 228–242. CSL

- 2012, Leibniz International Proceedings in Informatics, Fontainebleau, France (Sep 2012)
17. Fortier, J., Santocanale, L.: Cuts for circular proofs: Semantics and cut elimination. In: 22nd Conference on Computer Science Logic. LIPIcs, vol. 23, pp. 248–262 (2013)
 18. Gay, S.J., Hole, M.: Subtyping for session types in the π -calculus. *Acta Informatica* **42**(2–3), 191–225 (2005)
 19. Gay, S.J., Vasconcelos, V.T.: Linear type theory for asynchronous session types. *Journal of Functional Programming* **20**(1), 19–50 (Jan 2010)
 20. Gentzen, G.: Untersuchungen über das logische Schließen. *Mathematische Zeitschrift* **39**, 176–210, 405–431 (1935), english translation in M. E. Szabo, editor, *The Collected Papers of Gerhard Gentzen*, pages 68–131, North-Holland, 1969
 21. Girard, J.Y., Lafont, Y.: Linear logic and lazy computation. In: Ehrig, H., Kowalski, R., Levi, G., Montanari, U. (eds.) *Proceedings of the International Joint Conference on Theory and Practice of Software Development*. vol. 2, pp. 52–66. Springer-Verlag LNCS 250, Pisa, Italy (Mar 1987)
 22. Griffith, D.: Polarized Substructural Session Types. Ph.D. thesis, University of Illinois at Urbana-Champaign (Apr 2016)
 23. Guenot, N.: Session types, solos, and the computational contents of the sequent calculus (Dec 2014), talk at the Types Meeting
 24. Halstead, R.H.: Multilisp: A language for parallel symbolic computation. *ACM Transactions on Programming Languages and Systems* **7**(4), 501–539 (Oct 1985)
 25. Herbelin, H.: A lambda-calculus structure isomorphic to Gentzen-style sequent calculus structure. In: Pacholski, L., Tiuryn, J. (eds.) *8th International Workshop on Computer Science Logic*. pp. 61–75. Springer LNCS 933, Kazimierz, Poland (Sep 1994)
 26. Hilbert, D., Bernays, P.: *Grundlagen der Mathematik*. Springer-Verlag, Berlin (1934)
 27. Honda, K.: Types for dyadic interaction. In: *4th International Conference on Concurrency Theory*. pp. 509–523. CONCUR’93, Springer LNCS 715 (1993)
 28. Honda, K., Tokoro, M.: An object calculus for asynchronous communication. In: America, P. (ed.) *Proceedings of the European Conference on Object-Oriented Programming (ECOOP’91)*. pp. 133–147. Springer-Verlag LNCS 512, Geneva, Switzerland (Jul 1991)
 29. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: *7th European Symposium on Programming Languages and Systems*. pp. 122–138. ESOP’98, Springer LNCS 1381 (1998)
 30. Howard, W.A.: The formulae-as-types notion of construction (1969), unpublished note. An annotated version appeared in: *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, 479–490, Academic Press (1980)
 31. Jespersen, T.B.L., Munksgaard, P., Larsen, K.F.: Session types for Rust. In: Bahr, P., Erdweg, S. (eds.) *Proceedings of the 11th Workshop on Generic Programming (WGP 2015)*. pp. 13–22. ACM, Vancouver, Canada (Aug 2015)
 32. Kouzapas, D., Dardha, O., Perera, R., Gay, S.J.: Typechecking protocols with Mungo and StMungo: A session type toolchain for Java. *Science of Computer Programming* **155**, 52–75 (Apr 2018)
 33. Laneve, C., Victor, B.: Solos in concert. *Mathematical Structures in Computer Science* **13**(5), 657–683 (Oct 2003)
 34. Levy, P.B.: Call-by-Push-Value. Ph.D. thesis, University of London (2001)

35. Licata, D.R., Shulman, M.: Adjoint logic with a 2-category of modes. In: International Symposium on Logical Foundations of Computer Science (LFCS). pp. 219–235. Springer LNCS 9537 (Jan 2016)
36. Licata, D.R., Shulman, M., Riley, M.: A fibrational framework for substructural and modal logics. In: Miller, D. (ed.) Proceedings of the 2nd International Conference on Formal Structures for Computation and Deduction (FSCD’17). pp. 25:1–25:22. LIPIcs, Oxford, UK (Sep 2017)
37. Lindley, S., Morris, J.G.: Talking bananas: Structural recursion for session types. In: Garrigue, J., Keller, G., Sumii, E. (eds.) Proceedings of the 21st International Conference on Functional Programming. pp. 434–447. ACM, Nara, Japan (Sep 2016)
38. Miller, H., Haller, P., Müller, N., Boullier, J.: Function passing: A model for typed, distributed functional programming. In: Visser, E., Murphy-Hill, E., Lopes, C. (eds.) International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! 2016). pp. 82–97. ACM (Nov 2016)
39. Munch-Maccagnoni, G.: Syntax and Models of Non-Associative Composition. Ph.D. thesis, Université Paris Diderot – Paris 7 (Dec 2013)
40. Nollet, R., Saurin, A., Tasson, C.: Local validity for circular proofs in linear logic with fixed points. In: Ghica, D., Jung, A. (eds.) 27th Annual Conference on Computer Science Logic (CSL 2018). pp. 35:1–35:23. LIPIcs 119 (2018)
41. 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)
42. Pfenning, F.: Substructural operational semantics and linear destination-passing style. In: Chin, W.N. (ed.) Proceedings of the 2nd Asian Symposium on Programming Languages and Systems (APLAS’04). p. 196. Springer-Verlag LNCS 3302, Taipei, Taiwan (Nov 2004), abstract of invited talk
43. Pfenning, F., Griffith, D.: Polarized substructural session types. In: Pitts, A. (ed.) Proceedings of the 18th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2015). pp. 3–22. Springer LNCS 9034, London, England (Apr 2015), invited talk
44. Pfenning, F., Simmons, R.J.: Substructural operational semantics as ordered logic programming. In: Proceedings of the 24th Annual Symposium on Logic in Computer Science (LICS 2009). pp. 101–110. IEEE Computer Society Press, Los Angeles, California (Aug 2009)
45. Prawitz, D.: Natural Deduction. Almqvist & Wiksell, Stockholm (1965)
46. Pruiksmas, K., Chargin, W., Pfenning, F., Reed, J.: Adjoint logic (Apr 2018), <http://www.cs.cmu.edu/~fp/papers/adjoint18b.pdf>, unpublished manuscript
47. Pruiksmas, K., Pfenning, F.: A message-passing interpretation of adjoint logic. In: Martins, F., Orchard, D. (eds.) Workshop on Programming Language Approaches to Concurrency and Communication-Centric Software (PLACES). pp. 60–79. EPTCS 291, Prague, Czech Republic (Apr 2019)
48. Reed, J.: A judgmental deconstruction of modal logic (May 2009), <http://www.cs.cmu.edu/~jcreed/papers/jdml2.pdf>, unpublished manuscript
49. Scalas, A., Yoshida, N.: Lightweight session programming in scala. In: Proceedings of the 30th European Conference on Object-Oriented Programming (ECOOP 2016). pp. 21:1–21:28. LIPIcs 56, Rome, Italy (Jul 2016)
50. Shaikhha, A., Fitzgibbon, A., Jones, S.P., Vytiniotis, D.: Destination-passing style for efficient memory management. In: Tring, P., Oancea, C. (eds.) Proceedings of the 6th International Workshop on Functional High-Performance Computing (FHPC 2017). pp. 12–23. ACM (Sep 2017)

51. Toninho, B.: A Logical Foundation for Session-based Concurrent Computation. Ph.D. thesis, Carnegie Mellon University and Universidade Nova de Lisboa (May 2015), available as Technical Report CMU-CS-15-109
52. Toninho, B., Caires, L., Pfenning, F.: Higher-order processes, functions, and sessions: A monadic integration. In: M.Felleisen, P.Gardner (eds.) Proceedings of the European Symposium on Programming (ESOP'13). pp. 350–369. Springer LNCS 7792, Rome, Italy (Mar 2013)
53. Toninho, B., Caires, L., Pfenning, F.: Corecursion and non-divergence in session-typed processes. In: Maffei, M., Tuosto, E. (eds.) Proceedings of the 9th International Symposium on Trustworthy Global Computing (TGC 2014). pp. 159–175. Springer LNCS 8902, Rome, Italy (Sep 2014)
54. Wadler, P.: Propositions as sessions. In: Proceedings of the 17th International Conference on Functional Programming. pp. 273–286. ICFP 2012, ACM Press, Copenhagen, Denmark (Sep 2012)
55. Willsey, M., Prabhu, R., Pfenning, F.: Design and implementation of Concurrent C0. In: Fourth International Workshop on Linearity. pp. 73–82. EPTCS 238 (Jun 2016)

A Appendix

A.1 Purely Linear Multiplicatives

We provide a short explanation and example for linear implication in its asynchronous form derived from the semi-axiomatic sequent calculus. Reasoning along similar lines as for internal choice, we assign process expressions to $\multimap R$ and $\multimap L^0$.

$$\frac{\Gamma, u : A \vdash P :: (y : B)}{\Gamma \vdash \mathbf{case} x (\langle u, y \rangle \Rightarrow P) :: (x : A \multimap B)} \multimap R$$

$$\frac{}{w : A, x : A \multimap B \vdash x.\langle w, y \rangle :: (y : B)} \multimap L^0$$

The process for the left rule corresponds to a message $x.\langle w, y \rangle$ that sends channel $w : A$ along x , together with a continuation $y : B$. For the right rule we have used a **case** construct so that all message receipt rules can have a uniform syntax: u stands for the received channel of type A and $y : B$ for the continuation channel. The operational rule is then rather straightforward:

$$\mathbf{proc}(c, \mathbf{case} c (\langle u, y \rangle \Rightarrow P)), \mathbf{proc}(d, c.\langle e, d \rangle) \mapsto \mathbf{proc}(d, [e/u, d/y]P)$$

As an example, consider a proof of $x : A \multimap B, y : B \multimap C \vdash z : A \multimap C$:

$$x : A \multimap B, y : B \multimap C \vdash \mathbf{case} z (\langle u, z' \rangle \Rightarrow y' \leftarrow x.\langle u, y' \rangle ; y.\langle y', z' \rangle) :: (z : A \multimap C)$$

In this example, $u : A$, and the primed variables represent continuation channels $z' : C$, and $y' : B$.

Multiplicative conjunction is again dual, just reversing the role of sender and receiver. Rules can be found in Figure 6 (or Figure 8 for the purely linear form).

A.2 Purely Linear ADJ0

The process assignment and typing rules for the linear semi-axiomatic sequent calculus (plus recursive processes) can be found in Figure 8. Alternatively, we can think of this as adjoint logic with a single mode without structural rules and shifts, augmented with recursion.

A.3 Typing of the Translation of SILL

We verify the typing of the translations from linear adjoint logic to SILL, first for the introduction

$$\frac{\Psi ; \cdot \vdash P :: (x : S)}{\Psi \vdash x \leftarrow \{P\} : \{S\}} \{\}I \rightsquigarrow \frac{\llbracket \Psi \rrbracket_s \vdash \llbracket P \rrbracket :: (x_c : \llbracket S \rrbracket)}{\llbracket \Psi \rrbracket_s \vdash \mathbf{case} y_s (\mathbf{shift}(x_c) \Rightarrow \llbracket P \rrbracket) :: (y_s : \uparrow_c^s \llbracket S \rrbracket)} \uparrow R$$

$$\begin{array}{c}
 \frac{}{y : A \vdash x \leftarrow y :: (x : A)} \text{id} \quad \frac{\Delta \vdash P :: (x : A) \quad \Delta', x : A \vdash Q :: (z : C)}{\Delta, \Delta' \vdash (x \leftarrow P ; Q) :: (z : C)} \text{cut} \\
 \\
 \frac{\overline{y : B} \vdash p :: (x : A) \in \Sigma}{\overline{w : B} \vdash z \leftarrow p \leftarrow \overline{w} :: (z : A)} \text{call} \\
 \\
 \frac{(i \in L)}{y : A^i \vdash x.i(y) :: (x : \oplus\{\ell : A^\ell\}_{\ell \in L})} \oplus R^0 \quad \frac{\Gamma, y : A^\ell \vdash Q_\ell :: (z : C) \quad (\text{for all } \ell \in L)}{\Gamma, x : \oplus\{\ell : A^\ell\}_{\ell \in L} \vdash \mathbf{case} x (\ell(y) \Rightarrow Q_\ell)_{\ell \in L} :: (z : C)} \oplus L \\
 \\
 \frac{}{w : A, y : B \vdash x.\langle w, y \rangle :: (x : A \otimes B)} \otimes R^0 \quad \frac{\Gamma, w : A, y : B \vdash P :: (z : C)}{\Gamma, x : A \otimes B \vdash \mathbf{case} x (\langle w, y \rangle \Rightarrow P) :: (z : C)} \otimes L \\
 \\
 \frac{}{\cdot \vdash x.\langle \rangle :: (x : \mathbf{1})} \mathbf{1}R^0 \quad \frac{\Gamma \vdash P :: (z : C)}{\Gamma, x : \mathbf{1} \vdash \mathbf{case} x (\langle \rangle \Rightarrow P) :: (z : C)} \mathbf{1}L \\
 \\
 \frac{\Gamma \vdash P_\ell :: (y : A^\ell) \quad (\text{for all } \ell \in L)}{\Gamma \vdash \mathbf{case} x (\ell(y) \Rightarrow P_\ell)_{\ell \in L} :: (x : \&\{\ell : A^\ell\}_{\ell \in L})} \&R \quad \frac{(i \in L)}{x : \&\{\ell : A^\ell\}_{\ell \in L} \vdash x.i(y) :: (y : A^i)} \&L^0 \\
 \\
 \frac{\Gamma, w : A \vdash P :: (y : B)}{\Gamma \vdash \mathbf{case} x (\langle w, y \rangle \Rightarrow P) :: (x : A \multimap B)} \multimap R \quad \frac{}{w : A, x : A \multimap B \vdash x.\langle w, y \rangle :: (y : B)} \multimap L^0
 \end{array}$$

Fig. 8. ADJ0 (purely linear fragment)

and then for the elimination

$$\frac{\Psi \vdash M : \{S\} \quad \Psi ; \Gamma, x : S \vdash Q :: (z : T)}{\Psi ; \Gamma \vdash x \leftarrow M ; Q :: (z : T)} \{ \} E \quad \rightsquigarrow \\
 \\
 \frac{\begin{array}{c} \vdots \\ \mathcal{D} \end{array} \quad \frac{[\Psi]_S, [\Gamma]_C, x_C : [S] \vdash Q :: (z_C : [T])}{[\Psi]_S, [\Gamma]_C \vdash (x_C \leftarrow (y_S \leftarrow [M]_y ; y_S.\mathbf{shift}(x_C)) ; [Q]) :: (z_C : [T])} \text{cut}}{\text{cut}}$$

where

$$\mathcal{D} = \frac{[\Psi]_S \vdash [M]_y :: (y_S : \uparrow_c^S [S]) \quad \frac{}{y_S : \uparrow_c^S [S] \vdash y_S.\mathbf{shift}(x_C) :: (x_C : [S])} \uparrow L}{[\Psi]_S \vdash (y_S \leftarrow [M]_y ; y_S.\mathbf{shift}(x_C)) :: (x_C : [S])} \text{cut}}$$

A.4 Reconstructing Fork/Join Parallelism

We use fork/join parallelism as a simple, yet practically highly successful paradigm that provides preliminary insights into futures. In our reconstruction, we can *fork*

a new concurrent computation and use *join* to synchronize on its value. In the context of functional languages we can think of fork/join parallelism as “second-class futures” in that we cannot encapsulate the future as a functional value.

The mode structure for this is quite straightforward: in addition to mode S for sequential computation we have C for concurrent computation, where $S > C$. The sequential layer remains unchanged, and in the concurrent layer we only have $A_C ::= \downarrow_C^S A_S$.

Now assume we have a sequential computation $\Gamma \vdash P :: (x_S : A_S)$ we would like to run concurrently, making its eventual value accessible in Q . We write this as

$$y_C \leftarrow (x_S \leftarrow P ; y_C.\mathbf{shift}(x_S)) ; Q$$

which is typed with the derived rule

$$\frac{\Gamma \vdash P :: (x_S : A_S) \quad \Delta, y_C : \downarrow_C^S A_S \vdash Q :: (w_S : C_S)}{\Gamma, \Delta \vdash (y_C \leftarrow (x_S \leftarrow P ; y_C.\mathbf{shift}(x_S)) ; Q) :: (w_S : C_S)}$$

Because the outer cut is concurrent, this spawns a new thread executing $(x_S \leftarrow P ; d_C.\mathbf{shift}(x_S))$ for a freshly chosen d_C , the eventual location of the result. At the same time it continues with the execution of Q . In the first thread, we have a sequential cut, so we execute P with a new destination e_S , leaving the cell $\mathbf{cell}(d_C, _)$ empty for now and creating a continuation $\mathbf{cont}(e_S, d_C, d_C.\mathbf{shift}(e_S))$ waiting for the value of P to be returned to c_S .

Meanwhile, we may synchronize in the evaluation of Q . We accomplish this with the phrase

$$\mathbf{case} y_C (\mathbf{shift}(x_S) \Rightarrow R)$$

which makes the value of the promise y_C available as x_S in R . It is typed with

$$\frac{\Delta, x_S : A_S \vdash R :: (u_S : D_S)}{\Delta, y_C : \downarrow_C^S A_S \vdash \mathbf{case} y_C (\mathbf{shift}(x_S) \Rightarrow R) :: (u_S : D_S)} \downarrow L$$

In our operational narrative, y_C will have been replaced by the promise d_C when this expression is reached. It will now block until $\mathbf{cell}(d_C, _)$ has been filled.

When the executing P returns a value to c_S the continuation $d_C.\mathbf{shift}(c_S)$ is invoked, which writes the value $\mathbf{shift}(c_S)$ into cell d_C . This unlocks the $\mathbf{case} d_C$ and continues with the evaluation of $[c_S/x_S]R$.

In our formal semantics:

$$\begin{aligned} & \mathbf{eval}(w_S, y_C \leftarrow (x_S \leftarrow P ; y_C.\mathbf{shift}(x_S)) ; Q) \\ \mapsto & \mathbf{thread}(d_C, x_S \leftarrow P ; d_C.\mathbf{shift}(x_S), \mathbf{cell}(d_C, _), \mathbf{eval}(w_S, [d_C/y_C]Q)) \\ \mapsto & \mathbf{eval}(e_S, [e_S/x_S]P), \mathbf{cont}(e_S, d_C, d_C.\mathbf{shift}(e_S)), \mathbf{cell}(d_C, _), \mathbf{eval}(w_S, [d_C/y_C]Q) \end{aligned}$$

We see that continuation for e_S is the bridge between the process P and Q which may decide to wait for its value. When the thread computing Q joins the parallel thread d_C we have the situation

$$\mathbf{eval}(e_S, [e_S/x_S]P), \mathbf{cont}(e_S, d_C, d_C.\mathbf{shift}(e_S)), \mathbf{cell}(d_C, _), \mathbf{eval}(w_S, \mathbf{case} d_C (\mathbf{shift}(x_S) \Rightarrow R))$$

so evaluation of the **case** $d_c (\dots)$ blocks until d_c has been filled. This happens when the computation of P returns a value to e_s .

$$\begin{aligned} & \text{retn}(e_s, W), \text{cont}(e_s, d_c, d_c.\mathbf{shift}(e_s)), \text{cell}(d_c, _), \text{eval}(v_s, \mathbf{case } d_c (\mathbf{shift}(x_s) \Rightarrow R)) \\ \mapsto & \text{cell}(e_s, W), \text{thread}(d_c, d_c.\mathbf{shift}(e_s)), \text{cell}(d_c, _), \text{eval}(v_s, \mathbf{case } d_c (\mathbf{shift}(x_s) \Rightarrow R)) \\ \mapsto & \text{cell}(e_s, W), \text{cell}(d_c, \mathbf{shift}(e_s)), \text{eval}(v_s, \mathbf{case } d_c (\mathbf{shift}(x_s) \Rightarrow R)) \\ \mapsto & \text{cell}(e_s, W), \text{eval}(v_s, [e_s/x_s]R) \end{aligned}$$

A.5 Typing of the Translation of Futures

We show how to type the key components of the translation of futures into ADJ0.

$$\begin{array}{c} \frac{}{z_s : A_s \vdash u_s \leftarrow z_s :: (u_s : A_s)} \text{id} \\ \frac{\Gamma \vdash P :: (z_s : A_s) \quad z_s : A_s \vdash \mathbf{case } y_c (\mathbf{shift}(u_s) \Rightarrow u_s \leftarrow z_s) :: (y_c : \uparrow_s^c A_s)}{\Gamma \vdash z_s \leftarrow P ; \mathbf{case } y_c (\mathbf{shift}(u_s) \Rightarrow u_s \leftarrow z_s) :: (y_c : \uparrow_s^c A_s)} \uparrow R \\ \mathcal{D} = \frac{\vdots \quad \frac{}{y_c : \uparrow_s^c A_s \vdash x_s.\mathbf{shift}(y_c) :: (x : \downarrow_s^c \uparrow_s^c A_s)} \downarrow R^0}{\Gamma \vdash (y_c \leftarrow (z_s \leftarrow P ; \mathbf{case } y_c (\mathbf{shift}(u_s) \Rightarrow u_s \leftarrow z_s)) ; x_s.\mathbf{shift}(y_c)) :: (x_s : \downarrow_s^c \uparrow_s^c A_s)} \text{cut} \end{array}$$

$$\frac{\frac{\frac{}{y_c : \uparrow_s^c A_s \vdash y_c.\mathbf{shift}(z_s) :: (z_s : A_s)} \uparrow L^0 \quad \Gamma, z_s : A_s \vdash Q :: (w : C)}{\Gamma, y_c : \uparrow_s^c A_s \vdash z_s \leftarrow y_c.\mathbf{shift}(z_s) ; Q :: (w : C)} \text{cut}}{\Gamma, x_s : \downarrow_s^c \uparrow_s^c A_s \vdash \mathbf{case } x_s (\mathbf{shift}(y_c) \Rightarrow (z_s \leftarrow y_c.\mathbf{shift}(z_s) ; Q)) :: (w : C)} \downarrow L$$

A.6 Simulation of the Dynamics of Futures

We illustrate how the dynamics of futures is simulated in ADJ0. First, the creation of the future.

$$\begin{aligned} & \text{eval}(c_s, y_c \leftarrow (z_s \leftarrow P ; \mathbf{case } y_c (\mathbf{shift}(u_s) \Rightarrow u_s \leftarrow z_s)) ; c_s.\mathbf{shift}(y_c)) \\ \mapsto & \text{eval}(d_c, z_s \leftarrow P ; \mathbf{case } d_c (\mathbf{shift}(u_s) \Rightarrow u_s \leftarrow z_s), \text{eval}(c_s, c_s.\mathbf{shift}(d_c)) \\ \mapsto^2 & \text{eval}(e_s, P), \text{cont}(e_s, d_c, \mathbf{case } d_c (\mathbf{shift}(u_s) \Rightarrow u_s \leftarrow e_s)), \text{retn}(c_s, \mathbf{shift}(d_c)) \end{aligned}$$

At this point, P executes sequentially with destination e_s and the promise d_c has been returned to c_s . Meanwhile, an expression touching the future c_s executes as follows:

$$\begin{aligned} & \text{eval}(w, \mathbf{case } c_s (\mathbf{shift}(y_c) \Rightarrow (z_s \leftarrow y_c.\mathbf{shift}(z_s) ; Q))), \text{cell}(c_s, \mathbf{shift}(d_c)) \\ \mapsto & \text{eval}(w, z_s \leftarrow d_c.\mathbf{shift}(z_s) ; Q) \\ \mapsto & \text{eval}(f_s, d_c.\mathbf{shift}(f_s)), \text{cont}(f_s, w, [f_s/z_s]Q) \end{aligned}$$

Now $\text{eval}(f_s, d_c.\mathbf{shift}(f_s))$ is blocked until the cell d_c is filled, which takes place when $\text{eval}(e_s, P)$ finishes and invokes its continuation. Note that an eval object can only block if it attempts to read from a cell with concurrent mode, since otherwise the sequential semantics enforces that cells are always filled before they can be referenced. When the value V of the future has been computed, we continue with

$$\begin{aligned}
& \text{eval}(f_s, d_c.\mathbf{shift}(f_s)), \text{cell}(d_c, \mathbf{shift}(u_s) \Rightarrow u_s \leftarrow e_s), \text{cont}(f_s, w, [f_s/z_s]Q), \text{cell}(e_s, V) \\
\mapsto & \text{eval}(f_s, f_s \leftarrow e_s), \text{cont}(f_s, w, [f_s/z_s]Q), \text{cell}(e_s, V) \\
\mapsto & \text{eval}(w, [e_s/f_s][f_s/z_s]Q), \text{cell}(e_s, V) \\
= & \text{eval}(w, [e_s/z_s]Q), \text{cell}(e_s, V)
\end{aligned}$$

So we can see that this is the expected configuration given by the direct semantics.