

The Session Abstract Machine

Luís Caires Bernardo Toninho
Técnico Lisboa / INESC-ID NOVA FCT / NOVA LINCS

We build on a fine-grained analysis of session-based interaction as provided by the linear logic typing disciplines to introduce the SAM, a specialised abstract machine for mechanically executing session-typed processes. A remarkable feature of the SAM’s design is its ability to naturally segregate and coordinate sequential with concurrent session behaviours. In particular, implicitly sequential parts of session programs may be efficiently executed by deterministic sequential application of SAM transitions, amenable to compilation, and dispensing any concurrent synchronisation mechanisms. We provide an intuitive discussion of the SAM structure and its underlying design, and state and prove its correctness for executing programs in a session calculus corresponding to full classical linear logic CLL. We also discuss extensions and applications of the SAM to the execution of linear and session-based programming languages.

1 Introduction

In this work, we build on the linear logic based foundation for session types [13, 15, 69] to construct SAM, an abstract machine specially designed for executing session processes typed by (classical) linear logic CLL. Although motivated by the session type discipline, which originally emerged in the realm of concurrency and distribution [29, 31, 26, 32], a basic motivation for designing the SAM was to provide an efficient deterministic execution model for the implicitly sequential session-typed program idioms that often proliferate in concurrent session-based programming. It is well-known that in a world of fine-grained concurrency, building on many process-based encodings of concepts such as (abstract) data types, functions, continuations, and effects [47, 67, 62, 63, 10, 65, 54], large parts of the code turn out to be inherently sequential, further justifying the foundational and practical relevance of our results. A remarkable feature of the SAM’s design is therefore its potential to efficiently coordinate sequential with full-fledged concurrent behaviours in session-based programming.

Leveraging early work relating linear logic with the semantics of linear and concurrent computation [1, 6, 2], the proposition-as-types (PaT) interpretation [70] of linear logic proofs as a form of well-behaved session-typed nominal calculus has motivated many developments since its inception [12, 5, 65, 64]. We believe that, much how the λ -calculus is deemed a canonical typed model for functional (sequential) computation with pure values, the session calculus can be accepted as a fairly canonical typed model for stateful concurrent computation with linear resources, well-rooted in the trunk of “classical” Type Theory. The PaT interpretation of session processes also establishes a bridge between more classical theories of computation and process algebra via logic. It also reinstates Robin Milner’s view of computation as interaction [46], “data-as-processes” [47] and “functions-as-processes” [45], now in the setting of a tightly

typed world, based on linear logic, where types may statically ensure key properties like deadlock-freedom, termination, and correct resource usage in stateful programs. Session calculi are motivating novel programming language design, bringing up new insights on typeful programming [16] with linear and behavioral types, e.g., [22, 59, 18, 5]. Most systems of typed session calculi have been formulated in process algebraic form [29, 31, 26], or on top of concurrent λ -calculi with an extra layer of communication channels (e.g., [27]), logically inspired systems such as the those discussed in this paper (e.g., [13, 15, 69, 21, 37, 57, 25, 59]) are defined by a logical proof / type system where proof rules are seen as witnesses for the typing of process terms, proofs are read as processes, structural equivalence is proof conversion and computation corresponds to cut reduction. These formulations provide a fundamental semantic foundation to study the model’s expressiveness and meta-theory, but of course do not directly support the concrete implementation of programming languages based on them.

Although several programming language implementations of nominal calculi based languages have been proposed for some time (e.g. [55]), with some introducing abstract machines as the underlying technology (e.g., [66, 44]), we are not aware of any prior design proposal for an abstract machine for reducing session processes exploiting deep properties of a source session calculus, as e.g., the CAM [19] the LAM [39], or the KM [38], which also explore the Curry-Howard correspondences, may reclaim to be, respectively for call-by-value cartesian-closed structures, linear logic, and the call-by-name λ -calculus.

The SAM reduction strategy explores a form of “asynchronous” interaction that essentially expresses that, for processes typed by the logical discipline, sessions are always pairwise causally independent, in the sense that immediate communication on some session is never blocked by communication on a different session. This property is captured syntactically by prefix commutation equations, valid commuting conversions in the underlying logic: adding equations for such laws explicitly to process structural congruence keeps observational equivalence of CLL processes untouched [50]. Combined with insights related to focalisation and polarisation in linear logic [4, 52, 42], we realize that all communication in any session may be operationally structured as the exchange of bundles of positive actions from sender to receiver, where the roles sender/receiver flip whenever the session type swaps polarity. Communication may then be mediated by message buffers, first filled up by the sender (“write-biased” scheduling), and at a later time emptied by the receiver. Building on these observations and on key properties of linear logic proofs leveraged in well-known purely structural proofs of progress [13, 15, 59], we identify a sequential and deterministic reduction strategy for CLL typed processes, based on a form of co-routining where continuations are associated to session queues, and “context switching” occurs whenever polarity flips. That such strategy works at all, preserving all the required correctness properties of the CLL language does not seem immediately obvious, given that each processes may sequentially perform multiple actions on many different sessions, meaning that multiple context switches must be interleaved. The bulk of our paper is then devoted to establishing all such properties in a precise tech-

nical sense. We believe that the SAM may provide a principled foundation for safe execution environments for programming languages combining functional, imperative and concurrent idioms based on session and linear types, as witnessed in practice for Rust [35], (Linear) Haskell [43, 8, 36], Move [9], and in research languages [58, 34, 22]. To further substantiate these views we have developed an implementation of the SAM, integrated in a language for realistic session-based shared-state programs [?].

Outline and Contributions. In Section 2 we briefly review the session-typed calculus CLL, which exactly corresponds to (classical) Linear Logic with mix. In Section 3 we discuss the motivation and design principles of the core SAM, gradually presenting its structure for the language fragment corresponding to session types without the exponentials, which will be introduced later. Even if the core SAM structure and transition rules are fairly simple, the proofs of correctness are more technically involved, and require progressive build up. Therefore, we first bridge between CLL and SAM via an intermediate logical language CLLB, introducing explicit queues in cuts, presented in Section 4. We show preservation (Theorem 4.1) and progress (Theorem 4.2) for CLLB, and prove that there is two way simulation between CLLB and CLL via a strong operational correspondence (Theorem 4.3). Given this correspondence, in Section 5 we state and prove the adequacy of the SAM for executing CLL processes, showing soundness wrt. CLLB (Theorem 5.1) and CLL (Theorem 5.2), and progress / deadlock absence (Theorem 5.3). In Section 6 modularly extend the previous results to the exponentials and mix, and revise the core SAM by introducing explicit environments, stating the associated adequacy results (Theorem 6.1 and Theorem 6.2). We also discuss how to accommodate concurrency, and other extensions in the SAM. We conclude by a discussion of related work and additional remarks. Additional definitions and proofs can be found in the companion technical report [?].

2 Background on CLL, the core language and type system

We start by revisiting the language and type system of CLL, and its operational semantics. The system is based on a PaT interpretation of classical linear logic (we follow the presentations of [15, 11, 58])

Definition 2.1 (Types). *Types A, B are defined by*

$$A, B ::= \mathbf{1} \mid \perp \mid A \wp B \mid A \otimes B \mid \&_{\ell \in L} A_\ell \mid \oplus_{\ell \in L} A_\ell \mid !A \mid ?A$$

Types comprise of the units ($\mathbf{1}, \perp$), multiplicatives (\otimes, \wp), additives ($\oplus_{\ell \in L} A_\ell, \&_{\ell \in L} A_\ell$) and exponentials ($!, ?$). We adopt here a labeled version of the additives, where the linear logic sum type $A_{\#inl} \oplus A_{\#inr}$ is defined by $\oplus_{\ell \in \{\#inl, \#inr\}} A_\ell$. The *positive types* are $\mathbf{1}, \otimes, \oplus$, and $!$, while the *negative types* are $\perp, \wp, \&$ and $?$. We abbreviate $\overline{A} \wp B$ by $A \multimap B$. We write A^+ (resp. A^-) to assert that A is a positive (resp. negative) type. Type *duality* $A \mapsto \overline{A}$ corresponds to negation:

$$\overline{\mathbf{1}} = \perp \quad \overline{A \otimes B} = \overline{A} \wp \overline{B} \quad \overline{\oplus_{\ell \in L} A_\ell} = \&_{\ell \in L} \overline{A}_\ell \quad \overline{!A} = ?\overline{A}$$

Duality captures the symmetry of behaviour in binary process interaction, as manifest in the cut rule. *Typing judgements* have the form $P \vdash \Delta; \Gamma$, where P

$$\begin{array}{c}
\frac{}{0 \vdash \emptyset; \Gamma} \text{[T0]} \quad \frac{P \vdash \Delta'; \Gamma \quad Q \vdash \Delta; \Gamma}{P \parallel Q \vdash \Delta', \Delta; \Gamma} \text{[Tmix]} \\
\frac{}{\text{fwd } x \ y \vdash x : \bar{A}, y : A; \Gamma} \text{[Tfwd]} \quad \frac{P \vdash \Delta', x : A; \Gamma \quad Q \vdash \Delta, x : \bar{A}; \Gamma}{\text{cut } \{P \mid x : A \mid Q\} \vdash \Delta', \Delta; \Gamma} \text{[Tcut]} \\
\frac{}{\text{close } x \vdash x : \mathbf{1}; \Gamma} \text{[T1]} \quad \frac{Q \vdash \Delta; \Gamma}{\text{wait } x; Q \vdash \Delta, x : \perp; \Gamma} \text{[T}\perp\text{]} \\
\frac{P_\ell \vdash \Delta, x : A_\ell; \Gamma \quad (\text{all } \ell \in L)}{\text{case } x \ \{|\#\ell \in L:P_\ell\} \vdash \Delta, x : \&_{\ell \in L} A_\ell; \Gamma} \text{[T}\&\text{]} \quad \frac{Q \vdash \Delta', x : A_{\#\ell}; \Gamma \quad \#\ell \in L}{\#\ell \ x; Q \vdash \Delta', x : \oplus_{\ell \in L} A_\ell; \Gamma} \text{[T}\oplus\text{]} \\
\frac{P_1 \vdash \Delta_1, y : A; \Gamma \quad P_2 \vdash \Delta_2, x : B; \Gamma}{\text{send } x(y.P_1); P_2 \vdash \Delta_1, \Delta_2, x : A \otimes B; \Gamma} \text{[T}\otimes\text{]} \\
\frac{Q \vdash \Delta, z : A, x : B; \Gamma}{\text{recv } x(z); Q \vdash \Delta, x : A \wp B; \Gamma} \text{[T}\wp\text{]} \\
\frac{P \vdash y : A; \Gamma}{!x(y); P \vdash x : !A; \Gamma} \text{[T!]} \quad \frac{Q \vdash \Delta; \Gamma, x : A}{?x; Q \vdash \Delta, x : ?A; \Gamma} \text{[T?]} \\
\frac{P \vdash y : A; \Gamma \quad Q \vdash \Delta; \Gamma, x : \bar{A}}{\text{cut! } \{y.P \mid x : A \mid Q\} \vdash \Delta; \Gamma} \text{[Tcut!]} \quad \frac{Q \vdash \Delta, z : A; \Gamma, x : A}{\text{call } x(z); Q \vdash \Delta; \Gamma, x : A} \text{[Tcall]}
\end{array}$$

Fig. 1: Typing Rules of CLL.

is a process and the *typing context* $\Delta; \Gamma$ is dyadic [4, 7, 51, 13]: both Δ and Γ assign types to names, the context Δ is handled linearly (no implicit contraction or weakening) while the exponential context Γ is unrestricted. The typing rules of CLL are presented in Fig. 1.

The type system exactly corresponds, via a propositions-as-types correspondence, to the canonical proof system of Classical Linear Logic with Mix. The basic process composition rules are [Tmix] and [Tcut], which correspond to the mix and cut rules of Linear Logic. [Tmix] types a parallel composition $P \parallel Q$, where P and Q run in parallel without interference. On the other hand, [Tcut] types interactive composition $\text{cut } \{P \mid x : A \mid Q\}$ where processes P and Q execute concurrently, communicating through a single private linear session channel x . Session endpoints are typed by dual types A/\bar{A} , so x is used by P at type A and by Q at type \bar{A} . When a cut type annotation is easily inferred, we may omit it and write $\text{cut } \{P \mid x \mid Q\}$.

The basic process constructs correspond to the logical introduction rules. The multiplicatives \otimes/\wp type session send / receive actions, the additives $\oplus/\&$ type choice / offer, and the exponentials $!/?$ type replicated session servers and their invocation. We call *action* any process that is either a forwarder or realizes an introduction rule, and denote by \mathcal{A} the set of all actions, by $\mathcal{A}(x)$ the set of action with subject x (the subject of an action is the channel name in which

$\text{fwd } x y \equiv \text{fwd } y x$	[fwd]
$\text{cut } \{P \mid x : A \mid Q\} \equiv \text{cut } \{Q \mid x : \bar{A} \mid P\}$	[com]
$P \parallel 0 \equiv P \quad P \parallel Q \equiv Q \parallel P \quad P \parallel (Q \parallel R) \equiv (P \parallel Q) \parallel R$	[par]
$\text{cut } \{P \mid x \mid (Q \parallel R)\} \equiv (\text{cut } \{P \mid x \mid Q\}) \parallel R$	[CM]
$\text{cut } \{P \mid x \mid (\text{cut } \{Q \mid y \mid R\})\} \equiv \text{cut } \{(\text{cut } \{P \mid x \mid Q\}) \mid y \mid R\}$	[CC]
$\text{cut } \{P \mid z \mid (\text{cut}! \{y.Q \mid !x \mid R\})\} \equiv \text{cut}! \{y.Q \mid !x \mid (\text{cut } \{P \mid z \mid R\})\}$	[CC!]
$\text{cut}! \{y.Q \mid !x \mid (P \parallel R)\} \equiv P \parallel (\text{cut}! \{y.Q \mid !x \mid R\})$	[C!M]
$\text{cut}! \{y.P \mid !x \mid (\text{cut}! \{w.Q \mid !z \mid R\})\} \equiv \text{cut}! \{w.Q \mid !z \mid (\text{cut}! \{y.P \mid !x \mid R\})\}$	[C!C!]
$\text{cut}! \{y.P \mid !x \mid (Q \mid * \mid R)\} \equiv \text{cut}! \{y.P \mid !x \mid Q\} \mid * \mid \text{cut}! \{y.P \mid !x \mid R\}$	[C!*]
$a(x); Q \mid * \mid R \equiv a(x); (Q \mid * \mid R)$	[C+*]
$a_1(x); a_2(y); P \equiv a_2(y); a_1(x); P$	[Ci]

Provisos: in [CM] $x \in \text{fn}(Q)$; in [CC] $x, y \in \text{fn}(Q)$; in [CC!], [C!M] $x \notin \text{fn}(P)$; in [C!C!], $x \notin \text{fn}(Q)$ and $z \notin \text{fn}(P)$. In [Ci], $x \neq y$ and $\text{bn}(a_1(x)) \cap \text{bn}(a_2(y)) = \emptyset$

Fig. 2: Structural congruence $P \equiv Q$.

it interacts [47]). An action is deemed *positive* (resp. *negative*) if its associated type is positive (resp. negative). The set of positive (resp. negative) actions is denoted by \mathcal{A}^+ (resp. \mathcal{A}^-). We sometimes use, e.g., \mathcal{A} or $\mathcal{A}^+(x)$ to denote a process in the set. The CLL operational semantics is given by a *structural congruence* relation \equiv that captures static identities on processes, corresponding to commuting conversions in the logic, and a *reduction* relation \rightarrow that captures process interaction, and corresponds to cut-elimination steps.

Definition 2.2 ($P \equiv Q$). *Structural congruence \equiv is the least congruence on processes closed under α -conversion and the \equiv -rules in Fig. 2.*

The definition of \equiv reflects expected static laws, along the lines of the structural congruences / conversions in [13, 68]. The binary operators forwarder, cut, and mix are commutative. The set of processes modulo \equiv is a commutative monoid with operation the parallel composition ($- \parallel -$) and identity given by inaction 0 ([par]). Any static constructs commute, as expressed by the laws [CM]-[C!sC!]. The unrestricted cut distributes over all the static constructs by law [C*], where $- \mid * \mid -$ stands for either a mix, linear or unrestricted cut. The laws [C+*] and [C+] denote sound proof equivalences in linear logic and bring explicit the independence of linear actions (noted $a(x)$), in different sessions x [50]. These conversions are not required to obtain deadlock freedom. However, they are necessary for full cut elimination (e.g., see [68]), and expose more redexes, thus more non-determinism in the choice of possible reductions. Perhaps surprisingly, this extra flexibility is important to allow the deterministic sequential evaluation strategy for CLL programs adopted by the SAM to be expressed.

$$\begin{array}{ll}
\text{cut } \{\text{fwd } x \ y \ |y| \ P\} \rightarrow \{x/y\}P & [\text{fwd}] \\
\text{cut } \{\text{close } x \ |x| \ \text{wait } x; P\} \rightarrow P & [\mathbf{1}\perp] \\
\text{cut } \{\text{send } x(y.P); Q \ |x| \ \text{recv } x(z); R\} \rightarrow Q \ |x| \ (P \ |y| \ \{y/z\}R) & [\otimes\otimes] \\
\text{cut } \{\text{case } x \ \{\#\ell \in L: P_{\#\ell}\} \ |x| \ \#\ell \ x; R\} \rightarrow \text{cut } \{P_{\#\ell} \ |x| \ R\} & [\&\oplus i] \\
\text{cut } \{!x(y); P \ |x| \ ?x; Q\} \rightarrow \text{cut}! \{y.P \ |!x| \ Q\} & [!?] \\
\text{cut}! \{y.P \ |!x| \ \text{call } x(z); Q\} \rightarrow \text{cut} \{\{z/y\}P \ |z| \ (\text{cut}! \{y.P \ |!x| \ Q\})\} & [\text{call}]
\end{array}$$

Fig. 3: Reduction $P \rightarrow Q$.

Definition 2.3 (Reduction \rightarrow). *Reduction \rightarrow is defined by the rules of Fig. 3.*

We denote by \Rightarrow the reflexive-transitive closure of \rightarrow . Reduction includes the set of principal cut conversions, i.e. the redexes for each pair of interacting constructs. It is closed by structural congruence ($[=]$), in rule [cong] we consider that \mathcal{C} is a static context, i.e. a process context in which the single hole is covered only by the static constructs mix or cut. The forwarding behaviour is implemented by name substitution [fwd] [14]. All the other reductions act on a principal cut between two dual actions, and eliminate it on behalf of cuts involving their subprocesses. CLL satisfies basic safety properties [13] listed below, and also confluence, and termination [58, 59]. In particular we have:

Theorem 2.1 (Type Preservation). *Let $P \vdash \Delta; \Gamma$. (1) If $P \equiv Q$, then $Q \vdash \Delta; \Gamma$. (2) If $P \rightarrow Q$, then $Q \vdash \Delta; \Gamma$.*

A process P is *live* if and only if $P = \mathcal{C}[Q]$, for some static context \mathcal{C} (the hole lies within the scope of static constructs mix and cut) and Q is an active process (a process with a topmost action prefix).

Theorem 2.2 (Progress). *Let $P \vdash \emptyset; \emptyset$ be live. Then $P \rightarrow Q$ for some Q .*

3 A Core Session Abstract Machine

In this section we develop the key insights that guide the construction of our session abstract machine (SAM) and introduce its operational rules in an incremental fashion. We omit the linear logic exponentials for the sake of clarity of presentation, postponing their discussion for Section 6.

One of the main observations that drives the design of the SAM is the nature of proof dynamics in (classical) linear logic, and thus of process execution dynamics in the CLL system of Section 2. The proof dynamics of linear logic are derived from the computational content of the cut elimination proof, which defines a proof simplification strategy that removes (all) instances of the cut rule from a proof. However, the strategy induced by cut elimination is *non-deterministic* insofar as multiple simplification steps may apply to a given proof. Transposing this observation to CLL and other related systems, we observe that their operational semantics does not prescribe a rigid evaluation order for processes. For

S	$::= (P, H)$	State
H	$::= \text{SessionRef} \rightarrow \text{SessionRec}$	Heap
R	$::= x\langle q, P \rangle y$	Session Record
q	$::= \text{nil} \mid \text{Val}@q$	Queue
Val	$::= \checkmark$	Close token
	$\mid \#l$	Choice label
	$\mid \text{clos}(x, P)$	Process Closure

Fig. 4: core SAM Components

instance, in the process $\text{cut } \{P \mid x \mid Q\}$, reduction is allowed in both P and Q . This is of course in line with reduction in process calculi (e.g., [47]). However, in logical-based systems this amounts to *don't care* non-determinism since, regardless of the evaluation order, confluence ensures that the same outcomes are produced (in opposition to *don't know* non-determinism which breaks confluence and is thus disallowed in purely logical systems). The design of the SAM arises from attempting to fix a purely sequential reduction strategy for CLL processes, such that only *one* process is allowed to execute at any given point in time, in the style of coroutines. To construct such a strategy, we forego the use of purely synchronous communication channels, which require a handshake between two concurrently executing processes, and so consider session channels as a kind of *buffered* communication medium (this idea has been explored in the context of linear logic interpretations of sessions in [23]), or queue, where one process can asynchronously write messages so that another may, subsequently, read. To ensure the correct directionality of communication, the queue has a write endpoint (on which a process may only write) and a read endpoint (along which only reads may be performed), such that at any given point in time a process can only hold one of two endpoints of a queue. Moreover, our design takes inspiration from insights related to polarisation and focusing in linear logic, grouping communication in sequences of positive (i.e. write) actions.

Allowing session channels to buffer message sequences, we may then model process execution by alternating between writer processes (that inject messages into the respective queues) and corresponding reader processes. Thus, the SAM must maintain a *heap* that tracks the queue contents of each session (and its endpoints), as well as the suspended processes. The construction of the core of the SAM is given in Figure 4. An execution state is simply a pair consisting of *the* running process P and the heap H . For technical reasons that are made clear in Sections 4 and 5, the process language used in the SAM differs superficially from that of CLL, but for the purposes of this overview we will use CLL process syntax. Later we show the two languages are equivalent in a strong sense.

A heap is a mapping between session identifiers and *session records* of the form $x\langle q, Q \rangle y$, denoting a session with write endpoint x and read endpoint y , with queue contents q and a suspended process Q , holding one of the two endpoints. If Q holds the read endpoint then it is suspended waiting for the process holding

the write endpoint to fill the queue with data for it to read. If Q holds the write endpoint, then Q has been suspended *after* filling the queue and is now waiting for the reader process on y to empty the queue.

We adopt the convention of placing the write endpoint on the left and the read endpoint on the right. In general, session records in the SAM support a form of coroutines through their contained processes, which are called on and returned from multiple times over the course of the execution of the machine. A queue can either be empty (**nil**) or holding a sequence of values. A value is either a close session token (\checkmark), identifying the last output on a session; a choice label $\#l$ or a process closure $\text{clos}(x, P)$, used to model session send and receive. We overload the $@$ notation to also denote concatenation of queues.

Cut. We begin by considering how to execute a cut of the form $\text{cut } \{P \mid x : A^+ \mid Q\}$ where x is a positive type (in the sense of polarized logic [28]) in P . A positive type corresponds to a type denoting an *output* (or write) action, whereas a negative type denotes an *input* (or read) action. We maintain the invariant that in such a cut, P holds the write endpoint and Q the read endpoint. This means that the next action performed by P on the session will be to push some value onto the queue and, dually, the next action performed by Q on the session will be to read a value from the queue. In general, the holder of the write and read endpoint can change throughout execution.

Given the choice of either scheduling P or Q , we are effectively *forced* to schedule P *before* Q . Given that the cut introduces the (unique) session that is shared between the two processes, the only way for Q to exercise its read capability on the session successfully is to wait for P to have exercised (at least some of) its write capability. If we were to schedule Q before P , the process might attempt to read a value from an empty queue, resulting in a stuck state of the SAM. Thus, the SAM execution rule for cut is:

$$(\text{cut } \{P \mid x : A^+ \mid Q\}, H) \Rightarrow (P, H[x\langle \text{nil}, \{y/x\}Q \rangle y]) \quad [\text{SCut}]$$

The rule states that P is the process that is to be scheduled, adding the session record $x\langle \text{nil}, Q \rangle y$ to the heap, which effectively suspends the execution of Q until P has exercised some of its write capabilities on the new session. Note that, in general, both P and Q can interact along many different sessions as both readers and writers before exercising any action on x (resp. y). However, they alone hold the freshly created endpoints x and y and so the next value sent along the session must come from P and Q is its intended receiver.

Channel Output. To execute an output of the form $\text{send } x(z.R); Q$ in the SAM we simply lookup the session record for x and add to the queue a *process closure* containing R (which interacts along z), continuing with the execution of Q :

$$(\text{send } x(z.R); Q, H[x\langle q, P \rangle y]) \Rightarrow (Q, H[x\langle q @ \text{clos}(z, R), P \rangle y]) \quad [\text{S}\otimes]$$

Note that the SAM *eagerly* continues to execute Q instead of switching to P , the holder of the read endpoint of the queue. This allows for the running process to perform all available writes before a context switch occurs.

Session Closure. Executing a `close` follows a similar spirit, but no continuation process exists and so execution switches to the process P holding the *read* endpoint y of the queue:

$$(\text{close } x, H[x\langle q, P \rangle y]) \Rightarrow (P, H[x\langle q @ \checkmark, 0 \rangle y]) \quad [\text{S1}]$$

The process P will eventually read the termination mark from the queue (triggering the deallocation of the session record from the heap):

$$(\text{wait } y; P, H[x\langle \checkmark, 0 \rangle y]) \Rightarrow (P, H) \quad [\text{S}\perp]$$

Note the requirement that \checkmark be the final element of the queue.

Negative Action on Write Endpoint. As hinted above for the case of executing a `cut`, the SAM has a kind of *write bias* insofar as the process chosen to execute in a cut is that which holds the write endpoint for the newly created session. Since CLL processes use channels bidirectionally, the role of writer and reader on a channel (and thus the holder of the write and read endpoints of the queue) may be exchanged during execution. For instance, a process P may wish to send a value v to Q and then receive a response on the same channel. However, when considering a queue-based semantics, the execution of the input action *must not* obtain the value v , intended for Q . Care is therefore needed to ensure that v is received by the holder of the read endpoint of the queue *before* P is allowed to execute its input action (and so taking over the read endpoint). This notion is captured by the following rule, where \mathcal{A}^- denotes any process performing a negative polarity action (i.e., a `wait`, `recv`, `case` or, as we discuss later, a `fwd` x y when x is a write endpoint with a negative polarity type):

$$(\mathcal{A}^-(x), H[x\langle q, Q \rangle y]) \Rightarrow (Q, H[x\langle q, \mathcal{A}^-(x) \rangle y]) \quad [\text{S}-]$$

If the executing process is to perform a negative polarity action on a write endpoint x , the SAM context switches to Q , the holder of the read endpoint y of the session, and suspends the previously running process. This will now allow for Q to perform the appropriate inputs before execution of the action \mathcal{A}^- resumes.

Channel Input. The rules for `recv` actions are as follows:

$$\begin{aligned} (\text{recv } y(w:+); Q, H[x\langle \text{clos}(z, R) @ q, P \rangle y]) &\Rightarrow (Q, H[w\langle \text{nil}, R \rangle z][x\langle q \rangle^s y]) \quad [\text{S}\otimes_+] \\ (\text{recv } y(w:-); Q, H[x\langle \text{clos}(z, R) @ q, P \rangle y]) &\Rightarrow (R, H[z\langle \text{nil}, Q \rangle w][x\langle q \rangle^s y]) \quad [\text{S}\otimes_-] \end{aligned}$$

where $x\langle q \rangle^s y \triangleq$ if $(q = \text{nil})$ then $y\langle q, P \rangle x$ else $x\langle q, P \rangle y$. The execution of an input action requires the corresponding queue to contain a process closure, denoting the process that interacts along the received channel w . In order to ensure that no inputs attempt to read from an empty queue, we must *branch* on the polarity of the communicated session (written $w:+$ and $w:-$ in the rules above): if the session has a *positive* type, then Q must take the *write* endpoint w of the newly generated queue (since Q uses the session with a dual type) and thus we execute Q and allocate a session record in the heap for the new session, with read endpoint z ; if the exchanged session has a *negative* type, the converse holds and

Q must take the *read* endpoint of the newly generated queue. In this scenario, we must execute R so that it may exercise its write capability on the queue and suspend Q in the new session record.

In either case, the session record for the original session is updated by removing the received message from the queue. Crucially, since processes are well-typed, if the resulting queue is empty then it must be the case that Q has no more reads to perform on the session, and so we *swap* the read and write endpoints of the session. This swap serves two purposes: first, it enables Q to perform writes if needed; secondly, and more subtly, it allows for the process, say, P , that holds the other endpoint of the queue to be resumed to perform its actions accordingly. To see how this is the case, consider that such a process will be suspended (due to rule [S−]) attempting to perform a negative action on the write endpoint of the queue. After the swap, the endpoint of the suspended process now matches its intended action. Since Q now holds the write endpoint, it will perform some number of positive actions on the session which end either in a *close*, which context switches to P , or until it attempts to perform a negative action on the write endpoint, triggering rule [S−] and so context switching to P . **Choice and Selection.** The treatment of the additive constructs in the SAM is straightforward:

$$\begin{aligned} (\#l \ x; Q, H[x\langle q, P \rangle y]) &\Rightarrow (Q, H[x\langle q@\#l, P \rangle y]) && \text{[S}\oplus\text{]} \\ (\text{case } y \ \{ \#l \in L: Q_\ell \}, H[x\langle \#l@q, P \rangle y]) &\Rightarrow (Q_{\#l}, H[x\langle q \rangle^s y]) && \text{[S}\&\text{]} \end{aligned}$$

Sending a label $\#l$ simply adds the $\#l$ to the corresponding queue and proceeds with the execution, whereas executing a *case* reads a label from the queue and continues execution of the appropriate branch. Since removing the label may empty the queue, we perform the same adjustment as in rules [S \otimes $_+$] and [S \otimes $_-$]. **Forwarding.** Finally, let us consider the execution of a forwarder (we overload the $@$ notation to also denote concatenation of queues):

$$(\text{fwd } x^- \ y^+, H[z\langle q_1, Q \rangle x][y\langle q_2, P \rangle w]) \Rightarrow (P, H[z\langle q_2@q_1, Q \rangle w]) \quad \text{[Sfwd]}$$

A forwarder denotes the merging of two sessions x and y . Since the forwarder holds the read and write endpoints x and y , respectively, Q has written (through z) the contents of q_1 , whereas the previous steps of the currently running process have written q_2 . Thus, P is waiting to read $q_2@q_1$, justifying the rule above.

The reader may then wonder about other possible configurations of the SAM heap and how they interact with the forwarder. Specifically, what happens if y is of a positive type but a read endpoint of a queue, or, dually, if x is of a negative type but a write endpoint. The former case is *ruled out* by the SAM since the heap satisfies the invariant that any session record of the form $x:A\langle q, P \rangle y:A \in H$ is such that A must be of negative polarity or P is the inert process (which cannot be forwarded). The latter case is possible and is handled by rule [S−], since such a forward $\text{fwd } x^- \ y^+$ stands for a process that wants to perform a *negative polarity* action on a *write* endpoint (or a positive action on a read endpoint).

3.1 On the Write-Bias of the SAM

Consider the following CLL process:

$$\begin{aligned}
P &\triangleq \text{cut } \{P_1 \mid a : \mathbf{1} \otimes \mathbf{1} \mid \{a/b\}Q_1\} \\
P_1 &\triangleq \text{send } a(y.P_2); P_3 & Q_1 &\triangleq \text{recv } b(x); Q_2 \\
P_2 &\triangleq \text{close } y & Q_2 &\triangleq \text{wait } x; Q_3 \\
P_3 &\triangleq \text{close } a & Q_3 &\triangleq \text{wait } b; 0
\end{aligned}$$

Let us walk through the execution trace of P :

$$\begin{array}{ll}
(1) (P, \emptyset) \Rightarrow & \text{by [SCut]} \\
(2) (P_1, a\langle \text{nil}, Q_1 \rangle b) \Rightarrow & \text{by [S}\otimes\text{]} \\
(3) (P_3, a\langle \text{clos}(y, P_2), Q_1 \rangle b) \Rightarrow & \text{by [S1]} \\
(4) (Q_1, a\langle \text{clos}(y, P_2) @ \checkmark, 0 \rangle b) \Rightarrow & \text{by [S}\wp\text{-]} \\
(5) (P_2, y\langle \text{nil}, Q_2 \rangle x, a\langle \checkmark, 0 \rangle b) \Rightarrow & \text{by [S1]} \\
(6) (Q_2, y\langle \checkmark, 0 \rangle x, a\langle \checkmark, 0 \rangle b) \Rightarrow & \text{by [S}\perp\text{]} \\
(7) (Q_3, a\langle \checkmark, 0 \rangle b) \Rightarrow & \text{by [S}\perp\text{]} \\
(8) (0, \emptyset) &
\end{array}$$

The SAM begins in the state on line (1) above, executing the cut. Since the type of a is positive, we execute P_1 , and allocate the session record, suspending Q_1 , resulting in the state on line (2). Since P_1 is a write action on a write endpoint, we proceed via the [S \otimes] rule, resulting in the SAM configuration in line (3), executing P_3 and adding a closure containing P_2 to the session queue with write endpoint a . Executing P_3 (3), a **close** action, requires adding the \checkmark to the queue and context switching to the process Q_1 , now ready to receive the sent value. The applicable rule is now (4) [S \wp -], and so execution will context switch to P_2 after creating the session record for the new session with endpoints y and x . P_2 will execute and the machine ends up in state (6) followed by (7), which consume the appropriate \checkmark and deallocate the session records.

Note how after executing the send action of P_1 we eagerly execute the positive action in P_3 rather than context switching to Q_1 . While in this particular process it would have been safe to execute the negative action in Q_1 , switch to P_2 and then back to Q_2 , we would now need to somehow context switch to P_3 *before* continuing with the execution of Q_3 , or execution would be stuck. However, the relationship between P_3 and Q_2 is unclear at best. Moreover, if the continuation of Q_1 were of the form **wait** b ; **wait** x ; 0, the context switch after the execution of P_2 would have to execute P_3 , or the machine would also be in a stuck state.

3.2 Illustrating Forwarding

To better illustrate the way in which **fwd** $x^- y^+$ effectively stands for a negative action, consider the following CLL process (to simplify the execution trace we assume the existence of output and input of integers typed as $\text{int} \otimes A$ and $\overline{\text{int}} \wp A$, respectively, eliding the need for process closures in this example):

$$P \triangleq \text{cut } \{P_1 \mid b : \overline{\text{int}} \wp \overline{\text{int}} \wp \mathbf{1} \mid \{b/c\} \text{cut } \{Q_1 \mid a : \text{int} \otimes \overline{\text{int}} \wp \mathbf{1} \mid \{a/d\} R_1\}\}$$

$$\begin{array}{lll}
P_1 \triangleq \text{recv } b(x); P_2 & Q_1 \triangleq \text{send } a(1); Q_2 & R_1 \triangleq \text{recv } d(y); R_2 \\
P_2 \triangleq \text{recv } b(z); P_3 & Q_2 \triangleq \text{send } c(3); Q_3 & R_2 \triangleq \text{send } d(2); R_3 \\
P_3 \triangleq \text{close } b & Q_3 \triangleq \text{fwd } a \ c & R_3 \triangleq \text{wait } d; 0
\end{array}$$

If we consider the execution of P we observe:

$$\begin{array}{ll}
(1) (P, \emptyset) \Rightarrow & \text{by [SCut]} \\
(2) (\text{cut } \{Q_1 \mid a\} \{a/d\} R_1, c\langle \text{nil}, P_1 \rangle b) \Rightarrow & \text{by [SCut]} \\
(3) (Q_1, a\langle \text{nil}, R_1 \rangle d, c\langle \text{nil}, P_1 \rangle b) \Rightarrow & \text{by [S}\otimes\text{]} \\
(4) (Q_2, a\langle 1, R_1 \rangle d, c\langle \text{nil}, P_1 \rangle b) \Rightarrow & \text{by [S}\otimes\text{]} \\
(5) (\text{fwd } a \ c, a\langle 1, R_1 \rangle d, c\langle 3, P_1 \rangle b) \Rightarrow & \text{by [S-]} \\
(6) (R_1, a\langle 1, Q_3 \rangle d, c\langle 3, P_1 \rangle b) \Rightarrow & \text{by [S}\otimes\text{]} \\
(7) (R_2, d\langle \text{nil}, Q_3 \rangle a, c\langle 3, P_1 \rangle b) \Rightarrow & \text{by [S}\otimes\text{]} \\
(8) (R_3, d\langle 2, Q_3 \rangle a, c\langle 3, P_1 \rangle b) \Rightarrow & \text{by [S-]} \\
(9) (\text{fwd } a \ c, d\langle 2, R_3 \rangle a, c\langle 3, P_1 \rangle b) \Rightarrow & \text{by [Sfwd]} \\
(10) (P_1, d\langle 3\otimes 2, R_3 \rangle b) \Rightarrow & \text{by [S}\otimes\text{]} \\
(11) (P_2, b\langle 2, R_3 \rangle d) \Rightarrow & \text{by [S}\otimes\text{]} \\
(12) (P_3, b\langle \text{nil}, R_3 \rangle d) \Rightarrow & \text{by [S1]} \\
(13) (R_3, b\langle \checkmark, R_3 \rangle d) \Rightarrow & \text{by [S}\perp\text{]} \\
(14) (0, \emptyset) &
\end{array}$$

The first four steps of the execution of P allocate the two session records and the writes by Q_1 and Q_2 takes place. We are now in configuration (5), where $Q_3 = \text{fwd } a^- \ c^+$ is to execute and a is a write endpoint of a queue assigned a negative type ($\overline{\text{int}} \otimes \mathbf{1}$). This forwarder stands for a process performing a negative action on a write endpoint (i.e., P_1) and so context switching is required, rule [S-] applies and the SAM context switches to R_1 , suspending Q_3 until the forward can be performed. After R_1 receives (6) and the queue endpoints a and d are swapped (7), R_2 executes and then rule [S-] applies (8), context switching back to Q_3 . Since the queue endpoints are now flipped, rule [Sfwd] now applies (9), collapsing the two session records (via queue concatenation) and proceeding with the execution of P_1 , P_2 , P_3 and R_3 (10-14). Note the correct ordering in which the sent values are dequeued, where 3 is read before 2, as intended.

Discussion. The core execution rules for the SAM are summarized in Figure 5. At this point, the reader may wonder just how reasonable the SAM's evaluation strategy is. Our evaluation strategy is devised to be a deterministic, sequential strategy, where exactly one process is executing at any given point in time, supported by a queue-based buffer structure for channels and a heap for session records. Moreover, taking inspiration from focusing and polarized logic, we adopt a write-biased stance and prioritize (bundles of) write actions over reads, where suspended processes hold the read endpoint of queues while waiting for the writer process to fill the queue, and hold write endpoints of queues *after* filling them, waiting for the reader process to empty the queue.

While this latter point seems like a reasonable way to ensure that inputs never get stuck, it is not immediately obvious that the strategy is sound wrt the more standard (asynchronous) semantics of CLL and related languages, given that processes are free to act on multiple sessions. Thus, the write-bias of the cut rule (and the overall SAM) does not necessarily mean that the process that

$(\text{cut } \{P \mid x : A^+ \mid Q\}, H) \Rightarrow (P, H[x\langle \text{nil}, \{y/x\}Q \rangle y])$	[SCut]
$(\text{fwd } x \ y, H[z\langle q_1, Q \rangle x][y\langle q_2, P \rangle w]) \Rightarrow (P, H[z\langle q_2 @ q_1, Q \rangle w])$	[Sfwd]
$(\text{close } x, H[x\langle q, P \rangle y]) \Rightarrow (P, H[x\langle q @ \surd, 0 \rangle y])$	[S1]
$(\text{wait } y; P, H[x\langle \surd, 0 \rangle y]) \Rightarrow (P, H)$	[S \perp]
$(\mathcal{A}^-(x), H[x\langle q, Q \rangle y]) \Rightarrow (Q, H[x\langle q, \mathcal{A}^-(x) \rangle y])$	[S-]
$(\text{send } x(z.R); Q, H[x\langle q, P \rangle y]) \Rightarrow (Q, H[x\langle q @ \text{clos}(z, R), P \rangle y])$	[S \otimes]
$(\text{recv } y(w : +); Q, H[x\langle \text{clos}(z, R) @ q, P \rangle y]) \Rightarrow (Q, H[w\langle \text{nil}, R \rangle z][x\langle q \rangle^s y])$	[S \wp]
$(\text{recv } y(w : -); Q, H[x\langle \text{clos}(z, R) @ q, P \rangle y]) \Rightarrow (R, H[z\langle \text{nil}, Q \rangle w][x\langle q \rangle^s y])$	[S \wp]
$(\#l \ x; Q, H[x\langle q, P \rangle y]) \Rightarrow (Q, H[x\langle q @ \#l, P \rangle y])$	[S \oplus]
$(\text{case } y \ \{\#\ell \in L:Q_\ell\}, H[x\langle \#l @ q, P \rangle y]) \Rightarrow (Q_{\#l}, H[x\langle q \rangle^s y])$	[S $\&$]
<i>N.B.</i> : $x\langle q \rangle^s y \triangleq$ if $(q = \text{nil})$ then $y\langle q, P \rangle x$ else $x\langle q, P \rangle y$	

Fig. 5: The core SAM Transition Rules

is chosen to execute will immediately perform a write action on the freshly cut session x . In general, such a process may perform multiple write or read actions on many other sessions before performing the write on x , meaning that multiple context switches may occur. Given this, it is not obvious that this strategy is adequate insofar as preserving the correctness properties of CLL in terms of soundness, progress and type preservation. The remainder of this paper is devoted to establishing this correspondence in a precise technical sense.

4 CLLB: A Buffered Formulation of CLL

There is a substantial gap between the language CLL, presented in an abstract algebraic style, and its operational semantics, defined by equational and rewriting systems, and an abstract machine as the SAM, a deterministic state machine manipulating several low level structures. Therefore, even if the core SAM structure and transition rules are fairly simple, proving its correctness is more challenging and technically involved, and require progressive build up. Therefore, we first bridge between CLL and SAM via an intermediate logical language CLLB, which extends CLL with a buffered cut construct.

$$\text{cut } \{P \mid a : A \ [q] \ b : B \mid Q\}$$

The buffered cut construct models interaction via a “message queue” with two polarised endpoints a and b , held respectively by the processes P and Q . A polarised endpoint has the form x or \bar{x} . The endpoint marked \bar{x} is the only allowing writes, the unmarked y is the only one allowing reads, exactly one of

$$\begin{array}{c}
\frac{P \vdash^B \Delta', x : \bar{A}; \Gamma \quad Q \vdash^B \Delta, y : A; \Gamma}{\text{cut} \{P \mid x : \bar{A} \text{ [nil]} \bar{y} : A \mid Q\} \vdash^B \Delta', \Delta; \Gamma} \text{ (A positive) [TCutB]} \\
\\
\frac{\text{cut} \{\text{close } x \mid \bar{x} : \mathbf{1} [q] y : B \mid Q\} \vdash^B \Delta; \Gamma}{\text{cut} \{0 \mid \bar{x} : \emptyset [q@\checkmark] y : B \mid Q\} \vdash^B \Delta; \Gamma} \text{ [Tcut-1]} \\
\\
\frac{\text{cut} \{\text{send } x(y.R); P \mid \bar{x} : T \otimes A [q] y : B \mid Q\} \vdash^B \Delta; \Gamma}{\text{cut} \{P \mid \bar{x} : A [q@\text{clos}(y, R)] y : B \mid Q\} \vdash^B \Delta; \Gamma} \text{ [Tcut-}\otimes\text{]} \\
\\
\frac{\text{cut} \{\#l x; P \mid \bar{x} : \oplus_{\ell \in L} A_\ell [q] y : B \mid Q\} \vdash^B \Delta; \Gamma}{\text{cut} \{P \mid \bar{x} : A_{\#l} [q@\#l] y : B \mid Q\} \vdash^B \Delta; \Gamma} \text{ [Tcut-}\oplus\text{]} \\
\\
\frac{\text{cut} \{!x(z); P \mid \bar{x} : !A [q] y : B \mid Q\} \vdash^B \Delta; \Gamma}{\text{cut} \{0 \mid \bar{x} : \emptyset [q@\text{clos}!(z, P)] y : B \mid Q\} \vdash^B \Delta; \Gamma} \text{ [Tcut!]}
\end{array}$$

Fig. 6: Additional typing rules for CLLB.

the two endpoints is marked. The endpoints types A, B are of course related but do not need to be exact duals, the type of the writer endpoint may be advanced in time wrt the type of the reader endpoint, reflecting the messages already enqueued but not yet consumed. If the queue is empty, we have $A = \bar{B}$. Thus a buffered cut with empty queue corresponds to the basic cut of CLL.

$$\text{cut} \{P \mid x : A \mid Q\} \equiv \text{cut} \{P \mid \bar{x} : A \text{ [nil]} y : \bar{A} \mid \{y/x\}Q\} \quad (A+)$$

The queue q stores values V defined by

$$\begin{array}{l}
V ::= \checkmark \quad (\text{Close token}) \quad | \quad \#l \quad (\text{Selection Label}) \\
| \quad \text{clos}(x, P) \quad (\text{Linear Closure}) \quad | \quad \text{clos}!(x, P) \quad (\text{Exponential Closure})
\end{array}$$

$$q ::= \text{nil} \mid V \mid q@q \quad (\text{Queue})$$

The operator $@$ denotes (associative) concatenation operation of queues, with unit **nil**. Enqueue operations occur at the queue rhs, and dequeue at the lhs.

The type system CLLB is obtained from CLL by replacing [TCut] with the typing rules (and symmetric ones) in Fig. 6. We distinguish the type judgements as $P \vdash \Delta; \Gamma$ for CLL and $P \vdash^B \Delta; \Gamma$ for CLLB. The [TCutB] rule sets the endpoints mode based in the cut type polarity, applicable whenever the queue is empty. The remaining rules relate queue contents with their corresponding (positive action) processes. For instance, rule [Tcut- \otimes] can be read bottom-up as stating that typing processes mediated by a queue containing a process closure $\text{clos}(y, R)$ amounts to typing the process that will emit the session y (bound to R), interacting with the queue with the closure removed. Rules [Tcut- \oplus] and [Tcut!] apply a similar principle to the other possible queue contents. In [Tcut-1] and [Tcut!] the write endpoint is typed \emptyset , as the sender has terminated (0).

$$\begin{aligned}
\text{cut } \{Q \mid a : A[q]b : B \mid P\} &\equiv^{\text{B}} \text{cut } \{Q \mid b : B[q]a : A \mid P\} && [\text{comm}] \\
\text{cut } \{P \mid x[q]y \mid (Q \parallel R)\} &\equiv^{\text{B}} (\text{cut } \{P \mid x[q]y \mid Q\}) \parallel R && [\text{CM}] \\
P \mid x[q]z \mid (\text{cut } \{Q \mid y[p]w \mid R\}) &\equiv^{\text{B}} \text{cut } \{(\text{cut } \{P \mid x[q]z \mid Q\}) \mid y[p]w \mid R\} && [\text{CC}] \\
\text{cut } \{P \mid z[q]w \mid (\text{cut}! \{y.Q \mid !x \mid R\})\} &\equiv^{\text{B}} \text{cut}! \{y.Q \mid !x \mid (\text{cut } \{P \mid z[q]w \mid R\})\} && [\text{CC}!] \\
\text{cut}! \{y.P \mid !x \mid (\text{cut } \{Q \mid z[q]w \mid R\})\} &\equiv^{\text{B}} && \\
\text{cut } \{(y.P \mid !x \mid Q) \mid z[q]w \mid (\text{cut}! \{y.P \mid !x \mid R\})\} &&& [\text{D-C}]
\end{aligned}$$

Fig. 7: Additional structural congruence rules for CLLB.

Structural congruence for **B** (noted \equiv^{B}) is obtained by extending \equiv with commutative conversions for the buffered cut, listed in Fig. 7. The following provisos apply: [CM] $y \in \text{fn}(Q)$; in [CC] $y, z \in \text{fn}(Q)$; in [CC!] $x \notin \text{fn}(P)$. Accordingly, reduction for **B** (noted \rightarrow^{B}) is obtained by replacing the \rightarrow rules [fwd], [$\mathbf{1}\perp$], [$\otimes\wp$] and [$\oplus\&$] by the rules in Fig. 8. Essentially each principal cut reduction rule of CLL is replaced by a pair of “positive” (\rightarrow_p) / “negative” (\rightarrow_n) reduction rules that allow processes to interact asynchronously via the queue, that is, positive process actions (corresponding to positive types) are non-blocking. For example, the rule [\otimes] for send appends a session closure to the tail of the queue (rhs) and the rule for receive pops a session closure from the head of the queue (lhs). Notice that positive rules are enabled only if the relevant endpoint is in write mode (\bar{x}), and negative rules are enabled only if the relevant endpoint is in read mode (y). In [\wp] above the target cuts endpoint polarities depends on the types of the composed processes. To uniformly express the appropriate marking of endpoint polarities we define some convenient abbreviations:

Definition 4.1 (Setting polarities).

$$\begin{aligned}
\text{cut } \{Q \mid a : A[\mathbf{nil}]b : B \mid P\}^{\text{p}} &\triangleq \text{if } \bar{+}A \text{ then } \text{cut } \{Q \mid \bar{a} : A[\mathbf{nil}]b : B \mid P\} \\
&\quad \text{else } \text{cut } \{Q \mid a : A[\mathbf{nil}]\bar{b} : B \mid P\} \\
\text{cut } \{Q \mid \bar{a} : A[q]b : B \mid P\}^{\text{p}} &\triangleq \text{cut } \{Q \mid \bar{a} : A[q]b : B \mid P\} \quad (q \neq \mathbf{nil})
\end{aligned}$$

The following definition then formalizes the intuition given above about how to encode processes of CLL into processes of CLLB.

Definition 4.2 (Embedding). Let $P \vdash \Delta; \Gamma$. P^\dagger is the **B** process such that

$$(\text{cut } \{P \mid x : A \mid Q\})^\dagger \triangleq \text{cut } \{P^\dagger \mid x : A \mid \mathbf{nil} \mid y : \bar{A} \mid (\{y/x\}Q)^\dagger\}^{\text{p}}$$

homomorphically defined in the remaining constructs. Clearly $P^\dagger \vdash^{\text{B}} \Delta; \Gamma$.

4.1 Preservation and Progress for CLLB

In this section, we prove basic safety properties of CLLB: Preservation (Theorem 4.1) and Progress (Theorem 4.2). To reason about type derivations involving

$$\begin{array}{l}
\text{cut } \{Q \mid \bar{x} [q_1] x \mid \text{fwd } x y \mid \bar{y} [q_2] w \mid P\} \rightarrow^{\text{B}} \text{cut } \{Q \mid \bar{x} [q_2 @ q_1] w \mid P\} \quad [\text{fwdp}] \\
\text{cut } \{\text{close } x \mid \bar{x} [q] y \mid Q\} \rightarrow^{\text{B}} \text{cut } \{0 \mid \bar{x} [q @ \checkmark] y \mid Q\} \quad [\mathbf{1}] \\
\text{cut } \{0 \mid \bar{x} [\checkmark] y \mid \text{wait } y; P\} \rightarrow^{\text{B}} P \quad [\perp] \\
\text{cut } \{\text{send } x(z.P); Q \mid \bar{x} [q] y \mid R\} \rightarrow^{\text{B}} \text{cut } \{Q \mid \bar{x} [q @ \text{clos}(z, P)] \mid y \mid R\} \quad [\otimes] \\
\text{cut } \{Q \mid \bar{x} [\text{clos}(z, P) @ q] y \mid \text{recv } y(w); R\} \rightarrow^{\text{B}} \\
\quad \text{cut } \{Q \mid \bar{x} [q] y \mid \text{cut } \{P \mid z \mid \text{nil} \mid w \mid R\}^{\text{P}}\}^{\text{P}} \quad [\otimes] \\
\text{cut } \{\#l x; P \mid \bar{x} [q] y \mid R\} \rightarrow^{\text{B}} \text{cut } \{Q \mid \bar{x} [q @ \#l] \mid y \mid R\} \quad [\oplus] \\
\text{cut } \{Q \mid \bar{x} [l @ q] y \mid \text{case } y \{ \#l \in L : P_l \} \} \rightarrow^{\text{B}} \text{cut } \{Q \mid x [q] y \mid P_{\#l}\}^{\text{P}} \quad [\&] \\
\text{cut } \{!x(z); P \mid \bar{x} [q] y \mid Q\} \rightarrow^{\text{B}} \text{cut } \{0 \mid \bar{x} [q @ \text{clos}!(z, P)] \mid y \mid Q\} \quad [!] \\
\text{cut } \{0 \mid \bar{x} [\text{clos}!(y, P)] y \mid ?y; Q\} \rightarrow^{\text{B}} \text{cut}! \{y.P \mid !x \mid Q\} \quad [?]
\end{array}$$

Fig. 8: Reduction $P \rightarrow^{\text{B}} Q$.

buffered cuts, we formulate some auxiliary inversion principles that allow us, by aggregating sequences of application of [TCut-*] rules of CLLB, to talk in a uniform way about typing of values in queues and typing of processes connected by queues. To assert typing of queue values c we use judgments the form $\Gamma; \Delta \vdash c : E$, where E is either a type or a one hole type context, defined by

$$E ::= \square \mid T \mid T \wp E \mid \&_{\ell \in L} E_{\ell}$$

where in $\&_{\ell \in L} E_{\ell}$ only branch type $E_{\#l}$ for some selected label $\#l \in L$ is a one hole context (to plug the continuation type); only the branch chosen by the selected label in a queue is relevant to type next queue values. We identify the selected branch in the type by tagging it with the corresponding label $\#l$ thus $\&_{\ell \in L} E_{\ell}[\#l]$. We then introduce the following typing rules for queue values.

Definition 4.3 (Typing of Queue Values).

$$\frac{}{\Gamma; \vdash_{\text{val}} \checkmark : \perp} \quad \frac{P \vdash^{\text{B}} \Delta, z : T; \Gamma}{\Gamma; \Delta \vdash_{\text{val}} \text{clos}(z, P) : \bar{T} \wp E} \\
\frac{\&_{\ell \in L} E_{\ell}[\#l]}{\Gamma; \Delta \vdash_{\text{val}} \#l : \&_{\ell \in L} E_{\ell}} \quad \frac{P \vdash^{\text{B}} z : A; \Gamma}{\Gamma; \vdash_{\text{val}} \text{clos}!(z, P) : ?\bar{A}}$$

Given a sequence of k one hole queue value types E_i and a type A , we denote by $\mathbf{E}_k; A$ the type $E_1[E_2[\dots E_k[A]]]$. Queue value types allow us to talk in a uniform way about the type a receiver processes compatible with the types of enqueued values, as characterized by the following Lemma 4.1 and Lemma 4.2.

Lemma 4.1 (Non-full). *For $P \neq 0$ the rule below is admissible and invertible:*

$$\frac{P \vdash^{\text{B}} \Delta_P, x : A; \Gamma \quad Q \vdash^{\text{B}} \Delta_Q, y : B; \Gamma \quad q = \bar{c}_k \quad B = \mathbf{E}_k; \bar{A} \quad \Gamma; \Delta_i \vdash c_i : E_i \quad -B}{\text{cut } \{P \mid \bar{x} : A [q] y : B \mid Q\} \vdash^{\text{B}} \Delta_P, \Delta_Q, \Delta_1, \dots, \Delta_k; \Gamma}$$

Notice that a session type, as defined by a CLL proposition, may terminate in either $\mathbf{1}$, \perp or an exponential type $!A/?A$. We then also have

Lemma 4.2 (Full). *The proof rules below are admissible:*

$$\frac{Q \vdash^{\mathbf{B}} \Delta_Q, y : B; \Gamma \quad \Gamma; \Delta_i \vdash c_i : E_i \quad B = \mathbf{E}_k; \perp \quad c_k = \checkmark \quad -B}{\text{cut } \{0 \mid \bar{x} : \emptyset \mid [\bar{c}_k] y : B \mid Q\} \vdash^{\mathbf{B}} \Delta_Q, \Delta_1, \dots, \Delta_k; \Gamma}$$

$$\frac{Q \vdash^{\mathbf{B}} \Delta_Q, y : B; \Gamma \quad \Gamma; \Delta_i \vdash c_i : E_i \quad B = \mathbf{E}_{k-1}; C \quad \Gamma \vdash c_k = \text{clos!}(z, R) : C \quad -B}{\text{cut } \{0 \mid \bar{x} : \emptyset \mid [\bar{c}_k] y : B \mid Q\} \vdash^{\mathbf{B}} \Delta_Q, \Delta_1, \dots, \Delta_k; \Gamma}$$

Moreover, one of them must apply for inverting the judgment in the conclusion.

Lemma 4.3 (Non-empty). *Let $\text{cut } \{P \mid \bar{x} : A \mid [q] y : B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$.*

If A is negative or void, then $q \neq \text{nil}$.

Proof. If A is void then immediate by Lemma 4.2. If $-A$, suppose $q = \text{nil}$. By Lemma 4.1 $k = 0$ and $\bar{A} = B$. But then B positive, contradiction. So $q \neq \text{nil}$.

Theorem 4.1 (Preservation). *Let $P \vdash^{\mathbf{B}} \Delta; \Gamma$.*

(1) If $P \equiv^{\mathbf{B}} Q$, then $Q \vdash^{\mathbf{B}} \Delta; \Gamma$. (2) If $P \rightarrow^{\mathbf{B}} Q$, then $Q \vdash^{\mathbf{B}} \Delta; \Gamma$.

Proof. We verify that rules for $\equiv^{\mathbf{B}}$ (Fig. 7) (resp. $\rightarrow^{\mathbf{B}}$ (Fig. 8)) are type preserving. We illustrate with (2) [fwdp].

(Case [fwdp]) $P = \text{cut } \{Q' \mid \bar{z} : A \mid [q_1] x : \bar{B} \mid \text{cut } \{\text{fwd } x y \mid \bar{y} : B \mid [q_2] w : C \mid P'\}\} \vdash^{\mathbf{B}} \Delta; \Gamma$. If $q_2 = \text{nil}$ and $B = \bar{C}$, hence $Q = \text{cut } \{Q' \mid \bar{z} : A \mid [q_1] w : C \mid P'\} \vdash^{\mathbf{B}} \Delta; \Gamma$.

Otherwise $q_2 \neq \text{nil}$. Let $F_2 = \text{cut } \{\text{fwd } x y \mid \bar{y} : B \mid [q_2] w : C \mid P'\}$ where $F_2 \vdash^{\mathbf{B}} \Delta_2, x : \bar{B}$ and $\Delta = \Delta_1, \Delta_2$, and $q_2 = \bar{c}_k$. By Lemma 4.1, $\text{fwd } x y \vdash^{\mathbf{B}} x : \bar{B}, y : B$, and $C = \mathbf{E}_k; \bar{B}, \Gamma; \Delta_i \vdash c_i : E_i, P \vdash w : C, \Delta_P; \Gamma$ and $\Delta_2 = \Delta_1, \dots, \Delta_k, \Delta_P$.

Let $F_1 = \text{cut } \{Q' \mid \bar{z} : A \mid [q_1] x : \bar{B} \mid F_2\}$ with $F_1 \vdash^{\mathbf{B}} \Delta_1, x : B; \Gamma$ and $q_1 = \bar{d}_l$. By Lemma 4.1, $Q' \vdash^{\mathbf{B}} \Delta_Q, z : A; \Gamma, \Gamma; \Delta'_j \vdash d_j : F_j$ and $\bar{B} = \mathbf{F}_l; \bar{A}$ and $\Delta_1 = \Delta'_1, \dots, \Delta'_l, \Delta_Q$. By Lemma 4.1 we get $\text{cut } \{Q' \mid z : A \mid [q_2 @ q_1] w : C \mid P'\} \vdash^{\mathbf{B}} \Delta; \Gamma$.

A process P is *live* if and only if $P = \mathcal{C}[Q]$, for some static context \mathcal{C} (the hole lies within the scope of static constructs mix, cut) and Q is an action process. We first show that a live process either reduces or offers an interaction on a free name. The observability predicate defined in Fig. 9 (cf. [61]) characterises interactions of a process with the environment.

Lemma 4.4 (Liveness). *Let $P \vdash^{\mathbf{B}} \Delta; \Gamma$ be live. Either $P \downarrow_x$ or $P \rightarrow$.*

Proof. By induction on the derivation for $P \vdash \Delta; \Gamma$, and case analysis on the last typing rule. We illustrate with one rule.

(Case of [TCut \otimes]) We have $P = \text{cut } \{P_1 \mid \bar{x} : A \mid [q] y : B \mid P_2\} \vdash^{\mathbf{B}} \Delta; \Gamma$ where $q = r @ \text{clos}(z_k; R_k)$ derived from $P_1 \vdash^{\mathbf{B}} x : A, \Delta_{P_1}; \Gamma$. By i.h. $P_1 \rightarrow$ or $P_1 \downarrow_{x_1}$. If $x \neq x_1$ then $P_1 \downarrow_{x_1}$ and $P \downarrow_{x_1}$. If A is positive, by the same reasoning as above for P_2 we conclude that $P \rightarrow$ or $P \downarrow_w$.

$$\begin{array}{c}
\frac{}{\text{fwd } x \ y \downarrow_x} \text{ [fwd]} \quad \frac{s(\mathcal{A}) = x}{\mathcal{A} \downarrow_x} [\mathcal{A}] \quad \frac{P \equiv Q \quad Q \downarrow_x}{P \downarrow_x} [\equiv] \quad \frac{P \downarrow_x}{(P \parallel Q) \downarrow_x} \text{ [mix]} \\
\frac{P \downarrow_x \quad x \neq y}{(P \mid y[q]x \mid Q) \downarrow_x} \text{ [cut]} \quad \frac{Q \downarrow_x \quad x \neq y}{(z.P \mid !y \mid Q) \downarrow_x} \text{ [cut!]}
\end{array}$$

Fig. 9: Observability Predicate $P \downarrow_x$.

Otherwise, A is negative. By i.h. $P_2 \rightarrow$ or $P_2 \downarrow_z$. If $z \neq y$ then $P \downarrow_x$. Otherwise, $P_2 \downarrow_y$. By Lemma 4.1, $B = \bar{T} \wp C$. By Lemma A.1 (2,6), either $P_2 \equiv \text{fwd } y \ v \ \mid * \ \mid Q'$ (a) or $P_2 \equiv \text{recv } y(w); Q' \ \mid * \ \mid Q''$ (b). For case (a) we conclude as in [TCutB] above that $P \rightarrow$ or $P \downarrow_v$, for case (b) $P \rightarrow$ by $[\wp]$.

Theorem 4.2 (Progress). *Let $P \vdash \emptyset; \emptyset$ be a live process. Then, $P \rightarrow$.*

Proof. Follows from Lemma 4.4 since $\Delta = \emptyset$.

4.2 Correspondence between CLL and CLLB

In this section we establish the correspondence between reduction in CLL and CLLB, proving that the two languages simulate each other in a tight sense. Intuitively, the correspondence shows that CLLB allows some positive actions to be buffered ahead of reception, while in CLL a single positive action synchronises with the corresponding dual in one step, or a forward reduction takes place.

We write a reduction $P \rightarrow^B Q$ as $P \rightarrow^{Bp} Q$ if the reduced action is positive, $P \rightarrow^{Bn} Q$ if the reduced action is negative (we consider [call] negative), $P \rightarrow^{Ba} Q$ if the reduced action is a forwarder, and $P \rightarrow^{Bap} Q$ if the reduced action is positive or a forwarder. We also write $P \rightarrow^{Br} Q$ for positive action followed by a matching negative action on the same cut with an initially empty queue.

Lemma 4.5. *The following commutations of reductions hold.*

1. Let $P_1 \rightarrow^{Bp} S \rightarrow^{Bn} P_2$. Either $P_1 \rightarrow^{Br} P_2$, or $P_1 \rightarrow^{Bn} S' \rightarrow^{Bp} P_2$ for some S' .
2. Let $P_1 \rightarrow^{Ba} S \rightarrow^{Bn} P_2$. Then $P_1 \rightarrow^{Bn} S' \rightarrow^{Ba} P_2$ for some S' .
3. If $P_1 \rightarrow^{Bap} S \rightarrow^{Bn} P_2$, either $P_1 \rightarrow^{Br} P_2$, or $P_1 \rightarrow^{Bn} S' \rightarrow^{Bap} P_2$ for some S' .
4. Let $P_1 \rightarrow^{Bap} N \xrightarrow{\epsilon}^{Ba} S \rightarrow^{Br} P_2$. Either $P_1 \xrightarrow{\epsilon}^{Ba} N$ or $P_1 \rightarrow^{Br} S' \rightarrow^{Bap} P_2$ for some S' .

Lemma 4.6 (Simulation). *Let $P \vdash \emptyset; \emptyset$. If $P \rightarrow Q$ then $P^\dagger \Rightarrow^B Q^\dagger$.*

Proof. Each cut reduction of CLL is either simulated by two reduction steps of B in sequence or by a [fwd] reduction.

The following lemma identifies that in CLLB, a sequence of positive actions (or forwards) followed by a negative action can always be commuted either by pulling out the negative action first, followed by the sequence of positive actions and forwards; having the negative action follow a positive action on the same channel and then performing the remaining actions; or by first performing a sequence of forward actions, the output and input on the relevant session and then the remaining actions.

Lemma 4.7 (Simulation). *Let $P \vdash^B \emptyset; \emptyset$. If $P \Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}} Q$ then (1) $P \rightarrow^{\text{Bn}} R$ and $R \Rightarrow^{\text{Bap}} Q$ for some R , or; (2) $P \rightarrow^{\text{Br}} R$ and $R \Rightarrow^{\text{Bap}} Q$ for some R , or; (3) $P \xrightarrow{\epsilon}^{\text{Ba}} \rightarrow^{\text{Br}} R$ and $R \Rightarrow^{\text{Bap}} Q$ for some R .*

Proof. By induction on $P(\rightarrow^{\text{Bap}})^* P'$.

Theorem 4.3 (Operational correspondence CLL-CLLB). *Let $P \vdash \emptyset; \emptyset$.*

1. *If $P \Rightarrow R$ then $P^\dagger \Rightarrow^B R^\dagger$.*
2. *If $P^\dagger(\Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}})^* Q$ then there is R such that $P \Rightarrow R$ and $R^\dagger \Rightarrow^{\text{Bap}} Q$.*

Due to the progress property for CLLB (Theorem 4.2) and because queues are bounded by the size of positive/negative sections in types, after a sequence of positive or forwarder reductions a negative reduction consuming a queue value must occur. Theorem 4.3(2) states that every reduction sequence in CLLB is simulated by a reduction sequence in CLL up to some anticipated forwarding and buffering of positive actions. Our results imply that every reduction path in CLLB maps to a reduction path in CLL in which every negative reduction step in the former is mapped, in order, to a cut reduction step in the latter.

5 Correctness of the core SAM

We now prove that every execution trace of the core SAM defined in Fig. 5 represents a correct process reduction sequence CLLB (and therefore of CLL, in the light of Theorem 4.3), first for the language without exponentials and mix, which will be treated in Section 6. In what follows, we annotate endpoints of session records with their types (e.g. as $x:A(q, P)y:B$), these annotations are not needed to guide the operation of the SAM, but convenient for the proofs; they will be omitted when not relevant or are obvious from the context. We first define a simple encoding of well-typed CLLB processes to SAM states.

Definition 5.1 (Encode). *Given $P \vdash^B \emptyset$ we define $\text{enc}(P) = C$ as $\text{enc}(P, \emptyset) \xRightarrow{\text{cut}^*} C$ where $\text{enc}(P, H) \xRightarrow{\text{cut}^*} C$ is defined by the rules*

$$\frac{\text{enc}(P(x), H[x:A(q, Q)y:B]) \xRightarrow{\text{cut}^*} C}{\text{enc}(\text{cut} \{P \mid \bar{x}:A[q] y:B\} \mid Q), H) \xRightarrow{\text{cut}^*} C} \quad (A+)$$

$$\frac{\text{enc}(Q(y), H[x:A(q, P)y:B]) \xRightarrow{\text{cut}^*} C}{\text{enc}(\text{cut} \{P \mid \bar{x}:A[q] y:B\} \mid Q), H) \xRightarrow{\text{cut}^*} C} \quad (A- \text{ or } P = 0)$$

$$\text{enc}(A, H) \xRightarrow{\text{cut}^*} (A, H) \quad (A \in \mathcal{A})$$

Notice that $\text{enc}(P)$ maximally applies the SAM execution rule for cut to (P, \emptyset) until an action is reached. Clearly, for any $P \vdash^B \emptyset$, if $\text{enc}(P) = C$ then $P \Rightarrow^* C$. Also, if all cuts in a state C have empty queues then there is a process Q of CLL such that $\text{enc}(Q^\dagger) = C$. We then have

Theorem 5.1 (Soundness wrt CLLB). *Let $P \vdash^B \emptyset$.*

If $\text{enc}(P) \Rightarrow D \xRightarrow{\text{cut}^} C$ then there is Q such that $P \rightarrow \cup \equiv Q$ and $C = \text{enc}(Q)$.*

Proof. Let $enc(P) \Rightarrow \mathcal{B}$. Let $enc(P) = (\mathcal{A}, H)$ for some action \mathcal{A} . Then $P \equiv E[\mathcal{A}]$ for some cut context $E[-]$. We illustrate one case.

(Case of [S1]) Let $enc(P) = (\text{close } x, H[x\langle q, R \rangle y]) \Rightarrow (R, H[x\langle q @ \checkmark, 0 \rangle y]) = D \xrightarrow{\text{cut}^*} C$. Therefore, $P \equiv E[\text{close } x] \equiv E'[\text{cut } \{\text{close } x \mid \bar{x} [q] y \mid R\}]$ and $P \rightarrow Q = E'[\text{cut } \{0 \mid \bar{x} [q @ \checkmark] y \mid R\}]$ and $enc(Q) = C$.

We can then combine soundness with the operational correspondence between CLL and CLLB (Theorem 4.3) to obtain an overall soundness result for the SAM with respect to CLL:

Theorem 5.2 (Soundness wrt CLL). *Let $P \vdash^B \emptyset$.*

1. *If $enc(P) \xrightarrow{\text{cut}^*} C$ there is Q such that $P \Rightarrow \cup \equiv Q$ and $C = enc(Q)$.*
- Let $P \vdash \emptyset$. *If $enc(P^\dagger) \xrightarrow{\text{cut}^*} enc(Q^\dagger)$ then $P \Rightarrow Q$.*

In Definition 5.2 we identify readiness, the fundamental invariant property of SAM states, key to prove progress of its execution strategy. Readiness means that any running process holding an endpoint of negative type, and thus attempting to execute a negative action (e.g., a receive or offer action) on it, will always find an appropriate value (resp. a closure or a label) to be read in the appropriate session queue. No busy waiting or context switching will be necessary since the sequential execution semantics of the SAM enforces that all actions corresponding to a positive section of a session type have always been enqueued by the “caller” process before the “callee” takes over. As discussed in Section 3 it might not seem obvious whether all such input endpoints, (including endpoints moved around via send / receive interactions), always refer to non-empty queues.

Readiness must also be maintained by processes suspended in session records, even if a suspended process waiting on a read endpoint will not necessarily have the corresponding queue already populated. Intuitively, a process P is (H, N) -ready if all its “reads” in the input channels (except those in N) will be matched by values already stored in the corresponding session queue.

Definition 5.2 (Ready).

Process P is H, N -ready if for all $y \in \text{fn}(P) \setminus N$ and $x : A\langle q, R \rangle y \in H$ then A is negative or void. We abbreviate H, \emptyset -ready by H -ready.

Heap H is ready if, for all $x\langle q, R \rangle y \in H$, the following conditions hold:

1. *if $R(y)$ then R is $H, \{y\}$ -ready*
2. *if $R(x)$ then R is H -ready*
3. *if $\text{clos}(z : -, R) \in q$, R is $H, \{z\}$ -ready.*
4. *if $\text{clos}(z : +, R) \in q$, R is H -ready.*

State $C = (P, H)$ is ready if H is ready and P is H -ready.

Lemma 5.1 (Readiness). *Let $P \vdash \emptyset$ and $(P, \emptyset) \xrightarrow{\text{cut}^*} S$. Then S is ready.*

Proof. The property trivially holds for (P, \emptyset) for $\text{fn}(P) = \emptyset$ and $H = \emptyset$. We proceed by transition induction, assuming that $S = (P, H)$ is ready, and $S \Rightarrow S' = (P, H')$, we check that S' is ready. We illustrate one case.

S	$::= (\mathcal{E}, P, H)$	State
H	$::= Ref \rightarrow SessionRec$	Heap
$SessionRec$	$::= x\langle q, \mathcal{E}, P \rangle y$	
q	$::= \mathbf{nil} \mid Val \mid q@q$	Queue
Val	$::= \checkmark$	Close token
	$\mid \#l$	Choice label
	$\mid \mathbf{clos}(x, \mathcal{E}, P)$	Linear Closure
	$\mid \mathbf{clos!}(x, \mathcal{E}, P)$	Exponential Closure
$\mathcal{E}, \mathcal{G}, \mathcal{F}$	$::= Name \rightarrow (Ref \cup Val)$	Environment x

Fig. 10: The SAM

(Case of [SCut]) Let $R = \mathbf{cut} \{P \mid \bar{x} : T [q] y : B \mid Q\}$ and $S = (R, H) \Rightarrow S' = (P, H')$, where $H' = H[x\langle \mathbf{nil}, Q \rangle y]$ and T positive. By i.h., R is H -ready hence $P(x)$ is $H, \{x\}$ -ready and $Q(y)$ is $H, \{y\}$ -ready. Then $P(x)$ is H' -ready, and H' is ready. We conclude that S' is ready.

Theorem 5.3 (Progress). *Let $P \vdash^B \emptyset$ and P live. Then $enc(P) \Rightarrow S'$.*

Proof. Since P is live then $enc(P) = S = (\mathcal{A}, H)$ for some action \mathcal{A} .

(Case of [\perp]) $\mathcal{A} = \mathbf{wait} \ y; R'$. If $H = H'[y:A\langle q, Q \rangle x]$, then $S \Rightarrow S'$ by [S-]. Otherwise, $H = H'[x:A\langle q, \mathcal{A} \rangle y]$. By Lemma 5.1, \mathcal{A} is H' -ready, so A is negative or void. By Lemma 4.3 $q \neq \mathbf{nil}$. By Lemma 4.2, $q = \checkmark$. So $S \Rightarrow S'$ by [S \perp].

6 The SAM for full CLL

In this section, we complete our initial presentation of the SAM, in particular, we introduce support for the exponentials, allowing the machine to compute with non-linear values, and a selective concurrency semantics. We have delayed the introduction of an environment structure for the SAM, to make the presentation easier to follow. However, this was done at the expense of a more abstract formalisation of the operational semantics, making use of α -conversion, and overloading language syntax names as heap references for allocated session records.

The SAM actually relies on environment-based implementation of name management, presented in Fig. 6. A SAM state is then a triple (\mathcal{E}, P, H) where \mathcal{E} is an environment that maps each free name of the code P into either a closure or a heap record endpoint. These heap references are freshly allocated and unique, thus avoiding any clashes and enforcing proper static scoping. Closures, representing suspended linear ($\mathbf{clos}(z, \mathcal{E}, P)$) and exponential behaviour ($\mathbf{clos!}(z, \mathcal{E}, P)$), pair the code in its environment, and we expect the following structural safety conditions for name bidding in configurations to hold.

Definition 6.1 (Closure).

A process P is (\mathcal{E}, N) -closed if $\text{fn}(P) \setminus N \subseteq \text{dom}(\mathcal{E})$, and \mathcal{E} -closed if (\mathcal{E}, \emptyset) -closed. Environment \mathcal{E} is H -closed if for all $x \in \text{dom}(\mathcal{E})$ if $\mathcal{E}(x)$ is a reference then $x \in H$, if $\mathcal{E}(x) = \text{clos!}(z, \mathcal{F}, R)$ then \mathcal{F} is H -closed and R is $(\mathcal{F}, \{z\})$ -closed. Heap H is closed if for all $x(q, \mathcal{G}, Q)y \in H$, \mathcal{G} is H -closed, Q is \mathcal{G} -closed, and for all $\text{clos}(z, \mathcal{F}, R) \in q$ and $\text{clos!}(z, \mathcal{F}, R) \in q$, \mathcal{F} is H closed and R is $(\mathcal{F}, \{z\})$ -closed. State (\mathcal{E}, P, H) is closed if H is closed, \mathcal{E} is H -closed, and P is \mathcal{E} -closed.

In Figure 6 we present the environment-based execution rules for the SAM. All rules except those for exponentials have already been essentially presented in Fig. 5 and discussed in previous sections. The only changes to those rules are due to the presence of environments, which at all times record the bindings for free names in the code. Overall, we have

Lemma 6.1. *Let $P \vdash^{\text{B}} \emptyset; \emptyset$. For all S such that $(P, \emptyset, \emptyset) \xrightarrow{*} S$, S is closed.*

We discuss the SAM rules for the exponentials. Values of exponential type are represented by exponential closures $\text{clos!}(z, \mathcal{F}, R)$. Recall that a session type may terminate in either type $\mathbf{1}$, type \perp or in an exponential type $!A/?A$ (cf. 4.2). So, the (positive) execution rule [S!] is similar to rule [S1]: it enqueues the closure representing the replicated process, and switches context, since the session terminates (cf. [!] Fig. 8). The execution rule [S?] is similar to rule [S?]: it pops a closure from the queue (which, in this case, always becomes empty), and instead of using it immediately, adds it to the environment to become persistently available to client code (cf. reduction rule [S?] Fig. 8). Any such closure representing a replicated process may be called by client code with transition rule [Scall], which essentially creates a new linear session composed by cut with the client code, similarly to [S?]. Rule [SCall] operates with some similarity to rule [S?]: instead of activating a linear closure popped from the queue, it activate an exponential closure fetched from the environment.

We extend the *enc* map to the exponential cut and environment states (\mathcal{E}, P, H) by adapting Definition 5.1, and adding the clause:

$$\frac{\text{enc}(\mathcal{E}\{\text{clos!}(y, \mathcal{E}, R)/x\}, P), H \xrightarrow{\text{cut}^*} C}{\text{enc}(\mathcal{E}, \text{cut!}\{y.R \mid x\} P), H \xrightarrow{\text{cut}^*} C}$$

We now update our meta-theoretical results for the complete SAM.

Theorem 6.1 (Soundness). *Let $P \vdash^{\text{B}} \emptyset; \emptyset$.*

If $\text{enc}(P) \Rightarrow D \xrightarrow{\text{cut}^} C$ then there is Q such that $P \rightarrow \cup \equiv Q$ and $C = \text{enc}(Q)$.*

Theorem 6.2 (Progress). *Let $P \vdash^{\text{B}} \emptyset; \emptyset$ and P live. Then $\text{enc}(P) \Rightarrow C$.*

6.1 Concurrent Semantics of Cut and Mix

Intuitively, the execution of mix $P \parallel Q$ consists in the parallel execution of (non-interfering) processes P and Q . We may execute $P \parallel Q$ by sequentialising P and Q in some arbitrary way, and this actually may be useful in some cases.

However, much more interesting is the accommodation in the SAM of interfering concurrency, as required to support full-fledged concurrent languages for

$$\begin{array}{l}
(\mathcal{E}, \text{cut } \{P \mid \bar{x} : A [\text{nil}] y : B \mid Q\}, H) \Rightarrow (\mathcal{G}, P, H[a\langle q, \mathcal{F}, Q \rangle b]) \quad [\text{SCut}] \\
a, b = \text{new}, \mathcal{G} = \mathcal{E}\{a/x\}, \mathcal{F} = \mathcal{E}\{b/y\} \\
(\mathcal{E}, \text{close } x, H[a\langle q, \mathcal{F}, P \rangle b]) \Rightarrow (\mathcal{F}, P, H[a\langle q @ \checkmark, \emptyset, 0 \rangle b]) \quad [\text{S1}] \\
a = \mathcal{E}(x) \\
(\mathcal{E}, \text{fwd } x \ y, H[c\langle q_1, \mathcal{G}, Q \rangle a][b\langle q_2, \mathcal{F}, P \rangle d]) \Rightarrow (\mathcal{F}, P, H[c\langle q_2 @ q_1, \mathcal{G}, Q \rangle d]) \quad [\text{Sfwd}] \\
a = \mathcal{E}(x), b = \mathcal{E}(y) \\
(\mathcal{E}, \text{wait } y; P, H[a\langle \checkmark, \emptyset, 0 \rangle b]) \Rightarrow (\mathcal{E}, P, H) \quad [\text{S}\perp] \\
b = \mathcal{E}(y) \\
(\mathcal{E}, \mathcal{A}^-(x), H[a\langle q, \mathcal{G}, Q \rangle b]) \Rightarrow (\mathcal{G}, Q, H[a\langle q, \mathcal{E}, \mathcal{A}^-(x) \rangle b]) \quad [\text{S-}] \\
a = \mathcal{E}(x) \\
(\mathcal{E}, \text{send } x(z.R); Q, H[a\langle q, P \rangle b]) \Rightarrow \quad [\text{S}\otimes] \\
(\mathcal{E}, Q, H[a\langle q @ \text{clos}(z, \mathcal{E}, R), P \rangle b]) \\
a = \mathcal{E}(x) \\
(\mathcal{E}, \text{recv } y(w:+); Q, H[a\langle \text{clos}(z, \mathcal{F}, R) @ q, \mathcal{G}, P \rangle b]) \Rightarrow \quad [\text{S}\otimes+] \\
(\mathcal{E}', Q, H[e\langle \text{nil}, \mathcal{F}', R \rangle f][a\langle q \rangle^s b]) \\
e, f = \text{new}, b = \mathcal{E}(y), \mathcal{E}' = \mathcal{E}\{e/w\}, \mathcal{F}' = \mathcal{F}\{f/z\} \\
(\mathcal{E}, \text{recv } y(w:-); Q, H[a\langle \text{clos}(z, \mathcal{F}, R) @ q, \mathcal{G}, P \rangle b]) \Rightarrow \quad [\text{S}\otimes-] \\
(\mathcal{F}', R, H[e\langle \text{nil}, \mathcal{E}', Q \rangle f][a\langle q \rangle^s b]) \\
e, f = \text{new}, b = \mathcal{E}(y), \mathcal{F}' = \mathcal{F}\{e/z\}, \mathcal{E}' = \mathcal{E}\{f/w\} \\
(\mathcal{E}, \#l \ x; Q, H[a\langle q, \mathcal{G}, P \rangle b]) \Rightarrow (\mathcal{E}, Q, H[a\langle q @ \#l, \mathcal{G}, P \rangle b]) \quad [\text{S}\oplus] \\
a = \mathcal{E}(x) \\
(\mathcal{E}, \text{case } y \{ \#l \in L : Q_\ell \}, H[a\langle \#l @ q, \mathcal{G}, P \rangle b]) \Rightarrow (\mathcal{E}, Q_{\#l}, H[a\langle q \rangle^s b]) \quad [\text{S}\&] \\
b = \mathcal{E}(y) \\
(\mathcal{E}, !x(z); Q, H[a\langle q, \mathcal{G}, P \rangle b]) \Rightarrow (\mathcal{G}, P, H[a\langle q @ \text{clos}(z, \mathcal{E}, Q), \emptyset, 0 \rangle b]) \quad [\text{S!}] \\
a = \mathcal{E}(x) \\
(\mathcal{E}, ?y; Q, H[a\langle \text{clos}(z, \mathcal{F}, R), \emptyset, 0 \rangle b]) \Rightarrow (\mathcal{E}', Q, H) \quad [\text{S?}] \\
b = \mathcal{E}(y), \mathcal{E}' = \mathcal{E}\{\text{clos}(z, \mathcal{F}, R)/y\} \\
(\mathcal{E}, \text{call } y(w:+); Q, H) \Rightarrow (\mathcal{E}', Q, H[a\langle \text{nil}, \mathcal{F}', R \rangle b]) \quad [\text{Scall+}] \\
a, b = \text{new}, \mathcal{E}' = \mathcal{E}\{a/w\}, \mathcal{F}' = \mathcal{F}\{b/z\} \\
\text{clos}(z, \mathcal{F}, R) = \mathcal{E}(y) \\
(\mathcal{E}, \text{call } y(w:-); Q, H) \Rightarrow (\mathcal{F}', R, H[a\langle \text{nil}, \mathcal{E}', Q \rangle b]) \quad [\text{Scall-}] \\
a, b = \text{new}, \mathcal{E}' = \mathcal{E}\{b/w\}, \mathcal{F}' = \mathcal{F}\{a/z\} \\
\text{clos}(z, \mathcal{F}, R) = \mathcal{E}(y)
\end{array}$$

$a\langle q \rangle^s b \triangleq$ if $(q = \text{nil})$ then $b\langle q, \mathcal{G}, P \rangle a$ else $a\langle q, \mathcal{G}, P \rangle b$

Fig. 11: SAM Transition Rules for the complete CLL

session-based programming. We sketch here the development, more details in Appendix A.5. First, we evolve the SAM from single threaded to multithreaded, where states now expose a multiset of processes P_i ready for execution by the basic SAM sequential transitions: $(\{P_1, P_2, \dots, P_n\}, H)$ and introduce an annotated variant **pcut** of the cut. It has the same CLL/CLLB semantics, but to be implemented as a fork construct where P and Q spawn concurrently, their interaction mediated by an atomic *concurrent session record* $x \langle q \rangle y$. The type system ensuring that concurrent channels may be forwarded only to concurrent channels. We extend the SAM with transition rule for multisets:

$$\frac{(P, H) \Rightarrow (P', H')}{(P \uplus T, H) \Rightarrow (P' \uplus T, H')} \text{ [Srun]}$$

$$(\text{pcut } \{P \mid \bar{x}:A [q] y:B \mid Q\} \uplus T, H) \Rightarrow (\{P, Q\} \uplus T, H[x \langle \text{nil} \rangle y]) \text{ [SCutp]}$$

$$((P \parallel Q) \uplus T, H) \Rightarrow (\{P, Q\} \uplus T, H[x \langle \text{nil} \rangle y]) \text{ [SMixp]}$$

Each individual thread executes locally according to the SAM sequential transitions presented before, until an action on a concurrent queue is reached. Concurrent process actions on concurrent queues are atomic, and defined as expected. Positive actions always progress by pushing a value into the queue, while negative actions will either pop off a value from the queue or block, waiting for a value to become available. We illustrate with the rules for $\mathbf{1}, \perp$ typed actions.

$$(\text{close } x, H[x \langle q \rangle y]) \Rightarrow (0, H[x \langle q @ \checkmark \rangle y]) \text{ [S1c]}$$

$$(\text{wait } y; P, H[x \langle \checkmark, y \rangle]) \Rightarrow (P, H) \text{ [S\perp c]}$$

Notice that, as in the case for **wait** $y; P$ above, any negative action in the thread queue is unable to progress if the corresponding queue is empty. It should be clear how to define transition rules for all other pairs of dual actions. Given an appropriate encoding enc^c of annotated CLLB processes in concurrent SAM states, and as consequence of typing and leveraging the proof scheme for progress in CLLB (Theorem 4.2), we have:

Theorem 6.3 (Soundness-c). *Let $P \vdash^B \emptyset; \emptyset$. If $enc^c(P) \Rightarrow D \xRightarrow{\text{cut}^*} C$ then there is Q such that $P \rightarrow \cup \equiv Q$ and $C = enc^c(Q)$.*

Theorem 6.4 (Progress-c). *Let $P \vdash^B \emptyset; \emptyset$ and P live. Then $enc_c(P) \Rightarrow C$.*

The extended SAM executes concurrent session programs, consisting in an arbitrary number of concurrent threads. Each thread deterministically executes sequential code, but can at any moment spawn new concurrent threads. The whole model is expressed in the common language of (classical) linear logic, statically ensuring safety, proper resource usage, termination, and deadlock absence by static typing.

7 Concluding Remarks and Related Work

We introduce the Session Abstract Machine, or SAM, an abstract machine for executing session processes typed by (classical) linear logic CLL, deriving a deterministic, sequential evaluation strategy, where exactly one process is executing

at any given point in time. In the SAM, session channels are implemented as single queues with a write and a read endpoint, which are written to, and read by executing processes. Positive actions are non-blocking, giving rise to a degree of asynchrony. However, processes in a session synchronise at polarity inversions, where they alternate execution, according to a fixed co-routining strategy. Despite its specific strategy, the SAM semantics is sound wrt CLL and satisfies the correctness properties of logic-based session type systems. We also present a conservative concurrent extension of the SAM, allowing the degrees of concurrency to be modularly expressed at a fine grain, ranging from fully sequential to fully concurrent execution. Indeed, a practical concern with the SAM design lies in providing a principled foundation for an execution environment for multi-paradigm languages, combining concurrent, imperative and functional programming. The overall SAM design as presented here may be uniformly extended to cover any other polarised language constructs that conservatively extend the PaT paradigm, such as polymorphism, affine types, recursive and co-recursive types, and shared state [53, 59]. We have implemented a SAM-based version of an open-source implementation of CLL in [60], and written many executable examples (e.g., realistic implementations of several concurrent ADTs), which we plan to submit as a companion artifact for this paper.

A machine model provides evidence of the algorithmic feasibility of a programming language abstract semantics, and illuminates its operational meaning from certain concrete semantic perspective. Since the seminal work of Landin on the SECD [41], several machines to support the execution of programs for a given programming language have been proposed. The SAM is then proposed herein in this same spirit of Cousineau, Curien and Mauny’s Categorical Abstract Machine for the call-by-value λ -calculus [19], Lafont’s Linear Abstract Machine for the linear λ -calculus [39], and Krivine’s Machine for the call-by-name λ -calculus [38] ; these works explored Curry-Howard correspondences to propose provably correct solutions. In [20], Danvy developed a deconstruction of the SECD based on a sequence of program transformations. The SAM is also derived from Curry-Howard correspondences for linear logic CLL [15, 69], and we also rely on program conversions, via the intermediate buffered language CLLB, as a key proof technique. We believe that the SAM is the first proposal of its kind to tackle the challenges of a process language, while building on several deep properties of its type structure towards a principled design. Among those, focusing [4] and polarisation [42, 30, 52] played an important role to achieve a deterministic sequential reduction strategy for session-based programming, perhaps our main initial motivation. That allows the SAM to naturally and efficiently integrate the execution of sequential and concurrent session behaviours, and suggests effective compilation schemes for mainstream virtual machines or compiler frameworks.

The adoption of session and linear types is clearly increasing in research (e.g., [24, 3, 56, 22, 71, 63, 52, 59]) and general purpose languages (e.g., Haskell [8, 36], Rust [40, 18] Ocaml [33, 49], F# [48], Move [9], among many others), which either require sophisticated encodings of linear typing via type-level computa-

tion or forego of some static correctness properties for usability purposes. Such developments typically have as a main focus the realization of the session typing discipline (or of a particular refinement of such typing), with the underlying concurrent execution model often offloaded to existing language infrastructure.

We highlight the work [17], which studies the relationship between synchronous session types and game semantics, which are fundamentally asynchronous. Their work proposes an encoding of synchronous strategies into asynchronous strategies by so-called call-return protocols. While their focus differs significantly from ours, the encoding via asynchrony is reminiscent of our own.

In future work, we plan to study the semantics of the SAM in terms of games (and categories), along the lines of [17, 19, 39]. We also plan to investigate the ways in which the evaluation strategy of the SAM can be leveraged to develop efficient compilation of fine-grained session-based programming, and its relationship with effect handlers, coroutines and delimited continuations.

References

1. Abramsky, S.: Computational Interpretations of Linear Logic. *Theoret. Comput. Sci.* **111**(1–2), 3–57 (1993)
2. Abramsky, S., Gay, S.J., Nagarajan, R.: Interaction categories and the foundations of typed concurrent programming. In: NATO ASI DPD. pp. 35–113 (1996)
3. Almeida, B., Mordido, A., Thiemann, P., Vasconcelos, V.T.: Polymorphic lambda calculus with context-free session types. *Inf. Comput.* **289**(Part), 104948 (2022)
4. Andreoli, J.M.: Logic Programming with Focusing Proofs in Linear Logic. *J. Log. Comput.* **2**(3), 297–347 (1992)
5. Balzer, S., Pfenning, F.: Manifest sharing with session types. *Proc. ACM Program. Lang.* **1**(ICFP) (Aug 2017)
6. Bellin, G., Scott, P.: On the π -calculus and linear logic. *Theoret. Comput. Sci.* **135**(1), 11–65 (1994)
7. Benton, P.N.: A mixed linear and non-linear logic: Proofs, terms and models. In: International Workshop on Computer Science Logic. pp. 121–135. Springer (1994)
8. Bernardy, J., Boespflug, M., Newton, R.R., Jones, S.P., Spiwack, A.: Linear haskell: practical linearity in a higher-order polymorphic language. *Proc. ACM Program. Lang.* **2**(POPL), 5:1–5:29 (2018)
9. Blackshear, S., Cheng, E., Dill, D.L., Gao, V., Maurer, B., Nowacki, T., Pott, A., Qadeer, S., Russi, D., Sezer, S., Zakian, T., Zhou, R.: Move: A Language with Programmable Resources. Tech. rep. (2019)
10. Caires, L., Pérez, J.A.: Linearity, control effects, and behavioral types. In: Yang, H. (ed.) Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017. Lecture Notes in Computer Science, vol. 10201, pp. 229–259. Springer (2017)
11. Caires, L., Pérez, J.A.: Linearity, control effects, and behavioral types. In: Proceedings of the 26th European Symposium on Programming Languages and Systems - Volume 10201. p. 229–259. Springer-Verlag, Berlin, Heidelberg (2017)
12. Caires, L., Pérez, J.A., Pfenning, F., Toninho, B.: Behavioral polymorphism and parametricity in session-based communication. In: Proceedings of the 22nd European Conference on Programming Languages and Systems. p. 330–349. ESOP’13, Springer-Verlag, Berlin, Heidelberg (2013)
13. Caires, L., Pfenning, F.: Session types as intuitionistic linear propositions. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010 - Concurrency Theory. pp. 222–236. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
14. Caires, L., Pfenning, F., Toninho, B.: Towards concurrent type theory. In: Proceedings of the 8th ACM SIGPLAN Workshop on Types in Language Design and Implementation. p. 1–12. TLDI ’12, Association for Computing Machinery, New York, NY, USA (2012)
15. Caires, L., Pfenning, F., Toninho, B.: Linear logic propositions as session types. *Mathematical Structures in Computer Science* **26**(3), 367–423 (2016)
16. Cardelli, L.: Typeful Programming. IFIP State-of-the-Art Reports: Formal Description of Programming Concepts pp. 431–507 (1991)
17. Castellan, S., Yoshida, N.: Two sides of the same coin: session types and game semantics: a synchronous side and an asynchronous side. *Proc. ACM Program. Lang.* **3**(POPL), 27:1–27:29 (2019)
18. Chen, R., Balzer, S., Toninho, B.: Ferrite: A judgmental embedding of session types in rust. In: Ali, K., Vitek, J. (eds.) 36th European Conference on Object-Oriented Programming, ECOOP 2022, June 6–10, 2022, Berlin, Germany. LIPIcs, vol. 222, pp. 22:1–22:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022)

19. Cousineau, G., Curien, P., Mauny, M.: The categorical abstract machine. *Sci. Comput. Program.* **8**(2), 173–202 (1987)
20. Danvy, O.: A rational deconstruction of landin’s SECD machine. In: Grellk, C., Huch, F., Michaelson, G., Trinder, P.W. (eds.) *Implementation and Application of Functional Languages, 16th International Workshop, IFL 2004, Lübeck, Germany, September 8-10, 2004, Revised Selected Papers. Lecture Notes in Computer Science*, vol. 3474, pp. 52–71. Springer (2004)
21. Dardha, O., Gay, S.J.: A new linear logic for deadlock-free session-typed processes. In: Baier, C., Lago, U.D. (eds.) *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings. Lecture Notes in Computer Science*, vol. 10803, pp. 91–109. Springer (2018)
22. Das, A., Pfenning, F.: Rast: A language for resource-aware session types. *Log. Methods Comput. Sci.* **18**(1) (2022)
23. DeYoung, H., Caires, L., Pfenning, F., Toninho, B.: Cut reduction in linear logic as asynchronous session-typed communication. In: *Computer Science Logic (2012)*
24. Franco, J., Vasconcelos, V.T.: A concurrent programming language with refined session types. In: Counsell, S., Núñez, M. (eds.) *Software Engineering and Formal Methods - SEFM 2013. Lecture Notes in Computer Science*, vol. 8368, pp. 15–28. Springer (2013)
25. Frumin, D., D’Osualdo, E., van den Heuvel, B., Pérez, J.A.: A bunch of sessions: a propositions-as-sessions interpretation of bunched implications in channel-based concurrency. *Proc. ACM Program. Lang.* **6**(OOPSLA2), 841–869 (2022)
26. Gay, S., Hole, M.: Subtyping for Session Types in the Pi Calculus. *Acta Informatica* **42**(2-3), 191–225 (2005)
27. Gay, S.J., Vasconcelos, V.T.: Linear type theory for asynchronous session types. *Journal of Functional Programming* **20**(1), 19–50 (2010)
28. Girard, J.: A new constructive logic: Classical logic. *Math. Struct. Comput. Sci.* **1**(3), 255–296 (1991)
29. Honda, K.: Types for dyadic interaction. In: Best, E. (ed.) *CONCUR’93*. pp. 509–523. Springer Berlin Heidelberg, Berlin, Heidelberg (1993)
30. Honda, K., Laurent, O.: An exact correspondence between a typed pi-calculus and polarised proof-nets. *Theor. Comput. Sci.* **411**(22-24), 2223–2238 (2010)
31. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) *Programming Languages and Systems*. pp. 122–138. Springer Berlin Heidelberg, Berlin, Heidelberg (1998)
32. Hüttel, H., Lanese, I., Vasconcelos, V.T., et. al., L.C.: Foundations of Session Types and Behavioural Contracts. *ACM Comput. Surv.* **49**(1), 3 (2016)
33. Imai, K., Neykova, R., Yoshida, N., Yuen, S.: Multiparty session programming with global protocol combinators. In: Hirschfeld, R., Pape, T. (eds.) *34th European Conference on Object-Oriented Programming, ECOOP 2020. LIPIcs*, vol. 166, pp. 9:1–9:30. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
34. Jacobs, J., Balzer, S.: Higher-order leak and deadlock free locks. *Proc. ACM Program. Lang.* **7**(POPL), 1027–1057 (2023)
35. Klabnik, S., Nichols, C.: *The Rust Programming Language* (2021)
36. Kokke, W., Dardha, O.: Deadlock-free session types in linear haskell. In: Hage, J. (ed.) *Haskell 2021: Proceedings of the 14th ACM SIGPLAN International Symposium on Haskell, Virtual Event, Korea, August 26-27, 2021*. pp. 1–13. ACM (2021)

37. Kokke, W., Montesi, F., Peressotti, M.: Better late than never: a fully-abstract semantics for classical processes. *Proc. ACM Program. Lang.* **3**(POPL), 24:1–24:29 (2019)
38. Krivine, J.: A call-by-name lambda-calculus machine. *High. Order Symb. Comput.* **20**(3), 199–207 (2007)
39. Lafont, Y.: The linear abstract machine. *Theor. Comput. Sci.* **59**, 157–180 (1988)
40. Lagaillardie, N., Neykova, R., Yoshida, N.: Implementing Multiparty Session Types in Rust. In: *Coordination Models and Languages Coordination 2020. Lecture Notes in Computer Science*, vol. 12134, pp. 127–136. Springer (2020)
41. Landin, P.J.: The Mechanical Evaluation of Expressions. *The Computer Journal*, Volume 6, Issue 4, January 1964 **6**(4), 308–320 (1964)
42. Laurent, O.: Polarized proof-nets: Proof-nets for LC. In: Girard, J. (ed.) *Typed Lambda Calculi and Applications, 4th International Conference, TLCA'99, L'Aquila, Italy, April 7-9, 1999, Proceedings. Lecture Notes in Computer Science*, vol. 1581, pp. 213–227. Springer (1999)
43. Lindley, S., Morris, J.G.: Embedding session types in haskell. In: Mainland, G. (ed.) *Proceedings of the 9th International Symposium on Haskell, Haskell 2016, Nara, Japan, September 22-23, 2016*. pp. 133–145. ACM (2016)
44. Lopes, L.M.B., Silva, F.M.A., Vasconcelos, V.T.: A virtual machine for a process calculus. In: Nadathur, G. (ed.) *Principles and Practice of Declarative Programming, International Conference PPDP'99. Lecture Notes in Computer Science*, vol. 1702, pp. 244–260. Springer (1999)
45. Milner, R.: Functions as processes. *Math. Struct. Comput. Sci.* **2**(2), 119–141 (1992)
46. Milner, R.: Elements of interaction: Turing award lecture. *Communications of the ACM* **36**(1), 78–89 (1993)
47. Milner, R.: *Communicating and mobile systems - the Pi-calculus*. Cambridge University Press (1999)
48. Neykova, R., Hu, R., Yoshida, N., Abdeljallal, F.: A session type provider: compile-time API generation of distributed protocols with refinements in $\#$. In: Dubach, C., Xue, J. (eds.) *Proceedings of the 27th International Conference on Compiler Construction, CC 2018, February 24-25, 2018, Vienna, Austria*. pp. 128–138. ACM (2018)
49. Padovani, L.: A simple library implementation of binary sessions. *J. Funct. Program.* **27**, e4 (2017)
50. 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)
51. Pfenning, F.: Structural cut elimination. In: *Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science*. p. 156. LICS '95, IEEE Computer Society, USA (1995)
52. Pfenning, F., Griffith, D.: Polarized substructural session types. In: Pitts, A. (ed.) *Foundations of Software Science and Computation Structures*. pp. 3–22. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
53. Pfenning, F., Griffith, D.: Polarized Substructural Session Types. In: *Proc. of FoS-SaCS 2015. LNCS*, vol. 9034, pp. 3–22. Springer (2015)
54. Pfenning, F., Pruiksmá, K.: Relating message passing and shared memory, proof-theoretically. In: Jongmans, S., Lopes, A. (eds.) *Coordination Models and Languages - 25th IFIP WG 6.1 International Conference, COORDINATION 2023. Lecture Notes in Computer Science*, vol. 13908, pp. 3–27. Springer (2023)

55. Pierce, B.C., Turner, D.N.: Pict: a programming language based on the pi-calculus. In: Plotkin, G.D., Stirling, C., Tofte, M. (eds.) *Proof, Language, and Interaction, Essays in Honour of Robin Milner*. pp. 455–494. The MIT Press (2000)
56. Poças, D., Costa, D., Mordido, A., Vasconcelos, V.T.: System f^μ ω with context-free session types. In: Wies, T. (ed.) *Programming Languages and Systems - 32nd European Symposium on Programming, ESOP 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings*. Lecture Notes in Computer Science, vol. 13990, pp. 392–420. Springer (2023)
57. Qian, Z., Kavvos, G., Birkedal, L.: Client-server sessions in linear logic. *Proceedings of the ACM on Programming Languages* **5**(ICFP), 1–31 (2021)
58. Rocha, P., Caires, L.: Propositions-as-types and shared state. *Proceedings of the ACM on Programming Languages* **5**(ICFP), 1–30 (2021)
59. Rocha, P., Caires, L.: Safe session-based concurrency with shared linear state. In: Wies, T. (ed.) *Programming Languages and Systems - 32nd European Symposium on Programming, ESOP 2023*. LNCS, vol. 13990, pp. 421–450. Springer (2023)
60. Rocha, P., Caires, L.: Safe session-based concurrency with shared linear state (artifact) (January 2023). <https://doi.org/10.5281/zenodo.7506064>
61. Sangiorgi, D., Walker, D.: *PI-Calculus: A Theory of Mobile Processes*. Cambridge University Press, USA (2001)
62. Toninho, B., Caires, L., Pfenning, F.: Functions as Session-Typed Processes. In: *FoSSaCS'12*. No. 7213 in LNCS (2012)
63. Toninho, B., Caires, L., Pfenning, F.: Higher-order processes, functions, and sessions: A monadic integration. In: Felleisen, M., Gardner, P. (eds.) *Programming Languages and Systems*. pp. 350–369. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
64. Toninho, B., Caires, L., Pfenning, F.: A decade of dependent session types. In: Veltri, N., Benton, N., Ghilezan, S. (eds.) *PPDP 2021: 23rd International Symposium on Principles and Practice of Declarative Programming*. pp. 3:1–3:3. ACM (2021)
65. Toninho, B., Yoshida, N.: On polymorphic sessions and functions: A tale of two (fully abstract) encodings. *ACM Trans. Program. Lang. Syst.* **43**(2) (Jun 2021)
66. Turner, D.N.: *The polymorphic Pi-calculus : theory and implementation*. Ph.D. thesis, University of Edinburgh, UK (1996)
67. Vasconcelos, V.T.: Lambda and pi calculi, CAM and SECD machines. *J. Funct. Program.* **15**(1), 101–127 (2005)
68. Wadler, P.: Propositions as sessions. In: *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming*. p. 273–286. ICFP '12, Association for Computing Machinery, New York, NY, USA (2012)
69. Wadler, P.: Propositions as Sessions. *Journal of Functional Programming* **24**(2-3), 384–418 (2014)
70. Wadler, P.: Propositions as types. *Commun. ACM* **58**(12), 75–84 (2015)
71. Willsey, M., Prabhu, R., Pfenning, F.: Design and implementation of concurrent C0. In: Cervesato, I., Fernández, M. (eds.) *Proceedings Fourth International Workshop on Linearity, LINEARITY 2016*. EPTCS, vol. 238, pp. 73–82 (2016)

A Appendix - Supplementary Material

A.1 Proofs of Section 4.1: Preservation and Progress for CLLB

Lemma 4.1 (Non-full). *For $P \neq 0$ the rule below is admissible and invertible:*

$$\frac{P \vdash^{\mathbf{B}} \Delta_P, x:A; \Gamma \quad Q \vdash^{\mathbf{B}} \Delta_Q, y:B; \Gamma \quad q = \overline{c_k} \quad B = \mathbf{E}_k; \overline{A} \quad \Gamma; \Delta_i \vdash c_i: E_i \quad -B}{\text{cut} \{P \mid \overline{x}:A \ [q] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta_P, \Delta_Q, \Delta_1, \dots, \Delta_k; \Gamma}$$

Proof. To check that the rule is admissible, directly derive the conclusion using k applications of [TCut \otimes] or [TCut \oplus]. To check inversion, we proceed by induction in the derivation of $\text{cut} \{P \mid \overline{x}:A \ [\overline{c_k}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$.

(Case [TCutB]) We have $k = 0$ so $q = \text{nil}$, A positive and B negative.

(Case [TCut-1]) Not applicable, since $c_k \neq \checkmark$.

(Case [TCut \otimes]) We have $\text{cut} \{P \mid \overline{x}:A \ [\overline{c_{k-1}}] \ @\text{clos}(z_k; R_k) \mid y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$ derived from $\text{cut} \{\text{send } x(z_k.R_k); P \mid \overline{x}:T_k \otimes A \ [\overline{c_{k-1}}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$. By i.h. $\text{send } x(z_k.R_k); P \vdash^{\mathbf{B}} \Sigma_k, x:T_k \otimes A$ and $Q \vdash \Delta_Q, y:B$ where $B = \mathbf{E}_{k-1}; (\overline{T_k} \otimes \overline{A})$ and $\Delta_i \vdash^{\mathbf{B}} c_i : E_i$ for $i \leq k-1$. By inversion ([T \otimes]) $P \vdash^{\mathbf{B}} \Delta_P, x:A, \Sigma_k = \Delta_k, \Delta_P$ and $R_k \vdash^{\mathbf{B}} \Delta_k, z_k:T_k$, so $\Delta_k \vdash^{\mathbf{B}} c_k:E_k$ where $E_k = \overline{T_k} \wp \square$. We conclude $P \vdash^{\mathbf{B}} \Delta_P, x:A, Q \vdash \Delta_Q, y:B, B = \mathbf{E}_k; \overline{A}$ and $\Delta_i \vdash c_i:E_i$ for $i \leq k$.

(Case [TCut \oplus]) We have $\text{cut} \{P \mid \overline{x}:A_{\#1} \ [\overline{c_{k-1}}] \ @\#1 \mid y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$ derived from $\text{cut} \{\#1 \ x; P \mid \overline{x}:\oplus_{\ell \in L} A_\ell \ [\overline{c_{k-1}}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta, \Delta'; \Gamma$. We proceed by the i.h. as with [T \otimes] above, considering $B = \mathbf{E}_{k-1}; (\oplus_{\ell \in L} \overline{A}_\ell)$ and $\Delta_k \vdash^{\mathbf{B}} c_k : E_k$ where $E_k = \&_{\ell \in L} E_\ell$ and $E_{\#1} = \square$, so that $B = \mathbf{E}_k; \overline{A}$ and $\Delta_i \vdash c_i : E_i$ for $i \leq k$.

Lemma 4.2 (Full). *The proof rules below are admissible:*

$$\frac{Q \vdash^{\mathbf{B}} \Delta_Q, y:B; \Gamma \quad \Gamma; \Delta_i \vdash c_i:E_i \quad B = \mathbf{E}_k; \perp \quad c_k = \checkmark \quad -B}{\text{cut} \{0 \mid \overline{x}:\emptyset \ [\overline{c_k}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta_Q, \Delta_1, \dots, \Delta_k; \Gamma}$$

$$\frac{Q \vdash^{\mathbf{B}} \Delta_Q, y:B; \Gamma \quad \Gamma; \Delta_i \vdash c_i:E_i \quad B = \mathbf{E}_{k-1}; C \quad \Gamma \vdash c_k = \text{clos!}(z, R):C \quad -B}{\text{cut} \{0 \mid \overline{x}:\emptyset \ [\overline{c_k}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta_Q, \Delta_1, \dots, \Delta_k, \Gamma}$$

Moreover, one of them must apply for inverting the judgment in the conclusion.

Proof. We have $\text{cut} \{0 \mid \overline{x}:\emptyset \ [\overline{c_k}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$ derived from [TCut1] (a) or [TCut!] (b). In case (a) $\text{cut} \{\text{close } x \mid \overline{x}:1 \ [\overline{c_k}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$ with $c_k = \checkmark$. We conclude by Lemma 4.1. In case (b) $\text{cut} \{!x(z); R \mid \overline{x}!:A \ [\overline{c_k}] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$ with $c_k = \text{clos!}(z, R)$. By Lemma 4.1, $!x(z); R \vdash z!:A; \Gamma$, so $\Gamma; \vdash c_k = \text{clos!}(z, R):?\overline{A}$ and we conclude with $B = \mathbf{E}_{k-1}; C$ and $C = ?\overline{A}$.

Lemma 4.3 (Non-empty). *Let $\text{cut} \{P \mid \overline{x}:A \ [q] \ y:B \mid Q\} \vdash^{\mathbf{B}} \Delta; \Gamma$.*

If A is negative or void, then $q \neq \text{nil}$.

Proof. If A is void then immediate by Lemma 4.2. If A is negative, suppose $q = \text{nil}$. By Lemma 4.1 $k = 0$ and $\overline{A} = B$. But then B positive, contradiction. So $q \neq \text{nil}$.

Theorem 4.1 (Preservation). *Let $P \vdash^{\mathbf{B}} \Delta; \Gamma$.*

(1) If $P \equiv^{\mathbf{B}} Q$, then $Q \vdash^{\mathbf{B}} \Delta; \Gamma$. (2) If $P \rightarrow^{\mathbf{B}} Q$, then $Q \vdash^{\mathbf{B}} \Delta; \Gamma$.

Proof. We verify that rules for $\equiv^{\mathbf{B}}$ (Fig. 7) (resp. $\rightarrow^{\mathbf{B}}$ (Fig. 8)) are type preserving.

(Case [fwdp])

(Case [fwdp]) $P = \text{cut} \{Q' \mid \bar{z}:A [q_1] x:\bar{B} \mid \text{cut} \{\text{fwd } x y \mid \bar{y}:B [q_2] w:C \mid P'\} \} \vdash^{\mathbf{B}} \Delta; \Gamma$. If $q_2 = \text{nil}$ and $B = \bar{C}$, hence $Q = \text{cut} \{Q' \mid \bar{z}:A [q_1] w:C \mid P'\} \vdash^{\mathbf{B}} \Delta; \Gamma$.

Otherwise $q_2 \neq \text{nil}$. Let $F_2 = \text{cut} \{\text{fwd } x y \mid \bar{y}:B [q_2] w:C \mid P'\}$ where $F_2 \vdash^{\mathbf{B}} \Delta_2, x:\bar{B}$ and $\Delta = \Delta_1, \Delta_2$, and $q_2 = \bar{c}_k$. By Lemma 4.1, $\text{fwd } x y \vdash^{\mathbf{B}} x:\bar{B}, y:B$, and $C = \mathbf{E}_k; \bar{B}, \Gamma; \Delta_i \vdash c_i:E_i, P \vdash w:C, \Delta_P; \Gamma$ and $\Delta_2 = \Delta_1, \dots, \Delta_k, \Delta_P$.

Let $F_1 = \text{cut} \{Q' \mid \bar{z}:A [q_1] x:\bar{B} \mid F_2\}$ with $F_1 \vdash^{\mathbf{B}} \Delta_1, x:B; \Gamma$ and $q_1 = \bar{d}_l$. By Lemma 4.1, $Q' \vdash^{\mathbf{B}} \Delta_Q, z:A; \Gamma, \Gamma; \Delta'_j \vdash d_j : F_j$ and $\bar{B} = \mathbf{F}_l; \bar{A}$ and $\Delta_1 = \Delta'_1, \dots, \Delta'_l, \Delta_Q$. By Lemma 4.1 we get $\text{cut} \{Q' \mid \bar{z}:A [q_2 @ q_1] w:C \mid P'\} \vdash^{\mathbf{B}} \Delta; \Gamma$.

(Case of $[\otimes]$) Let $\text{cut} \{\text{send } x(z.P); Q \mid \bar{x} : A \otimes C [q] y : B \mid R\} \vdash^{\mathbf{B}} \Delta; \Gamma$. By Lemma 4.1, $\text{send } x(z.P); Q \vdash^{\mathbf{B}} \Delta_P, x : A$ and $R \vdash \Delta_R, y : B$ and $\Delta_i \vdash c_i : E_i$ for $1 \leq i \leq k$, and $B = \mathbf{E}_k; (A \otimes C)$. By inversion $[\text{T}\otimes]$, $\Delta_P = \Delta'_P, \Delta''_P$ and $Q \vdash^{\mathbf{B}} \Delta''_P, x : C; \Gamma$ and $P \vdash^{\mathbf{B}} \Delta'_P, z : A, \Gamma$. So $\Gamma; \Delta'_P \vdash c_{k+1} = \text{clos}(z, P) : \bar{A} \wp \bar{C}$. Let $E_{k+1} = \bar{A} \wp \square$ so $B = \mathbf{E}_{k+1}; \bar{C}$. By Lemma 4.1, $\text{cut} \{Q \mid \bar{x} : C [q @ \text{clos}(z, P)] \mid y : B \mid R\} \vdash^{\mathbf{B}} \Delta; \Gamma$.

(Case $[\wp]$) Let $\text{cut} \{P \mid \bar{x} : A [\bar{c}_k] y : B \mid \text{recv } y(w); Q\} \vdash^{\mathbf{B}} \Delta$ where $\bar{c}_k = \text{clos}(z, R) @ q$. By Lemma 4.1, $P \vdash \Delta_P, x : A$ and $\text{recv } y(w); Q \vdash^{\mathbf{B}} \Delta_Q, y : B$ and $\Delta_1 \vdash \text{clos}(z, R) : T_1 \wp \square$ and $\Delta_i \vdash^{\mathbf{B}} c_i : E_i$ for $2 \leq i \leq k$, and $B = \bar{T}_1 \wp E_2[..E_k[\bar{A}]]$. We then have $R \vdash^{\mathbf{B}} \Delta_1, z : T_1$. By inversion on $[\text{T}\wp]$, $Q \vdash \Delta_Q, w : \bar{T}_1, y : C$ where $C = E_2[..E_k[\bar{A}]]$. We consider the case where T_1 and A are positive, and check that $\text{cut} \{P \mid \bar{x} : A [q] y : C \mid (\text{cut} \{R_1 \mid \bar{z}_1 : T_1 [\text{nil}] w : \bar{T}_1 \mid Q\})\} \vdash^{\mathbf{B}} \Delta$ by $[\text{TCutB}]$ on R_1 and Q and Lemma 4.1. Other cases are handled in a similar way.

Lemma A.1 (Barbs inversion). *Let $P \vdash \Delta; \Gamma$ and $P \downarrow_x$.*

1. *If $x \in \Gamma$ then $P \equiv \text{call } x(y); Q \mid * \mid R$, Otherwise $x : A \in \Delta$ and*
2. *$P \equiv \text{fwd } x y \mid * \mid Q$, or*
3. *If $A = \mathbf{1}$ then $P \equiv \text{close } x$.*
4. *If $A = B \otimes C$ then $P \equiv \text{send } x(y.Q); R$.*
5. *If $A = \perp$ then $P \equiv \text{wait } x; Q \mid * \mid R$.*
6. *If $A = B \wp C$ then $P \equiv \text{recv } x(y); Q \mid * \mid R$.*
7. *If $A = ?B$ then $P \equiv ?x; Q \mid * \mid R$.*
8. *If $A = !B$ then $P \equiv !x(y); Q \mid * \mid R$.*

Lemma 4.4 (Liveness). *Let $P \vdash^{\mathbf{B}} \Delta; \Gamma$ be live. Either $P \downarrow_x$ or $P \rightarrow$.*

Proof. By induction on the derivation for $P \vdash \Delta; \Gamma$, and case analysis on the last typing rule.

(Case of $[\text{TcutB}]$) we have $P = \text{cut} \{P_1 \mid y : \bar{A} [\text{nil}] \bar{z} : A \mid P_2\} \vdash^{\mathbf{B}} \Delta', \Delta; \Gamma$, derived from $P_1 \vdash^{\mathbf{B}} \Delta', y : \bar{A}; \Gamma$ and $P_2 \vdash^{\mathbf{B}} \Delta, z : A; \Gamma$, where A is positive. By the i.h. we conclude that $P_2 \rightarrow$ or $P_2 \downarrow_{x_2}$. So if $P_2 \rightarrow$, then $P \rightarrow$. Otherwise, $P_2 \downarrow_{x_2}$. If $x_2 \neq z$ then $P \downarrow_{x_2}$ with $x_2 \in \Delta$.

Otherwise $x_2 = z$. By Lemma A.1 (1,2,3), either $P_2 \equiv \text{fwd } z \ v \ | * \ | \ R$ (a), or since A is positive, $P_2 \equiv \text{close } z$ (b) or $P_2 \equiv \text{send } z(w.Q); R$ (c).

In case (a) if $v \notin \Delta$, then P reduces by [fwdp] (both x_2 and v are bound in P by cuts), otherwise $v \in \Delta$ and $P_2 \downarrow_v$ and $P \downarrow_v$. In case (b) $P \rightarrow$ by [1] and in case (c) by [⊗].

(Case of [TCut!]) Let $\text{cut! } \{y.R \ | x : A \ | Q\} \vdash^B \Delta; \Gamma$ derived from $R \vdash^B y : A; \Gamma$ and $Q \vdash^B \Delta; \Gamma, x : \bar{A}$. By i.h. either $Q \rightarrow$ or $Q \downarrow_z$. If $z \neq x$ then $P \rightarrow$ or $P \downarrow_z$. If $z = x$ then by Lemma A.1 $Q \equiv \text{call } x(y); Q \ | * \ | \ R$ and $P \rightarrow$ by [call].

(Case of [TCut⊗]) We have $P = \text{cut } \{P_1 \ | \bar{x} : A \ [q] \ y : B \ | P_2\} \vdash^B \Delta; \Gamma$ where $q = r@\text{clos}(z_k; R_k)$ derived from $P_1 \vdash^B x : A, \Delta_{P_1}; \Gamma$. By i.h. $P_1 \rightarrow$ or $P_1 \downarrow_{x_1}$. If $x \neq x_1$ then $P_1 \downarrow_{x_1}$ and $P \downarrow_{x_1}$. If A is positive, by the same reasoning as above for P_2 we conclude that $P \rightarrow$ or $P \downarrow_w$.

Otherwise, A is negative. By i.h. $P_2 \rightarrow$ or $P_2 \downarrow_z$. If $z \neq y$ then $P \downarrow_x$. Otherwise, $P_2 \downarrow_y$. By Lemma 4.1, $B = \bar{T} \wp C$. By Lemma A.1 (2,6), either $P_2 \equiv \text{fwd } y \ v \ | * \ | \ Q'$ (a) or $P_2 \equiv \text{recv } y(w); Q' \ | * \ | \ Q''$ (b). For case (a) we conclude as in [TCutB] above that $P \rightarrow$ or $P \downarrow_v$, for case (b) $P \rightarrow$ by [⊗].

(Case of [TCut⊕]) Similar to [TCut⊗]

(Case of [TCut1]) Let $P = \text{cut } \{0 \ | \bar{x} : \emptyset \ [q] \ y : B \ | P_2\} \vdash^B \Delta; \Gamma$ where $q = r@\checkmark$ derived from $\text{cut } \{\text{close } x \ | \bar{x} : \mathbf{1} \ [r] \ y : B \ | P_2\} \vdash^B \Delta; \Gamma$. By i.h. $P_2 \rightarrow$ or $P_2 \downarrow_z$. If $z \neq y$ then $P \downarrow_x$. Otherwise, $P_2 \downarrow_y$. By Lemma 4.2, either $B = \bar{T} \wp C$ (a) or $B = \perp$ (b). If $P_2 \equiv \text{fwd } y \ v \ | * \ | \ Q'$, then as for [TCutB] we conclude that $P \rightarrow$ or $P \downarrow_v$. For (a), by Lemma A.1 (6), $P_2 \equiv \text{recv } y(w); Q' \ | * \ | \ Q''$, and $P \rightarrow$ by [⊗]. For (b), by Lemma A.1 (5), $P_2 \equiv \text{wait } y; Q' \ | * \ | \ Q''$, and $P \rightarrow$ by [⊥].

A.2 Proofs of Section 4.2: Correspondence between CLL and CLLB

We give some examples of commutations used in the proofs, those for other pairs $\rightarrow^{\text{Bp}} / \rightarrow^{\text{Bn}}$ are checked in a similar way.

Lemma A.2. *Commutation (pos-neg-promote):*

$$\begin{aligned} & \text{cut } \{\text{send } x(z.R); P \ | \bar{x} \ [c@q] \ y \ | \text{recv } y(w); Q\} \rightarrow^{\text{Bp}} \\ & \text{cut } \{P \ | \bar{x} \ [c@q@v] \ y \ | \text{recv } y(w); Q\} \rightarrow^{\text{Bn}} \\ & \text{cut } \{P \ | \bar{x} \ [q@v] \ y \ | \text{cut } \{S \ | u \ [w] \ Q\}\} \\ \\ & \text{cut } \{\text{send } x(z.R); P \ | \bar{x} \ [c@q] \ y \ | \text{recv } y(w); Q\} \rightarrow^{\text{Bn}} \\ & \text{cut } \{\text{send } x(z.R); P \ | \bar{x} \ [q] \ y \ | \text{cut } \{R \ | s \ [u] \ Q\}\} \rightarrow^{\text{Bp}} \\ & \text{cut } \{P \ | \bar{x} \ [q@v] \ y \ | \text{cut } \{S \ | u \ [w] \ Q\}\} \end{aligned}$$

Lemma A.3. *Commutation (pos-neg-seq-commute):*

$$\begin{aligned} & E_1[\text{cut } \{X \ | c \ [q_1] \ \bar{a} \ | \text{send } a(z.R); \text{recv } b(w); Q\}] \rightarrow^{\text{Bp}} \\ & E_2[\text{cut } \{Y \ | \bar{d} \ [c@q_2] \ b \ | \text{recv } b(w); Q\}] \rightarrow^{\text{Bn}} E_3[\text{cut } \{Y \ | \bar{x} \ [q@v] \ b \ | \text{cut } \{S \ | u \ [w] \ Q\}\}] \\ \\ & E_1[\text{cut } \{X \ | c \ [q_1] \ \bar{a} \ | \text{send } a(z.R); \text{recv } b(w); Q\}] \equiv \\ & E_2[\text{cut } \{X \ | \bar{d} \ [c@q_2] \ b \ | \text{recv } b(w); \text{send } a(z.R); Q\}] \rightarrow^{\text{Bn}} \\ & F_3[\text{cut } \{Z \ | c \ [q_1] \ \bar{a} \ | \text{cut } \{S \ | u \ [w] \ \text{send } a(z.R); Q\}\}] \rightarrow^{\text{Bp}} \\ & E_3[\text{cut } \{Y \ | \bar{x} \ [q@v] \ b \ | \text{cut } \{S \ | u \ [w] \ Q\}\}] \end{aligned}$$

Lemma A.4. *Commutation (fwd-neg-comm):*

$$\begin{aligned}
& E[\text{cut} \{P \mid \bar{z} [q_1] x \mid \text{fwd } x y \mid \bar{y} [c@q_2] b \mid \text{recv } b(w); Q\}] \rightarrow^{\text{Ba}} \\
& E[\text{cut} \{P \mid \bar{z} [c@q_2@q_1] b \mid \text{recv } b(w); Q\}] \rightarrow^{\text{Bn}} \\
& E[\text{cut} \{P \mid \bar{z} [q_2@q_1] b \mid \text{cut} \{S \mid u[]w \mid Q\}\}] \\
& \\
& E[\text{cut} \{P \mid \bar{z} [q_1] x \mid \text{fwd } x y \mid \bar{y} [c@q_2] b \mid \text{recv } b(w); Q\}] \rightarrow^{\text{Bn}} \\
& E[\text{cut} \{P \mid \bar{z} [q_1] x \mid \text{fwd } x y \mid \bar{y} [q_2] b \mid \text{cut} \{S \mid u[]w \mid Q\}\}] \rightarrow^{\text{Ba}} \\
& E[\text{cut} \{P \mid \bar{z} [q_2@q_1] b \mid \text{cut} \{S \mid u[]w \mid Q\}\}]
\end{aligned}$$

Lemma 4.5. *The following commutations of reductions hold.*

1. Let $P_1 \rightarrow^{\text{Bp}} S \rightarrow^{\text{Bn}} P_2$. Either $P_1 \rightarrow^{\text{Br}} P_2$, or $P_1 \rightarrow^{\text{Bn}} S' \rightarrow^{\text{Bp}} P_2$ for some S' .
2. Let $P_1 \rightarrow^{\text{Ba}} S \rightarrow^{\text{Bn}} P_2$. Then $P_1 \rightarrow^{\text{Bn}} S' \rightarrow^{\text{Ba}} P_2$ for some S' .
3. If $P_1 \rightarrow^{\text{Bap}} S \rightarrow^{\text{Bn}} P_2$, either $P_1 \rightarrow^{\text{Br}} P_2$, or $P_1 \rightarrow^{\text{Bn}} S' \rightarrow^{\text{Bap}} P_2$ for some S' .
4. Let $P_1 \rightarrow^{\text{Bap}} N \xrightarrow{\epsilon}^{\text{Ba}} S \rightarrow^{\text{Br}} P_2$. Either $P_1 \xrightarrow{\epsilon}^{\text{Ba}} N$ or $P_1 \rightarrow^{\text{Br}} S' \rightarrow^{\text{Bap}} P_2$ for some S' .

Proof. (1) Either (a) the reductions are in the same cut, or (b) the reductions are in different cuts. For (a), if reductions match we conclude. Otherwise they commute (by Lemma A.2) so $P_1 \rightarrow^{\text{Bn}} S \rightarrow^{\text{Bp}} P_2$ for some S . For (b), if reductions are independent (different threads), commute, and we conclude. If the reductions are dependent (same thread), they commute (by Lemma A.3) and we conclude.

(2) We consider two cases: either (a) the reductions are in the same cut, or (b) the reductions are in different cuts. For (a), the reductions commute (by Lemma A.4), so we have $P \rightarrow^{\text{Bn}} S \rightarrow^{\text{Ba}} Q$ for some S . For (b), the reductions are independent and commute.

(3) By (1) and (2).

(4) Assume $P_1 \rightarrow^{\text{Bp}} N$. Since the reductions in $S \rightarrow^{\text{Br}} P_2$ must act on the same cut with an initially empty queue, and all reductions $N \xrightarrow{\epsilon}^{\text{Ba}} S$ are on empty cuts, the reduction $P_1 \rightarrow^{\text{Bp}} N$ must act on a different cut of all those involved and thus commutes with $N \xrightarrow{\epsilon}^{\text{Ba}} S \rightarrow^{\text{Br}} P_2$, and we conclude (b). If $P_1 \rightarrow^{\text{Ba}} S$ then either this redex generates one empty cut, so $P_1 \xrightarrow{\epsilon}^{\text{Ba}} S$, and we conclude (a), or, as in (b), it must be independent of $N \xrightarrow{\epsilon}^{\text{Ba}} S \rightarrow^{\text{Br}} P_2$.

Lemma 4.6 (Simulation). *Let $P \vdash \emptyset; \emptyset$. If $P \rightarrow Q$ then $P^\dagger \xRightarrow{\text{B}} Q^\dagger$.*

Proof. Each cut reduction of CLL is simulated by two reduction steps of B in sequence: $[\otimes \wp]$ by $[\otimes]$ followed by $[\wp]$; $[\mathbf{1} \perp]$ by $[\mathbf{1}]$ followed by $[\perp]$; $[\text{!}]$ by $[\text{!}]$ followed by $[\text{?}]$. In a closed CLL process a [fwd] reduction has the form $\text{cut} \{Q \mid x : A \mid \text{fwd } x y \mid y : A \mid P\} \rightarrow \text{cut} \{Q \mid x : A \mid \{x/y\}P\}$, which, for $A+$, reduces by [fwdB] as $\text{cut} \{Q \mid \bar{x}:A \mid [\text{nil}] \mid z:\bar{A} \mid \text{fwd } z w \mid \bar{w} \mid [\text{nil}] \mid y \mid P\} \rightarrow_a \text{cut} \{Q \mid \bar{x}:A \mid [\text{nil}] \mid y:\bar{A} \mid P\}$.

Lemma 4.7 (Simulation). *Let $P \vdash^{\text{B}} \emptyset; \emptyset$. If $P \Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}} Q$ then (1) $P \rightarrow^{\text{Bn}} R$ and $R \Rightarrow^{\text{Bap}} Q$ for some R , or; (2) $P \rightarrow^{\text{Br}} R$ and $R \Rightarrow^{\text{Bap}} Q$ for some R , or; (3) $P \xrightarrow{\epsilon}^{\text{Ba}} \rightarrow^{\text{Br}} R$ and $R \Rightarrow^{\text{Bap}} Q$ for some R .*

Proof. By induction on $P(\rightarrow^{\text{Bap}})^*P'$.

(Base) We have $P \equiv P' \rightarrow^{\text{Bn}} Q$, hence (1).

(Inductive) Case $P \rightarrow^{\text{Bap}} P' \Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}} Q$. Let (r0) $P \rightarrow^{\text{Bap}} P'$.

By i.h. for P' there is R' so that (c1) $P' \rightarrow^{\text{Bn}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$, or (c2) $P' \rightarrow^{\text{Br}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$, or (c3) $P' \xrightarrow{\text{Ba}} \rightarrow^{\text{Br}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$.

Case (c1). We have $P \rightarrow^{\text{Bap}} P' \rightarrow^{\text{Bn}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$.

By Lemma 4.5 (3) on $P \rightarrow^{\text{Bap}} P' \rightarrow^{\text{Bn}} R'$, either $P \rightarrow^{\text{Br}} R'$ and we conclude (2) ($R = R'$) or there is R such that $P \rightarrow^{\text{Bn}} R \rightarrow^{\text{Bap}} R'$ and we conclude (1).

Case (c2). We have $P \rightarrow^{\text{Bap}} P' \rightarrow^{\text{Br}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$. By Lemma 4.5 (4) either (a) $P \xrightarrow{\text{Ba}} P' \rightarrow^{\text{Br}} R'$ or (b) $P \rightarrow^{\text{Br}} R'' \rightarrow^{\text{Bap}} R'$ for some R'' . In case (a) we conclude (3) ($R = R'$), in case (b) we conclude (2) ($R = R''$).

Case (c3). We have $P \rightarrow^{\text{Bap}} P' \xrightarrow{\text{Ba}} \rightarrow^{\text{Br}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$. By Lemma 4.5 (4) either (a) $P \xrightarrow{\text{Ba}} P'$ or (b) $P \rightarrow^{\text{Br}} R'' \rightarrow^{\text{Bap}} R'$. In case (a) $P \xrightarrow{\text{Ba}} \rightarrow^{\text{Br}} R'$ we conclude (3) ($R = R'$), in case (b) we conclude (2) ($R = R''$).

Theorem 4.3 (Operational correspondence CLL-CLLB). *Let $P \vdash \emptyset; \emptyset$.*

1. *If $P \Rightarrow R$ then $P^\dagger \Rightarrow^{\text{B}} R^\dagger$.*
2. *If $P^\dagger (\Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}})^* Q$ then there is R such that $P \Rightarrow R$ and $R^\dagger \Rightarrow^{\text{Bap}} Q$.*

Proof. 1. Iterating Lemma 4.6. 2. We first check (A) if $P^\dagger \Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}} Q$ then either there is R such that $P \rightarrow R$ and $R^\dagger \Rightarrow^{\text{Bap}} Q$ or there is R such that $P \rightarrow R$ and $R^\dagger \Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}} Q$. Subproof: Assume $P^\dagger \Rightarrow^{\text{Bap}} \rightarrow^{\text{Bn}} Q$. By Lemma 4.7 either (2) $P^\dagger \rightarrow^{\text{Br}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$ for some R' , or (3) $P^\dagger \xrightarrow{\text{Ba}} \rightarrow^{\text{Br}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$ for some R' ((1) cannot apply, since in P^\dagger all queues are empty). In case (a) we have $P \rightarrow R$ and $R' \Rightarrow^{\text{Bap}} Q$ with $R' = R^\dagger$ for some R since all cuts are empty in R' . In case (b) $P^\dagger \xrightarrow{\text{Ba}} P' \rightarrow^{\text{Br}} R'$ and $R' \Rightarrow^{\text{Bap}} Q$. Then $P \Rightarrow Q$ by [fwd] and $Q \rightarrow R'$, for some Q where $Q^\dagger = P'$. So $P \Rightarrow R'$ and $R' = R^\dagger$ for some R since all cuts are empty in R' . To conclude (2), we iterate (A).

A.3 Proofs of Section 5: Correctness of the core SAM

Lemma A.5 (Sanity).

1. *Let $(P, H) \Rightarrow (P', H')$. Then if $\text{fn}(P) \subseteq \text{dom}(H)$ then $\text{fn}(P') \subseteq \text{dom}(H')$.*

Lemma A.7 (2) below expresses a useful invariant of session records in machine states encoding well-typed processes: The right (reader) endpoint is of negative type. The left (writer) endpoint is of positive type unless $P = \mathbf{0}$ or $x \in \text{fn}(P)$.

Lemma A.6 (Soundness). *Let $P \vdash^{\text{B}} \emptyset$.*

1. *If $\text{enc}(P) \Rightarrow D \xRightarrow{\text{cut}^*} C$ then there is Q such that $P \rightarrow \cup \equiv Q$ and $C = \text{enc}(Q)$.*
2. *Let $P(H)$ denote property: If $x:A\langle q, P \rangle y:B \in H$ then $B-$ and if $(x \in \text{fn}(P)$ or $P = \mathbf{0})$ then $A+$.*

Then if $\text{enc}(P) = (\mathcal{A}, H) \Rightarrow (Q, H')$ and $P(H)$ then $P(H')$.

Proof. (1) Let $enc(P) \Rightarrow D \xrightarrow{\text{cut}^*} C$. Let $enc(P) = (\mathcal{A}, H)$ for some action \mathcal{A} . Then $P \equiv E[\mathcal{A}]$ for some cut context $E[-]$. We consider each reduction \Rightarrow .

(Case of [Sfwd]) Let $enc(P) = (\text{fwd } x \ y, H[z:A\langle q_1, P_1 \rangle x][y\langle q_2, P_2 \rangle w]) \Rightarrow D$ where $D = (P_2, H[z\langle q_2 @ q_1, P_1 \rangle w]) \xrightarrow{\text{cut}^*} C$. By (2) since $P_1(z)$, then A^- . Therefore $P \equiv E[\text{fwd } x \ y] \equiv E'[\text{cut } \{P_1 \mid \bar{z} [q_1] \ x\} \mid \text{cut } \{\text{fwd } x \ y \mid \bar{y} [q_1] \ w\} \mid P_2\}]$. Then $P \rightarrow Q \equiv E'[\text{cut } \{P_1 \mid \bar{z} [q_2 @ q_1] \ w\} \mid P_2\}]$ and $enc(Q) = C$.

(Case of [S1]) Let $enc(P) = (\text{close } x, H[x\langle q, R \rangle y]) \Rightarrow (R, H[x\langle q @ \checkmark, 0 \rangle y]) = D \xrightarrow{\text{cut}^*} C$. Therefore, $P \equiv E[\text{close } x] \equiv E'[\text{cut } \{\text{close } x \mid \bar{x} [q] \ y\} \mid R]$ and $P \rightarrow Q = E'[\text{cut } \{0 \mid \bar{x} [q @ \checkmark] \ y\} \mid R]$ and $enc(Q) = C$.

(Case of [S \perp]) Let $enc(P) = (\text{wait } y; R, H[x\langle \checkmark, 0 \rangle y]) \Rightarrow D$ and $D = (R, H)$.

Hence, $P \equiv E[\text{wait } y; R] \equiv E'[\text{cut } \{0 \mid \bar{x} : \emptyset [\checkmark] \ y\} \mid \text{wait } y; R] \rightarrow Q = E'[R]$ and $D \xrightarrow{\text{cut}^*} C$ and $enc(Q) = C$.

(Case of [S $-$]) Let $enc(P) = (\mathcal{A}^-(x), \mathcal{H}[x:A\langle q, R \rangle y]) \Rightarrow D$ with

$D = (R, H[x\langle q, \mathcal{A}^-(x) \rangle y])$. By typing A^- and $x \in \text{fn}(\mathcal{A}^-(x))$. Hence (2).

Hence, $P \equiv E[\mathcal{A}^-(x)] \equiv$

$E'[\text{cut } \{\mathcal{A}^-(x) \mid \bar{x} [q] \ y\} \mid R] \equiv Q$ and $D \xrightarrow{\text{cut}^*} C$ and $enc(Q) = C$.

(Case of [S \otimes]) Let $enc(P) = (\text{send } x(z.R); U, H[x\langle q, S \rangle y]) \Rightarrow D$

where $D = (U, H[x\langle q @ \text{clos}(z, R), S \rangle y])$.

So $P \equiv E[\text{send } x(z.R); U] \equiv E'[\text{cut } \{\text{send } x(z.R); U \mid \bar{x} [q] \ y\} \mid S] \rightarrow$

$E'[\text{cut } \{U \mid \bar{x} [q @ \text{clos}(z, R)] \ y\} \mid S] \equiv Q$ and $D \xrightarrow{\text{cut}^*} C$ and $enc(Q) = C$.

(Case of [S \otimes]) Let $enc(P) = (\text{recv } y(w:-); U, H[x\langle \text{clos}(z, R) @ q, S \rangle y]) \Rightarrow D$ where $D = (R, H[z\langle \text{nil}, U \rangle w][x\langle q \rangle^s y])$ and $S(x)$ with $x : A$ and A^- or $S = 0$.

So, $P \equiv E[\text{recv } y(w:-); U] \equiv E'[\text{cut } \{\text{recv } y(w:-); U \mid \bar{x} [\text{clos}(z, R) @ q] \ y\} \mid S] \rightarrow Q$. If $q \neq \text{nil}$ then $Q = E'[\text{cut } \{S \mid \bar{x} [q] \ y\} \mid \text{cut } \{R \mid \bar{z} [\text{nil}] \ w\} \mid U\}]$. If $q = \text{nil}$, then $Q = E'[\text{cut } \{S \mid x [\text{nil}] \ \bar{y}\} \mid \text{cut } \{R \mid \bar{z} [\text{nil}] \ w\} \mid U\}]$. In both cases, $B \xrightarrow{\text{cut}^*} C$ and $enc(Q) = C$.

(Case of [S \oplus]) Let $enc(P) = (\#l; R, H[x\langle q, S \rangle y]) \Rightarrow D = (R, H[x\langle q @ \#l, S \rangle y])$.

So $P \equiv E[\#l; R] \equiv E'[\text{cut } \{\#l; R \mid \bar{x} [q] \ y\} \mid S] \rightarrow E'[\text{cut } \{R \mid \bar{x} [q @ \#l] \ y\} \mid S] \equiv Q$, and $D \xrightarrow{\text{cut}^*} C$ and $enc(Q) = C$.

(Case of [S $\&$]) Similar to [S \otimes].

Lemma 5.1 (Readiness). *Let $P \vdash \emptyset$ and $(P, \emptyset) \xrightarrow{*} S$. Then S is ready.*

Proof. The property trivially holds for (P, \emptyset) for $\text{fn}(P) = \emptyset$ and $H = \emptyset$. We proceed by transition induction, assuming that $S = (P, H)$ is ready, and $S \Rightarrow S' = (P, H')$, we check that S' is ready.

(Case of [Sfwd]) Let $R = \text{fwd } x \ y$ and $S = (R, H') \Rightarrow S' = (P, H'')$, where $H' = H[z\langle q_1, Q \rangle x][y\langle q_2, P \rangle w]$ and $H'' = H[z\langle q_2 @ q_1, Q \rangle w]$. By i.h., R is H' -ready and H' is ready. Thus $z : A$ is negative or void, $Q(z)$ is H' -ready, and $P(w)$ is H' , $\{w\}$ -ready. Then $Q(z)$ is H'' -ready, and H'' is ready. Then $P(w)$ is H'' -ready and S' is ready.

(Case of [SCut]) Let $R = \text{cut } \{P \mid \bar{x} : T [q] \ y : B \mid Q\}$ and $S = (R, H) \Rightarrow S' = (P, H')$, where $H' = H[x\langle \text{nil}, Q \rangle y]$ and T positive. By i.h., R is H -ready hence $P(x)$ is H , $\{x\}$ -ready and $Q(y)$ is H , $\{y\}$ -ready. Then $P(x)$ is H' -ready, and H' is ready. We conclude that S' is ready.

(Case of [S1]) Let $R = \text{close } x$ and $S = (R, H') \Rightarrow (Q, H'') = S'$ where $H' = H[x : A\langle q, Q \rangle y]$ and $H'' = H[x : \emptyset\langle q @ \checkmark, 0 \rangle y]$. By the i.h, H' is ready and $Q(y)$ is $H', \{y\}$ -ready. Hence $Q(y)$ is H'' -ready, since $x : \emptyset$. Clearly, H'' is ready, since $0(x)$ is H'' -ready. We conclude S' ready.

(Case of [S⊥]) Let $R = \text{wait } y; U$, $H' = H[x\langle \checkmark, 0 \rangle y]$ and $S = (R, H') \Rightarrow (U, H) = S'$. By the i.h, R is H' -ready and H' is ready. Therefore U is H -ready and H' -ready. We conclude S' ready.

(Case of [S−]) Let $R = \mathcal{A}^-(x)$ and $S = (R, H') \Rightarrow (Q, H'') = S'$ where $H' = H[x : A\langle q, Q \rangle y]$ and $H'' = H[x\langle q, \mathcal{A}^-(x) \rangle y]$ with A^- . By i.h., $\mathcal{A}^-(x)$ is H' -ready and H' is ready. Hence, $Q(y)$ is $H', \{y\}$ -ready. We have $\mathcal{A}^-(x)$ is H'' -ready (x positive endpoint). Thus H'' is ready, Q is H'' -ready (since A is negative), and S' is ready.

(Case of [S⊗]) Let $R = \text{send } x(z.V); U$ and $S = (R, H') \Rightarrow (U, H'') = S'$ where $H' = H[x\langle q, Q \rangle y]$ and $H'' = H[x\langle q @ \text{clos}(z, V), Q \rangle y]$. By i.h, H' is ready and R is H' -ready. V is $H', \{z\}$ -ready and U is H' -ready. Then V is $H'', \{z\}$ -ready and U is H'' -ready. Also $Q(y)$ is $H', \{y\}$ -ready, and $Q(y)$ is $H'', \{y\}$ -ready. We conclude H'' -ready and thus S' ready.

(Case of [S⊗]) Let $R = \text{recv } y(w : -); U$ and $S = (R, H') \Rightarrow (V, H'') = S'$ where $H' = H[x\langle \text{clos}(z, V) @ q, Q \rangle y]$ and $H'' = H[z\langle \text{nil}, U \rangle w][x\langle q \rangle^s y]$. By i.h, V is H' -ready since z is positive, and thus H'' -ready. By typing we have $Q(x)$ and thus $Q(x)$ is H' -ready. Also R is H' -ready, so U is $H', \{w\}$ -ready.

If $q \neq \text{nil}$ then $H'' = H[z\langle \text{nil}, U \rangle w][x\langle q, Q \rangle y]$. Then U is $H'', \{w\}$ -ready, $Q(x)$ is H'' -ready, H'' is ready, and S' is ready.

If $q = \text{nil}$ then $H'' = H[z\langle \text{nil}, U \rangle w][y\langle \text{nil}, Q \rangle x]$. Then U is $H'', \{w\}$ -ready, $Q(x)$ is $H'', \{x\}$ -ready (x changed polarity), H'' is ready, and S' is ready. The case $w-$ is handled in the same way.

(Case of [S⊕]) Let $R = \#l x; U$ and $S = (R, H') \Rightarrow (U, H'') = S'$ where $H' = H[x\langle q, Q \rangle y]$ and $H'' = H[x\langle q @ \#l, Q \rangle y]$. By i.h, H' is ready and R is H' -ready. Therefore U is H'' -ready. We conclude that S' is ready.

(Case of [S&]) Let $R = \text{case } y \{ \#l \in L : P_\ell \}$ and $S = (R, H') \Rightarrow (P_{\#l}, H'') = S'$ where $H' = H[x\langle \#l @ q, Q \rangle y]$ and $H'' = H[x\langle q \rangle^s y]$. By i.h, H' is ready, R is H' -ready, so $Q(x)$ is H' -ready, and $P_{\#l}$ is H' -ready.

If $q \neq \text{nil}$ then $H'' = H[x\langle q, Q \rangle y]$ is ready and $P_{\#l}$ is H'' -ready since $Q(x)$ is H' -ready, so S' is ready.

If $q = \text{nil}$ then $H'' = H[y\langle \text{nil}, Q \rangle x]$. Notice that Q is $H'', \{x\}$ -ready (x changed polarity). Then H'' is ready. $P_{\#l}$ is H'' -ready, so S' is ready.

Theorem 5.3 (Progress). *Let $P \vdash^B \emptyset$ and P live. Then $\text{enc}(P) \Rightarrow S'$.*

Proof. Since P is live then $\text{enc}(P) = S = (\mathcal{A}, H)$ for some action \mathcal{A} .

(Case of [fwdB]) $\mathcal{A} = \text{fwd } x y$. Wlog, assume that $x:A$ is negative, so $y:\bar{A}$ is positive. By typing, x and y must be endpoints of different cuts (and session records). Case $H = H'[z\langle q_1, U(z) \rangle x][y\langle q_2, V(w) \rangle w]$. We have $S \Rightarrow S'$ by [Sfwd]. Case $H = H'[x\langle q_1, U(z) \rangle z][y\langle q_2, V(w) \rangle w]$. Then $\text{fwd } x y = \mathcal{A}^-(x)$, and $S \Rightarrow S'$ by [S−].

(Case of [1]) $\mathcal{A} = \text{close } x$ so $S \Rightarrow S'$ by [S1].

(Case of $[\perp]$) $\mathcal{A} = \text{wait } y; R'$. If $H = H'[y:A\langle q, Q\rangle x]$, then $S \Rightarrow S'$ by $[\text{S}\perp]$. Otherwise, let $H = H'[x:A\langle q, \mathcal{A}\rangle y]$. By Lemma 5.1, \mathcal{A} is H' -ready, so A is negative or void. By Lemma 4.3 $q \neq \text{nil}$. By Lemma 4.2, $q = \checkmark$. So $S \Rightarrow S'$ by $[\text{S}\perp]$.

(Case of $[\otimes]$) $\mathcal{A} = \text{send } x(z.R); Q$ so $S \Rightarrow S'$ by $[\text{S}\otimes]$.

(Case of $[\wp]$) $\mathcal{A} = \text{recv } y(w); U$. If $H = H'[y:A\langle q, Q\rangle x]$, then $S \Rightarrow S'$ by $[\text{S}\perp]$. Otherwise, $H = H'[x : A\langle q, Q\rangle y]$. By Lemma 5.1, \mathcal{A} is H' -ready, so A is negative or void. By Lemma 4.3 $q \neq \text{nil}$. By Lemma 4.1 or Lemma 4.2, we must have $q = \text{clos}(z, R)@q'$. So $S \Rightarrow S'$ by $[\text{S}\wp]$.

(Case of $[\oplus]$) By $[\text{S}\oplus]$.

(Case of $[\&]$) $\mathcal{A} = \text{case } y \{ \# \ell \in L: P_\ell \}$. If $H = H'[y:A\langle q, Q\rangle x]$, then $S \Rightarrow S'$ by $[\text{S}\perp]$. Otherwise, $H = H'[x:A\langle q, Q\rangle y]$. By Lemma 5.1, \mathcal{A} is H' -ready, so A is negative or void. By Lemma 4.3 $q \neq \text{nil}$. By Lemma 4.1 or Lemma 4.2, we must have $q = \#l@q'$ so $S \Rightarrow S'$ by $[\text{S}\&]$.

A.4 Proofs and Additional Definitions of Section 6: The SAM for full CLL

To alleviate the proofs we assume that only the unrestricted names are recorded in environment, for the linear names we use names as heap locations, and α -conversion as needed to generate fresh names.

Definition A.1 (Encode). Given $P \vdash^B \emptyset; \emptyset$ we define $\text{enc}(P) = C$ as $\text{enc}(\emptyset, P, \emptyset) \stackrel{\text{cut}^*}{\Rightarrow} C$ where $\text{enc}(\mathcal{E}, P, H) \stackrel{\text{cut}^*}{\Rightarrow} C$ is defined by the rules

$$\begin{aligned} & \frac{\text{enc}(P(x), H[x:A\langle q, Q\rangle y:B]) \stackrel{\text{cut}^*}{\Rightarrow} C}{\text{enc}(\text{cut} \{P \mid \bar{x}:A[q] y:B \mid Q\}, H) \stackrel{\text{cut}^*}{\Rightarrow} C} & (A+) \\ & \frac{\text{enc}(Q(y), H[x:A\langle q, P\rangle y:B]) \stackrel{\text{cut}^*}{\Rightarrow} C}{\text{enc}(\text{cut} \{P \mid \bar{x}:A [q] y:B \mid Q\}, H) \stackrel{\text{cut}^*}{\Rightarrow} C} & (A- \text{ or } P = 0) \\ & \frac{\text{enc}(\mathcal{E} \{ \text{clos}!(y, \mathcal{E}, R)/x \}, P), H) \stackrel{\text{cut}^*}{\Rightarrow} C}{\text{enc}(\mathcal{E}, \text{cut}! \{y.R \mid !x \mid P\}, H) \stackrel{\text{cut}^*}{\Rightarrow} C} \\ & \text{enc}(\mathcal{A}, H) \stackrel{\text{cut}^*}{\Rightarrow} (\mathcal{A}, H) & (A \in \mathcal{A}) \end{aligned}$$

Lemma A.7 (Soundness-!). Let $P \vdash^B \emptyset; \emptyset$.

1. If $\text{enc}(P) \Rightarrow D \stackrel{\text{cut}^*}{\Rightarrow} C$ then there is Q such that $P \rightarrow \cup \equiv Q$ and $C = \text{enc}(Q)$.
2. Let $P(H)$ denote property: If $x:A\langle q, P\rangle y:B \in H$ then $B-$ and if $(x \in \text{fn}(P)$ or $P = 0)$ then $A-$ else $A+$.

Then if $\text{enc}(P) = (\mathcal{E}, \mathcal{A}, H) \Rightarrow (\mathcal{E}, Q, H')$ and $P(H)$ then $P(H')$.

Proof. (1) Let $\text{enc}(P) \Rightarrow D \stackrel{\text{cut}^*}{\Rightarrow} C$. Let $\text{enc}(P) = (\mathcal{A}, H)$ for some action \mathcal{A} . Then $P \equiv E[\mathcal{A}]$ for some cut context $E[-]$. We consider each reduction \Rightarrow .

(Case of $[\text{S}!]$) Let $\text{enc}(P) = (\mathcal{E}, !x(z); U, H[x\langle q, R\rangle y]) \Rightarrow D$ where $D = (R, H[x\langle q @ \text{clos}!(z, \mathcal{E}, U), 0\rangle y]) \stackrel{\text{cut}^*}{\Rightarrow} C$.

Therefore, $P \equiv E[!x(z); U] \equiv E'[\text{cut } \{!x(z); U \mid \bar{x} [q] y \mid R\}]$ and $P \rightarrow Q = E'[\text{cut } \{0 \mid \bar{x} [q@\text{clos!}(z, \mathcal{E}, U)] y \mid R\}]$ and $\text{enc}(Q) = C$.

(Case of [S?]) Let $\text{enc}(P) = (\mathcal{E}, ?y; R, H[x\langle \text{clos!}(z, \mathcal{E}, U), 0 \rangle y]) \Rightarrow D$ and $D = (\mathcal{E}'\{y/\text{clos!}(z, \mathcal{E}, U)\}, R, H) \stackrel{\text{cut}^*}{\Rightarrow} C$.

Hence, $P \equiv E[?y; R] \equiv E'[\text{cut } \{0 \mid \bar{x} : \emptyset [\text{clos!}(z, \mathcal{E}, U)] y \mid ?y; R\}] \rightarrow Q = E'[R]$ and $D \stackrel{\text{cut}^*}{\Rightarrow} C$ and $\text{enc}(Q) = C$.

(Case of [Scall]) Let $\text{enc}(P) = (\mathcal{E}, \text{call } y(w:-); U, H) \Rightarrow D$ where $D = (R, H[z\langle \text{nil}, U \rangle w]) \stackrel{\text{cut}^*}{\Rightarrow} C$. Let $\mathcal{E}(y) = \text{clos!}(z, \mathcal{E}, R)$.

So, $P \equiv E[\text{call } y(w:-); U] \equiv E'[\text{cut! } \{z.R \mid !y \mid \text{call } y(w:-); U\}] \rightarrow Q$.

Then $Q = E'[\text{cut! } \{y.P \mid !x \mid \text{cut } \{R \mid z\langle \text{nil} \rangle w \mid U\}\}]$ and $\text{enc}(Q) = C$.

Lemma A.8 (Readiness-!). *Let $P \vdash \emptyset; \emptyset$ and $(\emptyset, P, \emptyset) \stackrel{*}{\Rightarrow} S$. Then S is ready.*

Proof. The property trivially holds for $(\emptyset, P, \emptyset)$ for $\text{fn}(P) = \emptyset$ and $H = \emptyset$. We proceed by transition induction, assuming that $S = (\mathcal{E}, P, H)$ is ready, and $S \Rightarrow S' = (\mathcal{E}', P, H')$, we check that S' is ready.

(Case of [SCut!]) Let $R = \text{cut! } \{y.R \mid !x \mid Q\}$ and $S = (\mathcal{E}, Q, H) \Rightarrow S' = (\mathcal{E}'P, H)$, where $E' = \mathcal{E}\{\text{clos!}(y, \mathcal{E}, R)/x\}$. By i.h., P is H -ready but then Q is H since x is unrestricted. We conclude that S' is ready. R is also $H, \{y\}$ -ready, since by typing the only free linear name in R is y .

(Case of [S!]) Let $R = !x(z); U$ and $S = (\mathcal{E}, R, H') \Rightarrow (\mathcal{G}, Q, H'') = S'$ where $H' = H[x : !A\langle q, \mathcal{G}, Q \rangle y]$ and $H'' = H[x : \emptyset\langle q@\text{clos!}(z, \mathcal{E}, U), \emptyset, 0 \rangle y]$.

By the i.h, H' is ready and $U(z)$ is $H', \{z\}$ -ready.

Hence $Q(y)$ is H'' -ready, since $x : \emptyset$.

Clearly, H'' is ready, since $0(x)$ is H'' -ready. We conclude S' ready.

(Case of [S?]) Let $R = ?y; Q$, $H' = H[x\langle \text{clos!}(z, \mathcal{E}, U), \emptyset, 0 \rangle y]$ and $S = (\mathcal{E}, R, H') \Rightarrow (\mathcal{E}'\{\text{clos!}(z, \mathcal{E}, U)/y\}, Q, H) = S'$. By the i.h, R is H' -ready and H' is ready. Therefore Q is H -ready and H' -ready. Also U is $H', \{z\}$ -ready (no linear names). We conclude S' ready.

(Case of [Scall]) Let $R = \text{call } y(w:-); U$ and $S = (\mathcal{E}, R, H) \Rightarrow (\mathcal{F}, V, H') = S'$ where $\mathcal{E}(y) = \text{clos!}(z, \mathcal{F}, V)$ and $H' = H[z\langle \text{nil}, \mathcal{E}, U \rangle w]$.

By i.h, U is $H, \{w\}$ -ready and V is H -ready since z is positive. Hence U is $H', \{w\}$ -ready, so H' is ready, and V is H' -ready. U is $H', \{w\}$ -ready. We conclude S' ready.

Theorem A.1 (Progress-!). *Let $P \vdash \emptyset; \emptyset$ and P live. Then $\text{enc}(P) = S \Rightarrow S'$.*

Proof. Since P is live then $\text{enc}(P) = S = (\mathcal{E}, \mathcal{A}, H)$ for some action \mathcal{A} . The interesting case is:

(Case of [?]) $\mathcal{A} = ?y; U$. If $H = H'[y:A\langle q, \mathcal{F}, Q \rangle x]$, then $S \Rightarrow S'$ by [S-]. Otherwise, $H = H'[x : A\langle q, Q \rangle y]$. By Lemma A.8, \mathcal{A} is H' -ready, so A is negative or void. By Lemma 4.3 $q \neq \text{nil}$. By Lemma 4.2, we must have $q = \text{clos!}(z, \mathcal{G}, R)$. So $S \Rightarrow S'$ by [S?].

A.5 Proofs of Section 6.1: Concurrent Semantics of Cut and Mix

Here we sketch the basic structure of progress for the SAM with concurrent cut and sequential threads. First we define the encoding of CLLB processes with annotated concurrent cuts (such annotation is silent for any purpose other than the concurrent execution strategy in the SAM). For simplicity we address core SAM (no exponentials) and concurrent cuts. The proof can be easily extended to any other pair of dual types types, and already allow the concurrent non-deterministic synchronisation of different threads.

$$\begin{array}{c}
\frac{enc(P, H) \stackrel{\text{cut}^*}{\Rightarrow} (A, H') \quad enc^c(T, H') \stackrel{\text{cut}^*}{\Rightarrow} (T', H'')}{enc^c(\{P\} \uplus T, H') \stackrel{\text{cut}^*}{\Rightarrow} (\{A\} \uplus T', H'')} \quad [\text{thr}] \\
\\
\frac{enc^c(T, H) \stackrel{\text{cut}^*}{\Rightarrow} C}{enc^c(\{0\} \uplus T, H) \stackrel{\text{cut}^*}{\Rightarrow} C} \quad [0] \\
\\
\frac{enc^c(\{P, Q\} \uplus T, H) \stackrel{\text{cut}^*}{\Rightarrow} C}{enc^c(\text{par}\{P \parallel Q\} \uplus T, H) \stackrel{\text{cut}^*}{\Rightarrow} C} \quad [\text{mix}] \\
\\
\frac{enc^c(\{P, Q\} \uplus T, H[x \langle q \rangle y]) \stackrel{\text{cut}^*}{\Rightarrow} C}{enc^c(\text{pcut}\{P \mid x : A^+[q]y : B \mid Q\} \uplus T, H) \stackrel{\text{cut}^*}{\Rightarrow} C} \quad [\text{pcut}]
\end{array}$$

Fig. 12: Encoding of annotated CLLB of into concurrent SAM

$$\begin{array}{l}
(\text{close } x, H[x \langle q \rangle y]) \Rightarrow (0, H[x \langle q @ \checkmark \rangle y]) \quad [\text{S1c}] \\
(\text{wait } y; P, H[x \langle \checkmark, y \rangle]) \Rightarrow (P, H) \quad [\text{S}\perp\text{c}] \\
(\text{send } x(z.R); Q, H[x \langle q \rangle y]) \Rightarrow (Q, H[x \langle q @ \text{clos}(z, R) \rangle y]) \quad [\text{S}\otimes\text{c}] \\
(\text{recv } y(w:-); Q, H[x \langle \text{clos}(z, R) @ q \rangle y]) \Rightarrow (R, H[z \langle \text{nil}, Q \rangle w][x \langle q \rangle^s y]) \quad [\text{S}\wp\text{-c}] \\
(\text{fwd } x y, H[z \langle q_1 \rangle x[y \langle q_2 \rangle w]]) \Rightarrow (0, H[z \langle q_2 @ q_1 \rangle w]) \quad [\text{SfwdPc}] \\
a \langle q \rangle^s b \triangleq \text{if } (q = \text{nil}) \text{ then } b \langle q \rangle a \text{ else } a \langle q \rangle b
\end{array}$$

Fig. 13: Transition rules for concurrent actions (sample)

We show that SAM execution on well-typed processes simulates CLLB reduction.

Theorem 6.3 (Soundness-c). *Let $P \vdash^B \emptyset; \emptyset$.*

$$\begin{array}{c}
\frac{(P, H) \Rightarrow (P', H')}{(P \uplus T, H) \Rightarrow (P' \uplus T, H')} \text{ [Srun]} \quad \frac{(T, H) \Rightarrow (T', H')}{(0 \uplus T, H) \Rightarrow (T', H')} \text{ [S0p]} \\
(\text{pcut } \{P \mid x : A^+ \mid Q\} \uplus T, H) \Rightarrow (\{P, Q\} \uplus T, H[x \langle \text{nil} \rangle y]) \text{ [SCutp]} \\
((P \parallel Q) \uplus T, H) \Rightarrow (\{P, Q\} \uplus T, H[x \langle \text{nil} \rangle y]) \text{ [SMixp]}
\end{array}$$

Fig. 14: Transition rules for concurrent actions

If $\text{enc}^c(P) \Rightarrow D \stackrel{\text{cut}^*}{\Rightarrow} C$ then there is Q such that $P \rightarrow \cup \equiv Q$ and $C = \text{enc}^c(Q)$.

Proof. Induction of $\text{enc}(P) \stackrel{\text{cut}^*}{\Rightarrow} B$.

To prove progress along the lines of Theorem 4.2 we introduce a simplified notion of observation (cf. Fig. 9) for our scenario of concurrent actions on heap allocated concurrent sessions, as follows.

$$P \downarrow_x^H \text{ iff } P \downarrow_x \text{ and } (x \langle q \rangle y \in H \text{ or } x \langle q \rangle y \in H).$$

Lemma A.9 (Liveness-c). Let $P \vdash \Delta$ and $\text{enc}^c(P, H) \stackrel{\text{cut}^*}{\Rightarrow} (T, H')$. For every non-empty $S \subseteq T$ there is $P_i \in S$ such that (a) $(P_i, H') \Rightarrow C$ or (b) $P_i \downarrow_y^{H'}$.

Proof. By induction on the derivation of $\text{enc}(P, H) \stackrel{\text{cut}^*}{\Rightarrow} (T, H)$.

(Case of [thr]) By Theorem 5.3 (Progress), we have $(P, H) \Rightarrow (A, H')$ for any action on a sequential cut. Otherwise P is an action in a concurrent cut, and we have (b). By i.h., the property holds for T', H'' . Hence it holds for $(\{A\} \uplus T', H'')$.

(Case of [mix]) Let $\text{enc}^c(\{P, Q\} \uplus T, H[x \langle q \rangle y]) \Rightarrow (T', H')$

By i.h. (a) holds for T', H' and we conclude.

(Case of [cutp]) Let $R = \text{pcut } \{P \mid \bar{x}:A [q] y:B \mid Q\} \uplus T$ and $\text{enc}(R, H) \Rightarrow (T', H')$. Let $U_Q = T \downarrow \text{fn}(Q)$. By i.h., there is $Q_i \in U_P$ such that (a) $(Q_i, H') \Rightarrow$ or (b) $Q_i \downarrow_z^{H'}$. In case (a) we conclude (a). Otherwise assume (b). If $z \neq y$ then (a) holds for T', H' . If $z = y$ then (by typing) B is negative and $Q_i = \mathcal{A}^B(y)$ is a negative action on y of type B and $w \langle q \rangle y \in H$.

If $q \neq \text{nil}$ then we have $(Q_i, H) \Rightarrow$ by [SBc] (we denote by [SBc] the negative reduction rule associated to the negative type B).

Otherwise $q = \text{nil}$ and $A = \bar{B}$ so $P \neq 0$.

Let $U_P = T \downarrow \text{fn}(P)$. By i.h., there is $P_i \in U_P$ such that $(P_i, H) \Rightarrow$ or (b) and $P_j \downarrow_u^{H'}$. In case (a) we conclude. Otherwise assume (b). If $u \neq x$ then (a) holds for T', H' . So $u = x = w$. Now, if A is negative or void, then by Lemma 4.3 $q \neq \text{nil}$, contradiction. So A is positive and $P_j = \mathcal{A}^A(x)$ is a positive action on x of type A . So we have $(P_j, H'[x \langle q \rangle y]) \Rightarrow (0, H'[x \langle q@c \rangle y])$ by [SAc] and $(Q_i, H'[x \langle q \rangle^s y]) \Rightarrow C$ by [SBc].