

Exceptional Asynchronous Session Types

Session Types without Tiers

SIMON FOWLER, The University of Edinburgh

SAM LINDLEY, The University of Edinburgh

J. GARRETT MORRIS, University of Kansas

SÁRA DECOVA, The University of Edinburgh

Session types statically guarantee that communication complies with a protocol. However, most accounts of session typing do not account for failure, which means they are of limited use in real applications—especially distributed applications—where failure is pervasive.

We present the first formal integration of asynchronous session types with exception handling in a functional programming language. We define a core calculus which satisfies preservation and progress properties, is deadlock free, confluent, and terminating.

We provide the first implementation of session types with exception handling for a fully-fledged functional programming language, by extending the Links web programming language; our implementation draws on existing work on effect handlers. We illustrate our approach through a running example of two-factor authentication, and a larger example of a session-based chat application where communication occurs over session-typed channels and disconnections are handled gracefully.

ACM Reference Format:

Simon Fowler, Sam Lindley, J. Garrett Morris, and Sára Decova. 2019. Exceptional Asynchronous Session Types: Session Types without Tiers. *Proc. ACM Program. Lang.* POPL, 1, Article 1 (November 2019), 56 pages. https://doi.org/10.475/123_4

1 INTRODUCTION

With the growth of the internet and mobile devices, as well as the failure of Moore’s law, concurrency and distribution have become central to many applications. Writing correct concurrent and distributed code requires effective tools for reasoning about communication protocols. While data types provide an effective tool for reasoning about the shape of data communicated, protocols also require us to reason about the order in which messages are transmitted.

Session types [Honda 1993; Honda et al. 1998] are types for protocols. They describe both the shape and order of messages. If a program type-checks according to its session type, then it is statically guaranteed to comply with the corresponding protocol. Alas, most accounts of session types do not handle failure, which means they are of limited use in distributed settings where failure is pervasive. Inspired by work of Mostrous and Vasconcelos [2014], we present the first account of asynchronous session types in a functional programming language, which smoothly handles both distribution and failure. We present both a core calculus enjoying strong

Authors’ addresses: Simon Fowler, The University of Edinburgh, simon.fowler@ed.ac.uk; Sam Lindley, The University of Edinburgh, sam.lindley@ed.ac.uk; J. Garrett Morris, University of Kansas, garrett@ittc.ku.edu; Sára Decova, The University of Edinburgh, sara.decova@gmail.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/11-ART1

https://doi.org/10.475/123_4

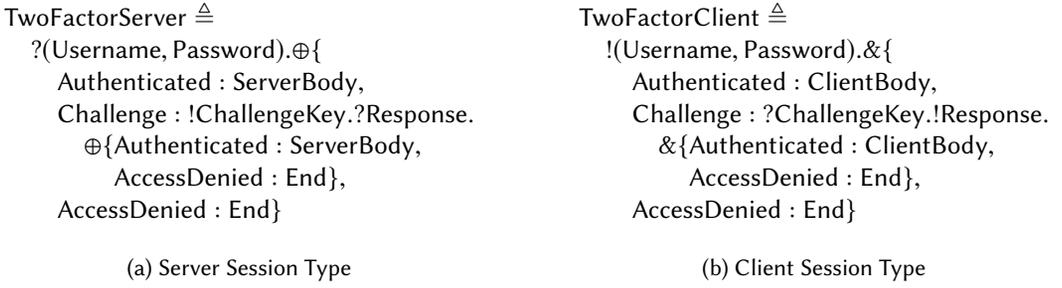


Fig. 1. Two-factor Authentication Session Types

metatheoretical correctness properties and a practical implementation as an extension of the Links web programming language [Cooper et al. 2007].

1.1 Session Types

We illustrate session types with a basic example of two-factor authentication. A user inputs their credentials. If the login attempt is from a known device, then they are authenticated and may proceed to perform privileged actions. If the login attempt is from an unrecognised device, then the user is sent a challenge code. They enter the challenge code into a hardware key which yields a response code. If the user responds with the correct response code, then they are authenticated.

A session type specifies the communication behaviour of one endpoint of a communication channel participating in a dialogue (or *session*) with the other endpoint of the channel. Fig. 1 shows the session types of two channel endpoints connecting a client and a server. Fig. 1a shows the session type for the server which first receives (?) a pair of a username and password from the client. Next, the server selects (\oplus) whether to authenticate the client, issue a challenge, or reject the credentials. If the server decides to issue a challenge, then it sends (!) the challenge string, awaits the response, and either authenticates or rejects the client. The *ServerBody* type abstracts over the remainder of the interactions, for example making a deposit or withdrawal.

Duality. The client implements the *dual* session type, shown in Fig. 1b. Whenever the server receives a value, the client sends a value, and vice versa. Whenever the server makes a selection, the client offers a choice (&), and vice versa. This *duality* between client and server ensures that each communication is matched by the other party. We denote duality with an overbar; thus $\overline{\text{TwoFactorClient}} = \text{TwoFactorServer}$ and $\overline{\text{TwoFactorServer}} = \text{TwoFactorClient}$.

Implementing Two-factor Authentication. Let us suppose we have constructs for sending and receiving along, and for closing, an endpoint.

send $M N : S$	where M has type A , and N is an endpoint with session type $!A.S$
receive $M : (A \times S)$	where M is an endpoint with session type $?A.S$
close $M : 1$	where M is an endpoint with session type End

Let us also suppose we have constructs for selecting and offering a choice:

select $\ell_j M : S_j$	where M is an endpoint with session type $\oplus\{\ell_i : S_i\}_{i \in I}$, and $j \in I$
offer $M \{\ell_i(x_i) \mapsto N_i\}_{i \in I} : A$	where M is an endpoint with session type $\&\{\ell_i : S_i\}_{i \in I}$, each x_i binds an endpoint with session type S_i , and each N_i has type A

We can now write a client implementation.

```

twoFactorClient : (Username × Password × TwoFactorClient) → 1
twoFactorClient(username, password, s) ≜
  let s = send (username, password) s in
  offer s {Authenticated(s) ↦ clientBody(s)
          Challenge(s)    ↦ let (key, s) = receive s in
                             let s = send (generateResponse(key)) s in
                             offer s {Authenticated(s) ↦ clientBody(s)
                                       AccessDenied(s) ↦ close s; loginFailed}
          AccessDenied(s) ↦ close s; loginFailed}

```

The `twoFactorClient` function takes the credentials and an endpoint of type `TwoFactorClient` as its arguments. The credentials are sent along the endpoint, then three choices are offered depending on whether the server authenticates the user, sends a two-factor challenge, or rejects the authentication attempt. If the server authenticates the user, then the program progresses to the main application (`clientBody(s)`). If the server sends a challenge, then the client receives the challenge key, and sends the response, calculated by `generateResponse`. Two choices are then offered according to whether the challenge response was successful. The rejection of an authentication attempt is part of the protocol and *not* exceptional behaviour. We can also write a server implementation.

```

twoFactorServer : TwoFactorServer → 1
twoFactorServer(s) ≜ let ((username, password), s) = receive s in
  if checkDetails(username, password) then
    let s = select Authenticated s in serverBody(s)
  else
    let s = select AccessDenied s in close s

```

The `twoFactorServer` function takes an endpoint of type `TwoFactorServer` along which it receives the credentials, which are checked using `checkDetails`. If the check passes, then the server proceeds to the application body (`serverBody(s)`); if not, then the server notifies the client by selecting the `AccessDenied` branch. This particular server implementation opts to never send a challenge request.

Statically checking session types demands a substructural type system. We discuss three options: linear types, affine types, and linear types with explicit cancellation.

1.2 Linear Types

Simply providing constructs for sending and receiving values, and for selecting and offering choices, is insufficient for safely implementing session types. Consider the following client:

```

wrongClient : TwoFactorClient → 1
wrongClient(s) ≜ let t = send ("Alice", "hunter2") s in
  let t = send ("Bob", "letmein") s in ...

```

Reuse of `s` allows a `(username, password)` pair to be sent along the same endpoint twice, violating the fundamental property of *session fidelity*, which states that in a well-typed program, communication over an endpoint matches its session type. To maintain session fidelity and ensure that all communication actions in a session type occur, session type systems typically require that each endpoint is used *linearly*—exactly once.

Exceptions. In practice, linear session types are unrealistic. Thus far, we have assumed `checkDetails` always succeeds, which may be plausible if checking against an in-memory store, but not if connecting to a remote database. One option would be for `checkDetails` to return false on

failure, but that would lose information. Instead, suppose we have an exception handling construct. As a first attempt, we might try to write:

```

148     exnServer1 : TwoFactorClient  $\rightarrow$  1
149
150     exnServer1(s)  $\triangleq$  let ((username, password), s) = receive s in
151
152         try if checkDetails(username, password) then
153             let s = select Authenticated s in serverBody(s)
154         else
155             let s = select AccessDenied s in close s
156         catch log("Database Error")
157

```

However, the above code does not type-check and is unsafe. Linear endpoint s is not used in the **catch** block and yet is still open if an exception is raised by checkDetails.

As a second attempt, we may decide to localise exception handling to the call to checkDetails. We introduce checkDetailsOpt, which returns Some(*result*) if the call is successful and None if not.

```

162     checkDetailsOpt : (Username  $\times$  Password)  $\rightarrow$  Option(Bool)
163     checkDetailsOpt(username, password)  $\triangleq$  try Some(checkDetails(username, password))
164         catch None
165

```

```

166     exnServer2 : TwoFactorServer  $\rightarrow$  1
167     exnServer2(s)  $\triangleq$  let ((username, password), s) = receive s in
168         case checkDetailsOpt(username, password) of
169             Some(res)  $\mapsto$  if res then let s = select Authenticated s in serverBody(s)
170             else let s = select AccessDenied s in close s
171             None  $\mapsto$  log("Database Error")
172

```

Still the code is unsafe as it does not use s in the None branch of the case-split. However, we do now have more precise information about the type of s , since it is unused in the **try** block. One solution could be to adapt the protocol by adding an **InternalError** branch:

```

176     TwoFactorServerExn  $\triangleq$  ?(Username, Password). $\oplus$ {
177         Authenticated : ServerBody,
178         Challenge : !ChallengeKey.Response. $\oplus$ {Authenticated : ServerBody, AccessDenied : End},
179         AccessDenied : End,
180         InternalError : End}
181

```

We could use **select** InternalError s in the None branch to yield a type-correct program, but doing so would be unsatisfactory as it clutters the protocol and the implementation with failure points.

Disconnection. The problem of failure is compounded by the possibility of disconnection. On a single machine it may be plausible to assume that communication always succeeds. In a distributed setting this assumption is unrealistic as parties may disconnect without warning. The problem is particularly acute in web applications as a client may close the browser at any point. In order to adequately handle failure we must incorporate some mechanism for detecting disconnection.

1.3 Affine Types

We began by assuming linear types—each endpoint must be used *exactly* once. One might consider relaxing linear types to *affine types*—each endpoint must be used *at most* once. Statically checked affine types form the basis of the existing Rust implementation of session types [Jespersen et al. 2015] and dynamically checked affine types form the basis of the OCaml FuSe [Padovani 2017] and Scala lchannels [Scalas and Yoshida 2016] session type libraries. Affine types present two

quandaries arising from endpoints being silently discarded. First, a developer receives no feedback if they *accidentally* forget to finish a protocol implementation. Second, if an exception is raised in an evaluation context that captures an open endpoint then the peer may be left waiting forever.

1.4 Linear Types with Explicit Cancellation

Mostrous and Vasconcelos [2014] address the difficulties outlined above through an *explicit* discard (or *cancellation*) operator. (They characterise their sessions as *affine*, but it is important not to confuse their system with affine type systems, as in §1.3, which allow variables to be discarded *implicitly*.) Their approach boils down to three key principles: endpoints can be explicitly discarded; an exception is thrown if a communication cannot succeed because a peer endpoint has been cancelled; and endpoint cancellations are propagated when endpoints become inaccessible due to an exception being thrown. They introduce a process calculus including the term $a\cancel{}$ (“cancel a ”), which indicates that endpoint a may no longer be used to perform communications. They provide an exception handling construct which attempts a communication action, running an exception handler if the action fails, and show that explicit cancellation is well-behaved: their calculus satisfies preservation and global progress (well-typed processes never get stuck), and is confluent.

Explicit cancellation neatly handles failure while ruling out accidentally incomplete implementations and providing a mechanism for notifying peers when an exception is raised. In this paper we take advantage of explicit cancellation to formalise and implement asynchronous session types with failure handling in a distributed functional programming language; this is not merely a routine adaptation of the ideas of Mostrous and Vasconcelos for the following reasons:

- They present a *process calculus*, but we work in a *functional programming language*.
- Communication in their system is *synchronous*, depending on a rendezvous between sender and receiver. We require *asynchronous* communication, which is more amenable to implementation in a distributed setting.
- Their exception handling construct is over a single communication action and does not allow nested exception handling. This design is difficult to reconcile with a functional language, as it is inherently *non-compositional*. Our exception handling construct is *compositional*.

We define a core concurrent λ -calculus, *Exceptional GV* (EGV), with asynchronous session-typed communication and exception handling. As with the calculus of Mostrous and Vasconcelos, an exception is raised when a communication action fails. But our compositional exception handling construct can be arbitrarily nested, and allows exception handling over multiple communication actions. Using EGV, we may implement the two factor authentication server as follows:

```

exnServer3 : TwoFactorServer  $\multimap$  1
exnServer3(s)  $\triangleq$  let ((username, password), s) = receive s in
  try checkDetails(username, password) as res in
    if res then let s = select Authenticated s in serverBody(s)
    else let s = select AccessDenied s in close s
  otherwise
    cancel s; log("Database Error")

```

Following Benton and Kennedy [2001], an exception handler **try** L **as** x **in** M **otherwise** N takes an explicit success continuation M as well as the usual failure continuation N . If checkDetails fails with an exception, then s is safely discarded using **cancel**, which takes an endpoint and returns the unit value. Disconnection is handled by cancelling all endpoints associated with a client. If a peer tries to read along a cancelled endpoint then an exception is thrown.

```

246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294

```

<pre> try let $s = \text{fork}(\lambda t. \text{cancel } t)$ in let $(res, s) = \text{receive } s$ in close $s; res$ as res in print ("Result: " + res) otherwise print "Error!" </pre> <p>(a) Cancellation and Exceptions</p>	<pre> let $s =$ fork $(\lambda t.$ let $(res, t) = \text{receive } t$ in close $t; res)$ in let $u = \text{fork}(\lambda v. \text{cancel } v)$ in let $u = \text{send } s$ u in close u </pre> <p>(b) Delegation</p>	<pre> let $f = (\lambda x. \text{send } x$ $s)$ in raise; $f(5)$ </pre> <p>(c) Closures</p>
--	---	---

Fig. 2. Failure Examples

We implement the constructs described by EGV as an extension to Links [Cooper et al. 2007], a functional programming language for the web. Our implementation is based on a minimal translation to effect handlers [Plotkin and Pretnar 2013].

1.5 Contributions

This paper makes five main contributions:

- (1) *Exceptional GV* (§2), a core linear lambda calculus extended with asynchronous session-typed channels and exception handling. We prove (§3) that the core calculus enjoys preservation, progress, a strong form of confluence called the *diamond property*, and termination.
- (2) Extensions to EGV supporting exception payloads, unrestricted types, and access points (§4).
- (3) The design and implementation of an extension of the Links web programming language to support tierless web applications which can communicate using session-typed channels (§5).
- (4) Client and server backends for Links implementing session typing with exception handling (§5.4), drawing on connections with effect handlers [Plotkin and Pretnar 2013].
- (5) Example applications using the infrastructure (§6). In addition to our two-factor authentication workflow we outline the implementation of a chat server.

Links is open-source and freely-available. The website can be found at <http://www.links-lang.org> and the source at <http://www.github.com/links-lang/links>. Users of the opam tool can install Links by invoking `opam install links`.

The rest of the paper is structured as follows: §2 presents Exceptional GV and §3 its metatheory; §4 discusses extensions to Exceptional GV; §5 describes the implementation; §6 presents a chat application written in Links; §7 discusses related work; and §8 concludes.

2 EXCEPTIONAL GV

In this section, we introduce Exceptional GV (henceforth EGV). GV is a core session-typed linear λ -calculus that has a tight correspondence with classical linear logic [Lindley and Morris 2015; Wadler 2014]. EGV is an asynchronous variant of GV with support for failure handling.

Due to GV's close correspondence with classical linear logic, EGV has a strong metatheory, enjoying preservation, global progress, the diamond property, and termination. Much like the simply-typed λ -calculus, this well-behaved core must be extended to be expressive enough to write larger applications. Nonetheless, the core calculus alone is expressive enough to support our two-factor authentication example, and to support server applications which gracefully handle disconnection. In §3, we show that cancellation is well-behaved, and does not violate any of the core properties of GV. In §4, following Lindley and Morris [2015, 2017], we extend EGV modularly

295	Types	$A, B, C ::= 1 \mid A \multimap B \mid A + B \mid A \times B \mid S$
296	Session Types	$S, T ::= !A.S \mid ?A.S \mid \text{End}$
297	Variables	x, y
298	Terms	$L, M, N ::= x \mid \lambda x.M \mid MN \mid () \mid \mathbf{let} () = M \mathbf{in} N \mid (M, N) \mid \mathbf{let} (x, y) = M \mathbf{in} N$
299		$\mid \mathbf{inl} M \mid \mathbf{inr} M \mid \mathbf{case} L \mathbf{of} \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \}$
300		$\mid \mathbf{fork} M \mid \mathbf{send} M N \mid \mathbf{receive} M \mid \mathbf{close} M$
301		$\mid \mathbf{cancel} M \mid \mathbf{raise} \mid \mathbf{try} L \mathbf{as} x \mathbf{in} M \mathbf{otherwise} N$
302	Type Environments	$\Gamma ::= \cdot \mid \Gamma, x : A$

Fig. 3. Syntax

with standard features of our implementation, some of which provide weaker guarantees. Channel cancellation and exceptions are orthogonal to these features.

2.1 Integrating Sessions with Exceptions, by Example

Integrating session types with failure handling into a higher-order functional language requires care. Fig. 2 illustrates three important cases: cancellation and exceptions, delegation, and closures. In order to initiate a session, we adopt the **fork** primitive of Lindley and Morris [2015]. Given a term M of type $S \multimap 1$, the term **fork** M of type \bar{S} creates a fresh channel with endpoints a of type S and b of type \bar{S} , forks a child thread that executes $M a$, and returns endpoint b .

Cancellation and Exceptions. Fig. 2a forks a thread which immediately cancels its endpoint. The parent attempts to receive, but the message can never arrive so an exception is raised and the **otherwise** clause is invoked.

Delegation. A central feature of π -calculus is *mobility* of names. In session calculi sending an endpoint is known as *session delegation*. The code in Fig. 2b begins by forking a thread and returning endpoint s . The child is passed endpoint t on which it blocks receiving. Next, the parent forks a second child, yielding endpoint u . The second child is passed endpoint v , which is immediately discarded using **cancel**. Now the parent thread sends endpoint s along u . Endpoint s will never be received as the peer endpoint v of u has been cancelled. In turn, this renders s irretrievable and an exception is thrown in the first child thread, as it can never receive a value.

Closures. It is crucial that cancellation plays nicely with closures. The code in Fig. 2c defines a function f which sends its argument x along s . The parent thread then raises an exception. As s appears in the closure bound to f , which appears in the continuation and is thus discarded, s must be cancelled.

2.2 Syntax and Typing Rules for Terms

Fig. 3 gives the syntax of EGV. Types include unit (1), linear functions ($A \multimap B$), linear sums ($A + B$), linear tensor products ($A \times B$), and session types (S).

Terms include variables (x) and the usual introduction and elimination forms for linear functions, unit, products, and sums. We write $M; N$ as syntactic sugar for **let** $() = M$ **in** N and **let** $x = M$ **in** N for $(\lambda x.N) M$. The standard session typing primitives [Lindley and Morris 2015] are as follows: **fork** M creates a fresh channel with endpoints a of type S and b of type \bar{S} , forks a child thread that executes $M a$, and returns endpoint b ; **send** $M N$ sends M along endpoint N ; **receive** M receives along endpoint M ; and **close** M closes an endpoint when a session is complete.

344	Term Typing	$\Gamma \vdash M : A$
345	T-VAR	
346	$\frac{}{x : A \vdash x : A}$	
347	T-ABS	
348	$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x. M : A \multimap B}$	
349	T-APP	
350	$\frac{\Gamma_1 \vdash M : A \multimap B \quad \Gamma_2 \vdash N : A}{\Gamma_1, \Gamma_2 \vdash M N : B}$	
351	T-LETUNIT	
352	$\frac{\Gamma_1 \vdash M : \mathbf{1} \quad \Gamma_2 \vdash N : A}{\Gamma_1, \Gamma_2 \vdash \mathbf{let} () = M \mathbf{in} N : A}$	
353	T-PAIR	
354	$\frac{\Gamma_1 \vdash M : A \quad \Gamma_2 \vdash N : B}{\Gamma_1, \Gamma_2 \vdash (M, N) : A \times B}$	
355	T-LETPAIR	
356	$\frac{\Gamma_1 \vdash M : A \times B \quad \Gamma_2, x : A, y : B \vdash N : C}{\Gamma_1, \Gamma_2 \vdash \mathbf{let} (x, y) = M \mathbf{in} N : C}$	
357	T-INL	
358	$\frac{\Gamma \vdash M : A}{\Gamma \vdash \mathbf{inl} M : A + B}$	
359	T-INR	
360	$\frac{\Gamma \vdash M : B}{\Gamma \vdash \mathbf{inr} M : A + B}$	
361	T-CASE	
362	$\frac{\Gamma_1 \vdash L : A + B \quad \Gamma_2, x : A \vdash M : C \quad \Gamma_2, y : B \vdash N : C}{\Gamma_1, \Gamma_2 \vdash \mathbf{case} L \mathbf{of} \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \} : C}$	
363	T-FORK	
364	$\frac{\Gamma \vdash M : S \multimap \mathbf{1}}{\Gamma \vdash \mathbf{fork} M : \bar{S}}$	
365	T-SEND	
366	$\frac{\Gamma_1 \vdash M : A \quad \Gamma_2 \vdash N : !A.S}{\Gamma_1, \Gamma_2 \vdash \mathbf{send} M N : S}$	
367	T-RECV	
368	$\frac{\Gamma \vdash M : ?A.S}{\Gamma \vdash \mathbf{receive} M : (A \times S)}$	
369	T-CLOSE	
370	$\frac{\Gamma \vdash M : \mathbf{End}}{\Gamma \vdash \mathbf{close} M : \mathbf{1}}$	
371	T-CANCEL	
372	$\frac{\Gamma \vdash M : S}{\Gamma \vdash \mathbf{cancel} M : \mathbf{1}}$	
373	T-TRY	
374	$\frac{\Gamma_1 \vdash L : A \quad \Gamma_2, x : A \vdash M : B \quad \Gamma_2 \vdash N : B}{\Gamma_1, \Gamma_2 \vdash \mathbf{try} L \mathbf{as} x \mathbf{in} M \mathbf{otherwise} N : B}$	
375	T-RAISE	
376	$\frac{}{\cdot \vdash \mathbf{raise} : A}$	
377	Duality	\bar{S}
378	$\overline{!A.S} = ?A.\bar{S}$	
379	$\overline{?A.S} = !A.\bar{S}$	
380	$\overline{\mathbf{End}} = \mathbf{End}$	

Fig. 4. Term Typing and Duality

We introduce three new term constructs to support session typing with failure handling: **cancel** M explicitly discards session endpoint M ; **raise** raises an exception; and **try** L **as** x **in** M **otherwise** N evaluates L , on success binding the result to x in M and on failure evaluating N .

Explicit success continuations. Benton and Kennedy [2001] argue that:

From the points of view of programming pragmatics, rewriting and operational semantics, the syntactic construct used for exception handling in ML-like programming languages, and in much theoretical work on exceptions, has subtly undesirable features.

Benton and Kennedy show that explicit success continuations avoid the subtly undesirable features they identify; correspondingly, we adopt their construct. Moreover, explicit success continuations align with the definition of handlers for algebraic effects [Plotkin and Pretnar 2013] that we use in our implementation (§5.4).

Branching and selection. Though our implementation supports **select** and **offer** directly, and we use them in examples, we omit them from the core calculus (following Lindley and Morris [2015, 2017]) as they can be encoded using sums and delegation [Dardha et al. 2017; Kobayashi 2003].

Typing. Fig. 4 gives the typing rules for EGV. As usual, linearity is enforced by splitting environments when typing subterms, ensuring T-VAR takes a singleton environment, and leaf rules T-UNIT and T-RAISE take an empty environment. We write Γ_1, Γ_2 to mean the disjoint union of Γ_1 and Γ_2 . The bulk of the rules are standard for a linear λ -calculus. Session types are related by *duality*. The T-FORK rule forks a thread connected by dual endpoints of a channel. The rules T-SEND, T-RECV, and T-CLOSE capture session-typed communication.

393	Runtime Types	$R ::= S \mid S^\#$
394	Names	a, b, c
395	Terms	$M ::= \dots \mid a$
396	Values	$U, V, W ::= a \mid \lambda x.M \mid () \mid (V, W) \mid \mathbf{inl} V \mid \mathbf{inr} V$
397	Configurations	$C, \mathcal{D}, \mathcal{E} ::= (va)C \mid C \parallel \mathcal{D} \mid \phi M \mid \mathbf{halt} \mid \zeta a \mid a(\vec{V}) \rightsquigarrow b(\vec{W})$
398	Thread Flags	$\phi ::= \bullet \mid \circ$
399	Top-level threads	$\mathcal{T} ::= \bullet M \mid \mathbf{halt}$
400	Auxiliary threads	$\mathcal{A} ::= \circ M \mid \zeta a \mid a(\vec{V}) \rightsquigarrow b(\vec{W})$
401	Type Environments	$\Gamma ::= \dots \mid \Gamma, a : S$
402	Runtime Type Environments	$\Delta ::= \cdot \mid \Delta, a : R$
403	Evaluation Contexts	$E ::= [] \mid EM \mid VE$ $\mid \mathbf{let} () = E \mathbf{in} M \mid (E, M) \mid (V, E) \mid \mathbf{let} (x, y) = E \mathbf{in} M$ $\mid \mathbf{inl} E \mid \mathbf{inr} E \mid \mathbf{case} E \mathbf{of} \{\mathbf{inl} x \mapsto M; \mathbf{inr} x \mapsto N\}$ $\mid \mathbf{fork} E \mid \mathbf{send} EM \mid \mathbf{send} VE \mid \mathbf{receive} E \mid \mathbf{close} E$ $\mid \mathbf{cancel} E \mid \mathbf{try} E \mathbf{as} x \mathbf{in} M \mathbf{otherwise} N$
404		
405	Pure Contexts	$P ::= [] \mid PM \mid VP$ $\mid \mathbf{let} () = P \mathbf{in} M \mid \mathbf{let} (x, y) = P \mathbf{in} M \mid (P, M) \mid (V, P)$ $\mid \mathbf{inl} P \mid \mathbf{inr} P \mid \mathbf{case} P \mathbf{of} \{\mathbf{inl} x \mapsto M; \mathbf{inr} x \mapsto N\}$ $\mid \mathbf{fork} P \mid \mathbf{send} PM \mid \mathbf{send} VP \mid \mathbf{receive} P \mid \mathbf{close} P$ $\mid \mathbf{cancel} P$
406		
407		
408		
409		
410		
411		
412	Thread Contexts	$\mathcal{F} ::= \phi E$
413	Configuration Contexts	$\mathcal{G} ::= [] \mid (va)\mathcal{G} \mid \mathcal{G} \parallel C$
414		
415	Syntactic Sugar	
416		$\zeta V \triangleq \zeta a_1 \parallel \dots \parallel \zeta a_n \text{ where } \text{fn}(V) = \{a_i\}_i$
417		$\zeta P \triangleq \zeta a_1 \parallel \dots \parallel \zeta a_n \text{ where } \text{fn}(P) = \{a_i\}_i$
418		$\zeta E \triangleq \zeta a_1 \parallel \dots \parallel \zeta a_n \text{ where } \text{fn}(E) = \{a_i\}_i$

Fig. 5. Runtime Syntax

As exceptions do not return values, the rule T-RAISE allows an exception to be given any type A . Rule T-TRY embraces explicit success continuations as advocated by [Benton and Kennedy \[2001\]](#), binding a result in M if L evaluates successfully. The T-CANCEL rule explicitly discards an endpoint. Naïvely implemented, cancellation violates progress: a thread could discard an endpoint, leaving a peer waiting forever. We avoid this pitfall by raising an exception when a communication action would wait forever due to cancellation.

2.3 Operational Semantics

We now give a small-step operational semantics for EGV.

Runtime Syntax. Fig. 5 shows the runtime syntax of EGV. We write $S^\#$ for the type of a channel which can be split into two endpoints of types S and \bar{S} . Runtime types R are either session types or channel types. We extend the syntax of terms to include names ranged over by a, b, c . Depending on context, a name a is variously used to identify a channel of type $S^\#$ and each of its endpoints of type S and \bar{S} . Values are standard. The semantics makes use of *configurations*, which are similar to π -calculus processes: $(va)C$ binds name a in configuration C , and $C \parallel \mathcal{D}$ is the parallel composition of configurations C and \mathcal{D} . Program threads take the form ϕM , where ϕ is a thread flag identifying whether the term is the *main thread* (\bullet), which returns a top-level result, or a *child thread* (\circ), which

442	Term Reduction	$M \rightarrow_M N$
443	E-LAM	$(\lambda x.M) V \rightarrow_M M\{V/x\}$
444	E-UNIT	$\mathbf{let} () = () \mathbf{in} M \rightarrow_M M$
445	E-PAIR	$\mathbf{let} (x, y) = (V, W) \mathbf{in} M \rightarrow_M M\{V/x, W/y\}$
446	E-INL	$\mathbf{case inl} V \mathbf{of} \{\mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N\} \rightarrow_M M\{V/x\}$
447	E-INR	$\mathbf{case inr} V \mathbf{of} \{\mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N\} \rightarrow_M N\{V/y\}$
448	E-VAL	$\mathbf{try} V \mathbf{as} x \mathbf{in} M \mathbf{otherwise} N \rightarrow_M M\{V/x\}$
449	E-LIFT	$E[M] \rightarrow_M E[M'], \text{ if } M \rightarrow_M M'$
450	Configuration Equivalence	$C \equiv D$
451		$C \parallel (D \parallel E) \equiv (C \parallel D) \parallel E \quad C \parallel D \equiv D \parallel C \quad (va)(vb)C \equiv (vb)(va)C$
452		$C \parallel (va)D \equiv (va)(C \parallel D), \text{ if } a \notin \text{fn}(C)$
453		$a(\vec{V}) \rightsquigarrow b(\vec{W}) \equiv b(\vec{W}) \rightsquigarrow a(\vec{V}) \quad \circ () \parallel C \equiv C \quad (va)(vb)(\zeta a \parallel \zeta b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)) \parallel C \equiv C$
454	Configuration Reduction	$C \rightarrow D$
455	E-FORK	$\mathcal{F}[\mathbf{fork}(\lambda x.M)] \rightarrow (va)(vb)(\mathcal{F}[a] \parallel \circ M\{b/x\} \parallel a(\epsilon) \rightsquigarrow b(\epsilon)), \text{ where } a, b \text{ are fresh}$
456	E-SEND	$\mathcal{F}[\mathbf{send} U a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W}) \rightarrow \mathcal{F}[a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W} \cdot U)$
457	E-RECEIVE	$\mathcal{F}[\mathbf{receive} a] \parallel a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W}) \rightarrow \mathcal{F}[U, a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})$
458	E-CLOSE	$(va)(vb)(\mathcal{F}[\mathbf{close} a] \parallel \mathcal{F}'[\mathbf{close} b] \parallel a(\epsilon) \rightsquigarrow b(\epsilon)) \rightarrow \mathcal{F}[\mathbf{close} a] \parallel \mathcal{F}'[\mathbf{close} b]$
459	E-CANCEL	$\mathcal{F}[\mathbf{cancel} a] \rightarrow \mathcal{F}[\mathbf{close} a] \parallel \zeta a$
460	E-ZAP	$\zeta a \parallel a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W}) \rightarrow \zeta a \parallel \zeta U \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})$
461	E-CLOSEZAP	$\mathcal{F}[\mathbf{close} a] \parallel \zeta b \parallel a(\epsilon) \rightsquigarrow b(\epsilon) \rightarrow \mathcal{F}[\mathbf{raise}] \parallel \zeta a \parallel \zeta b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)$
462	E-RECEIVEZAP	$\mathcal{F}[\mathbf{receive} a] \parallel \zeta b \parallel a(\epsilon) \rightsquigarrow b(\vec{W}) \rightarrow \mathcal{F}[\mathbf{raise}] \parallel \zeta a \parallel \zeta b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})$
463	E-RAISE	$\mathcal{F}[\mathbf{try} P[\mathbf{raise}] \mathbf{as} x \mathbf{in} M \mathbf{otherwise} N] \rightarrow \mathcal{F}[N] \parallel \zeta P$
464	E-RAISECHILD	$\circ P[\mathbf{raise}] \rightarrow \zeta P$
465	E-RAISEMAIN	$\bullet P[\mathbf{raise}] \rightarrow \mathbf{halt} \parallel \zeta P$
466	E-LIFTC	$\mathcal{G}[C] \rightarrow \mathcal{G}[D], \text{ if } C \rightarrow D$
467	E-LIFTM	$\phi M \rightarrow \phi M', \text{ if } M \rightarrow_M M'$

Fig. 6. Reduction and Equivalence for Terms and Configurations

does not, and must return the unit value. A configuration has at most one main thread. As well as program threads, configurations include three special forms of thread. A *zapper thread* (ζa) manages an endpoint a that has been cancelled, and is used to propagate failure. A *halted thread* (**halt**) arises when the main thread has crashed due to an uncaught exception. A *buffer thread* ($a(\vec{V}) \rightsquigarrow b(\vec{W})$) models asynchrony: \vec{V} and \vec{W} are sequences of values ready to be received along endpoints a and b respectively. We sometimes find it useful to distinguish top-level threads \mathcal{T} (main threads and halted threads) from auxiliary threads \mathcal{A} (child threads, zapper threads, and buffer threads).

Environments. We extend type environments Γ to include runtime names of session type and introduce runtime type environments Δ , which type both buffer endpoints of session type and channels of type $S^\#$ for some S , but not object variables.

Contexts. Evaluation contexts E are set up for standard left-to-right call-by-value evaluation. Pure contexts P are those evaluation contexts that include no exception handling frames. Thread

491 contexts \mathcal{F} support reduction in program threads. Configuration contexts \mathcal{G} support reduction
 492 under ν -binders and parallel composition.

493
 494 *Free Names.* We let the meta operation $\text{fn}(-)$ denote the set of free names in a term, type
 495 environment, buffer environment, value, configuration, pure context, or evaluation context.

496
 497 *Syntactic Sugar.* We follow the standard convention that parallel composition of configurations
 498 associates to the right. We write ζV , ζP , and ζE , as shorthand for the parallel composition of zipper
 499 threads for each free name in values V , pure contexts P , and evaluation contexts E , respectively.

500 Following prior work on linear functional languages with session types [Gay and Vasconcelos
 501 2010; Lindley and Morris 2015, 2016, 2017], we present the semantics of EGV via a deterministic
 502 reduction relation on terms (\longrightarrow_M), an equivalence relation on configurations (\equiv), and a nondeter-
 503 ministic reduction relation on configurations (\longrightarrow). We write \Longrightarrow for the relation $\equiv \longrightarrow \equiv$. Fig. 6
 504 presents reduction and equivalence rules for terms and configurations.

505
 506 *Term Reduction.* Reduction on terms is standard call-by-value β -reduction.

507
 508 *Configuration Equivalence.* A running program can make use of the standard structural π -calculus
 509 equivalence rules [Milner 1999] of associativity and commutativity of parallel composition, name
 510 restriction reordering, and scope extrusion. Formally, equivalence is defined as the smallest con-
 511 gruence relation satisfying the equivalence axioms in Figure 6. We incorporate a further rule to
 512 allow buffers to be treated symmetrically and two garbage collection rules, allowing completed
 513 child threads and cancelled empty buffers to be discarded.

514
 515 *Communication and Concurrency.* The E-FORK rule creates two fresh names for each endpoint of
 516 a channel, returning one name and substituting the other in the body of the spawned thread, as
 517 well as creating a channel with two empty buffers. The E-SEND and E-RECEIVE rules send to and
 518 receive from a buffer. The E-CLOSE rule discards an empty buffer once a session is complete.

519
 520 *Cancellation.* The E-CANCEL rule cancels an endpoint by creating a zipper thread. The E-ZAP rule
 521 ensures that when an endpoint is cancelled, all endpoints in the buffer of the cancelled endpoint are
 522 also cancelled: it dequeues a value from the head of the buffer and cancels any endpoints contained
 523 within the dequeued value (ζU). It is applied repeatedly until the buffer is empty.

524
 525 *Raising Exceptions.* Following Mostrous and Vasconcelos [2014], an exception is raised when it
 526 would be otherwise impossible for a communication action to succeed. The E-RECEIVEZAP rule
 527 raises an exception if an attempt is made to receive along an endpoint whose buffer is empty and
 528 whose peer endpoint has been cancelled. Similarly, E-CLOSEZAP raises an exception if an attempt is
 529 made to close a channel where the peer endpoint has been cancelled. There is no rule for the case
 530 where a thread tries to send a value along a cancelled endpoint; the free names in the communicated
 531 value must eventually be cancelled, but this is achieved through E-ZAP. We choose not to raise an
 532 exception in this case since to do so would violate confluence, which we discuss in more detail
 533 in §3.4. Not raising exceptions on message sends to dead peers is standard behaviour for languages
 534 such as Erlang.

534
 535 *Handling Exceptions.* The E-RAISE rule invokes the **otherwise** clause if an exception is raised,
 536 while also cancelling all endpoints in the enclosing pure context. If an unhandled exception occurs
 537 in a child thread, then all free endpoints in the evaluation context are cancelled and the thread
 538 is terminated (E-RAISECHILD). If the exception is in the main thread then all free endpoints are
 539 cancelled and the main thread reduces to **halt** (E-RAISEMAIN).

2.4 Synchrony

As we are interested in writing distributed applications, we consider asynchronous session types. However, our semantics adapts straightforwardly to the synchronous setting, where a send to a cancelled peer must also raise an exception:

$$\begin{array}{l}
 \text{E-SYNC COMM} \quad \mathcal{F}[\mathbf{send} \ V \ a] \parallel \mathcal{F}'[\mathbf{receive} \ a] \longrightarrow \mathcal{F}[a] \parallel \mathcal{F}'[(V, a)] \\
 \text{E-SYNC SEND ZAP} \quad \mathcal{F}[\mathbf{send} \ V \ a] \parallel \not\downarrow a \longrightarrow \mathcal{F}[\mathbf{raise}] \parallel \not\downarrow V \parallel \not\downarrow a \parallel \not\downarrow a \\
 \text{E-SYNC RECV ZAP} \quad \mathcal{F}[\mathbf{receive} \ a] \parallel \not\downarrow a \longrightarrow \mathcal{F}[\mathbf{raise}] \parallel \not\downarrow a \parallel \not\downarrow a \\
 (va)(\not\downarrow a \parallel \not\downarrow a) \parallel C \equiv C
 \end{array}$$

3 METATHEORY

Even in the presence of channel cancellation and exceptions, EGV retains GV's strong metatheory [Lindley and Morris 2015]. The central property of session-typed systems is session fidelity: all communication follows the prescribed session types. Session fidelity follows as a corollary of preservation of configuration typing under reduction.

Session calculi with roots in linear logic are deadlock-free as interpreting the logical cut rule as a combination of name restriction and parallel composition necessarily ensures acyclicity [Caires and Pfenning 2010]. It is also possible to use deadlock-freedom to derive a global progress result. We prove that global progress holds even in the presence of channel cancellation. (Our proof is direct, not requiring catalyser processes [Carbone et al. 2014; Mostrous and Vasconcelos 2014].) We also prove that EGV is confluent and terminating.

3.1 Runtime Typing

To state our main results we require typing rules for names and configurations. These are given in Fig. 7. The configuration typing judgement has the shape $\Gamma; \Delta \vdash^\phi C$, which states that under type environment Γ , runtime environment Δ , and thread flag ϕ , configuration C is well-typed. We additionally require that $\text{fn}(\Gamma) \cap \text{fn}(\Delta) = \emptyset$. Thread flags ensure that there can be at most one top-level thread which can return a value: \bullet denotes a configuration with a top-level thread and \circ denotes a configuration without. The main thread returns the result of running a program. Any configuration C such that $\Gamma; \Delta \vdash^\bullet C$ has exactly one main thread or halted thread as a subconfiguration. We write $\Gamma; \Delta \vdash^\bullet C : A$ whenever the derivation of $\Gamma; \Delta \vdash^\bullet C$ contains a subderivation of the form

$$\frac{\Gamma' \vdash M : A}{\Gamma'; \cdot \vdash^\bullet \bullet M} \quad \text{or} \quad \frac{}{\cdot; \cdot \vdash^\bullet \mathbf{halt}}$$

We say that a C is a *ground configuration* if there exists A such that $\cdot; \cdot \vdash^\bullet C : A$ and A contains no session types or function types.

The T-NU rule introduces a channel name; T-CONNECT₁ and T-CONNECT₂ connect two configurations over a channel; and T-MIX composes two configurations that share no channels. The latter three rules use the $+$ operator to combine the flags from subconfigurations. The T-MAIN and T-CHILD rules introduce main and child threads. Child threads always return the unit value. The T-HALT rule types the **halt** configuration, which signifies that an unhandled exception has occurred in the main thread. The T-ZAP rule types a zapper thread, given a single name in the type environment. The T-BUFFER rule ensures that buffers contain values corresponding to the session types of their endpoints. This is the only rule that consumes names from the runtime environment. Buffers rely on two auxiliary judgements. The queue typing judgement $\Gamma \vdash \vec{V} : \vec{A}$ states that under type environment Γ , the sequence of values \vec{V} have types \vec{A} . The session slicing operator S/\vec{A} captures reasoning about session types discounting values contained in the buffer. The session

<p>589 Term Typing $\boxed{\Gamma \vdash M : A}$</p> <p>590 T-NAME</p> <p>591 $\frac{}{a : S \vdash a : S}$</p> <p>592</p> <p>593 Configuration Typing</p> <p>594 $\boxed{\Gamma; \Delta \vdash^\phi C}$</p> <p>595 T-NU</p> <p>596 $\frac{\Gamma; \Delta, a : S^\# \vdash^\phi C}{\Gamma; \Delta \vdash^\phi (\nu a)C}$</p> <p>597</p> <p>598 T-MIX</p> <p>599 $\frac{\Gamma_1; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2 \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}}$</p> <p>600 T-CONNECT₁</p> <p>601 $\frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}}$</p> <p>602</p> <p>603 T-CONNECT₂</p> <p>604 $\frac{\Gamma_1; \Delta_1, a : \bar{S} \vdash^{\phi_1} C \quad \Gamma_2, a : S; \Delta_2 \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}}$</p> <p>605</p> <p>606 T-MAIN $\frac{\Gamma \vdash M : A}{\Gamma; \cdot \vdash^\bullet \bullet M}$</p> <p>607 T-CHILD $\frac{\Gamma \vdash M : 1}{\Gamma; \cdot \vdash^\circ \circ M}$</p> <p>608 T-HALT $\frac{}{\cdot; \cdot \vdash^\bullet \mathbf{halt}}$</p> <p>609 T-ZAP $\frac{}{a : S; \cdot \vdash^\circ \not\downarrow a}$</p> <p>610 T-BUFFER</p> <p>611 $\frac{S/\bar{A} = \overline{S'/\bar{B}} \quad \Gamma_1 \vdash \bar{V} : \bar{A} \quad \Gamma_2 \vdash \bar{W} : \bar{B}}{\Gamma_1, \Gamma_2; a : S, b : S' \vdash^\circ a(\bar{V}) \rightsquigarrow b(\bar{W})}$</p> <p>612 Flag Combination $\boxed{\phi_1 + \phi_2 = \phi_3}$</p> <p>613 $\bullet + \circ = \bullet \quad \circ + \bullet = \bullet$</p> <p>614 $\circ + \circ = \circ \quad \bullet + \bullet \text{ undefined}$</p> <p>615 Environment Reduction $\boxed{\Gamma; \Delta \longrightarrow \Gamma'; \Delta'}$</p> <p>616 $\frac{S \longrightarrow S'}{\Gamma, a : S; \Delta \longrightarrow \Gamma, a : S'; \Delta}$</p> <p>617 $\frac{S \longrightarrow S'}{\Gamma; \Delta, a : S \longrightarrow \Gamma; \Delta, a : S'}$</p> <p>618 $\frac{S \longrightarrow S'}{\Gamma; \Delta, a : S^\# \longrightarrow \Gamma; \Delta, a : S'^\#}$</p>	<p>Session Slicing $\boxed{S/\bar{A}}$</p> <p>$\frac{S/\epsilon = S}{!A.S/A \cdot \bar{A} = S/\bar{A}}$</p> <p>Queue Typing $\boxed{\Gamma \vdash \bar{V} : \bar{A}}$</p> <p>$\frac{\Gamma_1 \vdash V : A \quad \Gamma_2 \vdash \bar{V} : \bar{A}}{\Gamma_1, \Gamma_2 \vdash V \cdot \bar{V} : A \cdot \bar{A}}$</p> <p>$\frac{}{\cdot \vdash \epsilon : \epsilon}$</p>	
<p>619</p> <p>620</p> <p>621</p> <p>622</p> <p>623</p> <p>624</p> <p>625</p> <p>626</p> <p>627</p> <p>628</p> <p>629</p> <p>630</p> <p>631</p> <p>632</p> <p>633</p> <p>634</p> <p>635</p> <p>636</p> <p>637</p>		

Fig. 7. Runtime Typing

types of two buffer endpoints are compatible if they are dual up to values contained in the buffer. The partiality of the slicing operator coupled with the duality constraint ensures that at least one queue in a buffer is always empty. Appendix A shows an example configuration typing derivation.

3.2 Preservation

Preservation for the functional fragment of EGV is standard.

LEMMA 3.1 (PRESERVATION (TERMS)). *If $\Gamma \vdash M : A$ and $M \longrightarrow_M M'$, then $\Gamma \vdash M' : A$.*

Given a relation \mathcal{R} , we write $\mathcal{R}^?$ for its reflexive closure. We write Ψ for the restriction of type environments Γ to contain runtime names but no variables:

$$\Psi ::= \cdot \mid \Psi, a : S$$

Preservation of typing by configuration reduction holds only for closed configurations.

THEOREM 3.2 (PRESERVATION). *If $\Psi; \Delta \vdash^\phi C$ and $C \longrightarrow C'$, then there exist Ψ', Δ' such that $\Psi; \Delta \longrightarrow^? \Psi'; \Delta'$ and $\Psi'; \Delta' \vdash^\phi C'$.*

PROOF. By induction on the derivation of $C \longrightarrow C'$, making use of Lemma 3.1, and lemmas for subconfiguration typeability and replacement. The proof cases can be found in Appendix C.1. \square

638 *Typing and Configuration Equivalence.* As is common in logically-inspired session-typed func-
 639 tional languages [Lindley and Morris 2015, 2017], typeability of configurations is *not* preserved
 640 by equivalence. Consider $\Gamma; \Delta \vdash^\phi (va)(vb)(C \parallel (\mathcal{D} \parallel \mathcal{E}))$ with $a \in \text{fn}(C)$, $b \in \text{fn}(\mathcal{D})$, and
 641 $a, b \in \text{fn}(\mathcal{E})$. But $\Gamma; \Delta \not\vdash^\phi (va)(vb)((C \parallel \mathcal{D}) \parallel \mathcal{E})$. Fortunately this looseness of the equivalence
 642 relation is unproblematic: we may always safely re-associate parallel composition (for example,
 643 $\Gamma; \Delta \vdash^\phi (va)(vb)((C \parallel \mathcal{E}) \parallel \mathcal{D})$; see Appendix C.1), and any reduction sequence which uses ill-typed
 644 equivalences may be replaced by one that does not.

645 **THEOREM 3.3 (PRESERVATION MODULO EQUIVALENCE).** *If $\Psi; \Delta \vdash^\phi C$, $C \equiv \mathcal{D}$, and $\mathcal{D} \longrightarrow \mathcal{D}'$, then:*

646 (1) *There exists some $\mathcal{E} \equiv \mathcal{D}$ and some \mathcal{E}' such that $\Psi; \Delta \vdash^\phi \mathcal{E}$ and $\mathcal{E} \longrightarrow \mathcal{E}'$*

647 (2) *There exist Ψ' , Δ' such that $\Psi; \Delta \longrightarrow^? \Psi'; \Delta'$ and $\Psi'; \Delta' \vdash^\phi \mathcal{E}'$*

648 (3) *$\mathcal{D}' \equiv \mathcal{E}'$*

649 **PROOF.** The only non-trivial reductions are those involving a synchronisation with a buffer
 650 (E-SEND, E-RECEIVE, E-CLOSE, E-ZAP, E-CLOSEZAP, E-RECEIVEZAP). The only equivalence rule that
 651 can lead to an ill-typed configuration is associativity of parallel composition
 652

$$653 \quad C \parallel (\mathcal{D} \parallel \mathcal{E}) \equiv (C \parallel \mathcal{D}) \parallel \mathcal{E}$$

654 where both compositions arise from the T-CONNECT₁ and T-CONNECT₂ rules. The only reason to
 655 apply the associativity rule from left-to-right is to enable threads inside C and \mathcal{D} to synchronise.
 656 But for synchronisation to be possible there must exist a name a such that $a \in \text{fn}(C)$ and $a \in \text{fn}(\mathcal{D})$.
 657 Because the left-hand-side of the equation is well-typed, we know that C and \mathcal{E} have no names in
 658 common, that \mathcal{D} and \mathcal{E} share a name, and that the right-hand-side must be well-typed as there is
 659 still exactly one channel connecting each of the parallel compositions. The argument for applying
 660 the rule from right-to-left is symmetric. In summary, any ill-typed use of equivalence is useless. \square

662 3.3 Progress

663 To prove that EGV enjoys a strong notion of progress we identify a *canonical form* for configura-
 664 tions. We prove that every well-typed configuration is equivalent to a well-typed configuration
 665 in canonical form, and that ground configurations can always either reduce, or are equivalent to
 666 either a value or **halt**.

667 The functional fragment of EGV enjoys progress.

668 **LEMMA 3.4 (PROGRESS: OPEN TERMS).** *If $\Psi \vdash M : A$, then either:*

- 669 • *M is a value;*
- 670 • *there exists some M' such that $M \longrightarrow_M M'$; or*
- 671 • *M has the form $E[M']$, where M' is a session typing primitive of the form: **fork** V , **send** $V W$,*
 672 ***receive** V , **close** V , or **cancel** V .*

673 **PROOF.** By induction on the derivation of $\Psi \vdash M : A$. \square

674 To reason about progress of configurations, we characterise *canonical forms*, which make explicit
 675 the property that at most one name is shared between threads. Recall that \mathcal{A} ranges over auxiliary
 676 threads and \mathcal{T} over top-level threads (Fig. 5). Let \mathcal{M} range over configurations of the form:

$$677 \quad \mathcal{A}_1 \parallel \cdots \parallel \mathcal{A}_m \parallel \mathcal{T}$$

678 **Definition 3.5 (Canonical Form).** A configuration C is in *canonical form* if there is a sequence of
 679 names a_1, \dots, a_n , a sequence of configurations $\mathcal{A}_1, \dots, \mathcal{A}_n$, and a configuration \mathcal{M} , such that:

$$680 \quad C = (va_1)(\mathcal{A}_1 \parallel (va_2)(\mathcal{A}_2 \parallel \cdots \parallel (va_n)(\mathcal{A}_n \parallel \mathcal{M}) \dots))$$

681 where $a_i \in \text{fn}(\mathcal{A}_i)$ for each $i \in 1..n$.

687 The following lemma implies that communication topologies are always acyclic.

688 LEMMA 3.6. *If $\Gamma; \Delta \vdash^\phi C$ and $C = \mathcal{G}[\mathcal{D} \parallel \mathcal{E}]$, then $\text{fn}(\mathcal{D}) \cap \text{fn}(\mathcal{E})$ is either \emptyset or $\{a\}$ for some a .*

689
690 PROOF. By induction on the derivation of $\Gamma; \Delta \vdash^\phi C$; the only interesting rules are those for
691 parallel composition. As the environments are well-formed, $\text{fn}(\Gamma) \cap \text{fn}(\Delta) = \emptyset$. Thus, T-CONNECT₁
692 and T-CONNECT₂ allow exactly one name to be shared, whereas T-Mix forbids sharing of names. \square

693 All well-typed configurations can be written in canonical form.

694
695 THEOREM 3.7 (CANONICAL FORMS). *Given C such that $\Gamma; \Delta \vdash^\bullet C$, there exists some $\mathcal{D} \equiv C$ such
696 that $\Gamma; \Delta \vdash^\bullet \mathcal{D}$ and \mathcal{D} is in canonical form.*

697 PROOF. By induction on the count of ν -bound variables, following Lindley and Morris [2015] and
698 making use of Lemma 3.6. The additional features of EGV do not change the essential argument.
699 The full proof can be found in Appendix C.2. \square

700
701 Definition 3.8. We say that term M is *ready to perform an action on name a* if M is about to send
702 on, receive on, close, or cancel a . Formally:

$$703 \text{ready}(a, M) \triangleq \exists E. (M = E[\text{send } V \ a]) \vee (M = E[\text{receive } a]) \vee (M = E[\text{close } a]) \vee (M = E[\text{cancel } a])$$

704 Using the notion of a ready thread, we may classify a notion of progress for open configurations.

705
706 THEOREM 3.9 (PROGRESS: OPEN). *Suppose $\Psi; \Delta \vdash^\bullet C$, where C is in canonical form.*

707 *Let $C = (\nu a_1)(\mathcal{A}_1 \parallel (\nu a_2)(\mathcal{A}_2 \parallel \dots \parallel (\nu a_n)(\mathcal{A}_n \parallel \mathcal{M})) \dots)$.*

708 *Either there exists some C' such that $C \Longrightarrow C'$, or:*

709 (1) *For $1 \leq i \leq n$, each auxiliary thread \mathcal{A}_i is either:*

- 710 (a) *a child thread $\circ M$ for which there exists $a \in \{a_j \mid 1 \leq j \leq i\} \cup \text{fn}(\Psi)$ such that $\text{ready}(a, M)$;*
- 711 (b) *a zipper thread $\frac{1}{2} a_i$; or*
- 712 (c) *a buffer.*

713 (2) *$\mathcal{M} = \mathcal{A}'_1 \parallel \dots \parallel \mathcal{A}'_m \parallel \mathcal{T}$ such that for $1 \leq j \leq m$:*

- 714 (a) *\mathcal{A}'_j is either:*
 - 715 (i) *a child thread $\circ N$ with $N = ()$ or $\text{ready}(a, N)$ for some $a \in \{a_i \mid 1 \leq i \leq n\} \cup \text{fn}(\Psi) \cup \text{fn}(\Delta)$;*
 - 716 (ii) *a zipper thread $\frac{1}{2} a$ for some $a \in \{a_i \mid 1 \leq i \leq n\} \cup \text{fn}(\Psi) \cup \text{fn}(\Delta)$; or*
 - 717 (iii) *a buffer.*

- 718 (b) *Either $\mathcal{T} = \bullet N$, where N is either a value or $\text{ready}(a, N)$ for some $a \in \{a_i \mid 1 \leq i \leq$
719 $n\} \cup \text{fn}(\Psi) \cup \text{fn}(\Delta)$; or $\mathcal{T} = \mathbf{halt}$.*

720
721 PROOF. The result follows from a more verbose, but finer-grained, property which we prove by
722 induction on the derivation of $\Psi; \Delta \vdash^\bullet C$. Full details are in Appendix C.3. \square

723 This theorem tells us that open reduction cannot “go wrong”. A progress theorem states that
724 either reduction is possible or the configuration is a value. Conditions 1(a)(b)(c) and 2(a)(b) constitute
725 a suitable generalisation of ‘value’.

726 By restricting attention to closed environments, we obtain a tighter progress property.

727
728 THEOREM 3.10 (PROGRESS: CLOSED). *Suppose $\cdot; \cdot \vdash^\bullet C$ where C is in canonical form.*

729 *Let $C = (\nu a_1)(\mathcal{A}_1 \parallel (\nu a_2)(\mathcal{A}_2 \parallel \dots \parallel (\nu a_n)(\mathcal{A}_n \parallel \mathcal{M})) \dots)$.*

730 *Either there exists some C' such that $C \Longrightarrow C'$, or:*

731 (1) *For $1 \leq i \leq n$, each auxiliary thread \mathcal{A}_i is either:*

- 732 (a) *a child thread $\circ M$ for some M such that $\text{ready}(a_i, M)$; or*
- 733 (b) *a zipper thread $\frac{1}{2} a_i$; or*
- 734 (c) *a buffer.*

(2) Either $M = \bullet W$ for some value W , or $M = \mathbf{halt}$.

The above progress results do not specifically mention deadlock. However, Lemma 3.6 ensures deadlock-freedom. Nevertheless, communication can still be blocked if an endpoint appears in the value returned by the main thread. A conservative way of disallowing endpoints in the result is to insist that the return type of the program be free of session types and function types (closures may capture endpoints). All configurations of such programs are ground configurations.

THEOREM 3.11 (GLOBAL PROGRESS). *Suppose C is a ground configuration. Either there exists some C' such that $C \Longrightarrow C'$; or $C \equiv \bullet V$; or $C \equiv \mathbf{halt}$.*

PROOF. As a consequence of Theorem 3.10, either there exists some C' such that $C \Longrightarrow C'$, or $C \not\Longrightarrow$ and each thread \mathcal{A}_i must be a zapper, a buffer, or ready to perform an action. If $C \not\Longrightarrow$, since C is ground, by Lemma 3.6, we have that no thread can be ready to perform an action. Thus, each \mathcal{A}_i must be either $\circ()$, a zapper, or an empty buffer. The result then follows by the garbage collection congruences of Fig. 6. \square

3.4 Confluence

EGV enjoys a strong form of confluence known as the *diamond property* [Barendregt 1984].

THEOREM 3.12 (DIAMOND PROPERTY). *If $\Psi; \Delta \vdash^\phi C$, and $C \Longrightarrow \mathcal{D}_1$, and $C \Longrightarrow \mathcal{D}_2$, then either $\mathcal{D}_1 \equiv \mathcal{D}_2$, or there exists some \mathcal{D}_3 such that $\mathcal{D}_1 \Longrightarrow \mathcal{D}_3$ and $\mathcal{D}_2 \Longrightarrow \mathcal{D}_3$.*

PROOF. First, note that \longrightarrow_M is entirely deterministic and hence confluent due to the call-by-value, left-to-right ordering imposed by evaluation contexts. By linearity, we know that endpoints to different buffers may not be shared, so it follows that communication actions on different channels may be performed in any order. Asynchrony and cancellation introduce two critical pairs which may be resolved in a single step; see Appendix C.4 for details. \square

Remark. The system becomes non-confluent if we choose to raise an exception when sending to a cancelled buffer. Suppose that instead of the current semantics, we were to replace E-SEND with the following two rules:

$$\begin{aligned} (vb)(\mathcal{F}[\mathbf{send} \ U \ a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W}) \parallel \phi M) &\longrightarrow (vb)(\mathcal{F}[a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W}) \cdot U \parallel \phi M) \\ \mathcal{F}[\mathbf{send} \ U \ a] \parallel \cancel{b} \parallel a(\vec{V}) \rightsquigarrow b(\vec{W}) &\longrightarrow \mathcal{F}[\mathbf{raise}] \parallel \cancel{b} \parallel \cancel{U} \parallel a(\vec{V}) \rightsquigarrow b(\vec{W}) \end{aligned}$$

Then, sending and cancelling peer endpoints of a buffer results in a non-convergent critical pair:

$$\begin{array}{ccc} & (vb)(\mathcal{F}[\mathbf{send} \ U \ a] \parallel \mathcal{F}[\mathbf{cancel} \ b] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})) & \\ & \swarrow \qquad \qquad \qquad \searrow & \\ (vb)(\mathcal{F}[a] \parallel \mathcal{F}[\mathbf{cancel} \ b] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W}) \cdot U) & & (vb)(\mathcal{F}[\mathbf{send} \ U \ a] \parallel \mathcal{F}[\mathbf{cancel} \ b] \parallel \cancel{b} \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})) \\ \downarrow & & \downarrow \\ (vb)(\mathcal{F}[a] \parallel \mathcal{F}[\mathbf{cancel} \ b] \parallel \cancel{b} \parallel a(\vec{V}) \rightsquigarrow b(\vec{W}) \cdot U) & & (vb)(\mathcal{F}[\mathbf{raise}] \parallel \mathcal{F}[\mathbf{cancel} \ b] \parallel \cancel{b} \parallel \cancel{U} \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})) \end{array}$$

In either case, the endpoints contained in U will still eventually be cancelled, thus preservation and global progress still hold. However, the lack of confluence affects exactly *when* the exception is raised in context \mathcal{F} . This decision has practical significance, in that it characterises the race between sending a message and propagating a cancellation notification.

3.5 Termination

As EGV is linear, it has an elementary strong normalisation proof.

THEOREM 3.13 (STRONG NORMALISATION). *If $\Psi; \Delta \vdash^\phi C$, then there are no infinite \Longrightarrow reduction sequences from C .*

PROOF. Let the size of a configuration be the sum of the sizes of the abstract syntax trees of all of the terms contained in its main threads, child threads, and buffers, modulo exhaustively applying the garbage collection equivalences from left-to-right. The size of a configuration is invariant under \equiv and strictly decreases under \longrightarrow , hence \implies reduction must always terminate. \square

We conjecture that the strong normalisation result continues to hold in the presence of unrestricted types or shared channels for session initiation, but the proof technique is necessarily more involved. We believe that a logical relations argument along the lines of Pérez et al. [2012] or a CPS translation along the lines of Lindley and Morris [2016] would suffice.

4 EXTENSIONS

4.1 User-defined Exceptions with Payloads

In order to focus on the interplay between exceptions and session types we have thus far considered handling a single kind of exception. In practice it can be useful to distinguish between multiple kinds of user-defined exception, each of which may carry a payload.

Consider again handling the exception in checkDetails. An exception may arise if the database is corrupt, or if there are too many connections. We might like to handle each case separately:

```

802   exnServer4(s)  $\triangleq$ 
803     let ((username, password), s) = receive s in
804     try checkDetails(username, password) as res in
805       if res then let s = select Authenticated s in serverBody(s)
806       else let s = select AccessDenied s in close s
807     unless
808       DBCorrupt(y)  $\mapsto$  cancel s; log("Database Corrupt: " + y)
809       TooManyConnections(y)  $\mapsto$  cancel s; log("Too many connections: " + y)
810

```

An exception in checkDetails might be raised by the term **raise** DatabaseCorrupt(filename), for example. Our approach generalises straightforwardly to handle this example.

Syntax. Figure 8 shows extensions to EGV for exceptions with payloads. We introduce a type of exceptions, Exn. We assume a countably infinite set $X \in \mathbb{B}$ of exception names, and a type schema function $\Sigma(X) = A$ mapping exception names to payload types. We extend **raise** to take a term of type Exn as its argument. Finally, we generalise **tryLasxinMotherwiseN** to **tryLasxinMunlessH**, where H is an exception handler with clauses $\{X_i(y_i) \mapsto N_i\}_i$, such that X_i is an exception name; y_i binds the payload; and N_i is the clause to be evaluated when the exception is raised.

Typing Rules. The TP-EXN rule ensures that an exception's payload matches its expected type. The TP-RAISE and TP-TRY are the natural extensions of T-RAISE and T-TRY.

Semantics. Our presentation is similar to operational accounts of effect handlers; the formulation here is inspired by that of Hillerström et al. [2017]. To define the semantics of the generalised exception handling construct, we first introduce the auxiliary function handled(E), which defines the exceptions handled in a given evaluation context:

$$\begin{aligned} \text{handled}(P) &= \emptyset & \text{handled}(\mathbf{try } E \mathbf{ as } x \mathbf{ in } M \mathbf{ unless } H) &= \text{handled}(E) \cup \text{dom}(H) \\ \text{handled}(E) &= \text{handled}(E'), & \text{if } E \text{ is not a } \mathbf{try} \text{ and } E' \text{ is the immediate subcontext of } E \end{aligned}$$

The EP-RAISE rule handles an exception. The side conditions ensure that the exception is caught by the nearest matching handler and is handled by the appropriate clause. As with plain EGV, all free names are safely discarded. The EP-RAISECHILD and EP-RAISEMAIN rules cover the cases where an exception is unhandled. Due to the use of the handled function we no longer require pure contexts.

834	Syntax	
835	Types	$A, B ::= \dots \mid \text{Exn}$
836	Terms	$L, M, N ::= \dots \mid X(M) \mid \mathbf{raise} M \mid \mathbf{try} L \mathbf{as} x \mathbf{in} M \mathbf{unless} H$
837	Exception Handlers	$H ::= \{X_i(x_i) \mapsto N_i\}_i$
838	Runtime Syntax	
839	Evaluation Contexts $E ::= \dots \mid \mathbf{raise} E \mid \mathbf{try} E \mathbf{as} x \mathbf{in} M \mathbf{unless} H$	
840	Term typing	$\boxed{\Sigma(X) = A} \quad \boxed{\Gamma \vdash M : A}$
841		
842		TP-TRY
843	TP-EXN	$\frac{\Sigma(X) = A \quad \Gamma \vdash M : A}{\Gamma \vdash X(M) : \text{Exn}}$
844	TP-RAISE	$\frac{\Gamma \vdash M : \text{Exn}}{\Gamma \vdash \mathbf{raise} M : A}$
845		$\frac{\Gamma_1 \vdash L : A \quad \Gamma_2, x : A \vdash M : B \quad (\Gamma_2, y_i : \Sigma(X_i) \vdash N_i : B)_i}{\Gamma_1, \Gamma_2 \vdash \mathbf{try} L \mathbf{as} x \mathbf{in} M \mathbf{unless} \{X_i(y_i) \mapsto N_i\}_i : B}$
846	Term and Configuration Reduction	$\boxed{M \longrightarrow_M N} \quad \boxed{C \longrightarrow \mathcal{D}}$
847	EP-VAL	$\mathbf{try} V \mathbf{as} x \mathbf{in} M \mathbf{unless} H \longrightarrow_M M\{V/x\}$
848	EP-RAISE	$\mathcal{F}[\mathbf{try} E[\mathbf{raise} X(V)] \mathbf{as} x \mathbf{in} M \mathbf{unless} H] \longrightarrow \mathcal{F}[N\{V/y\}] \parallel \not\downarrow E$ where $X \notin \text{handled}(E)$
849		$(X(y) \mapsto N) \in H$
850	EP-RAISECHILD	$\circ E[\mathbf{raise} X(V)] \longrightarrow \not\downarrow E \parallel \not\downarrow V$ where $X \notin \text{handled}(E)$
851	EP-RAISEMAIN	$\bullet E[\mathbf{raise} X(V)] \longrightarrow \mathbf{halt} \parallel \not\downarrow E \parallel \not\downarrow V$ where $X \notin \text{handled}(E)$
852		
853		

Fig. 8. User-defined Exceptions with Payloads

All of EGV's metatheoretic properties (preservation, global progress, confluence, and termination) adapt straightforwardly to this extension.

4.2 Unrestricted Types and Access Points

Unrestricted (intuitionistic) types allow some values to be used in a non-linear fashion. Access points [Gay and Vasconcelos 2010] provide a more flexible method of session initiation than **fork**, allowing two threads to dynamically establish a session. Both features are useful in practice: unrestricted types because some data is naturally multi-use, and access points because they admit cyclic communication topologies supporting racey stateful servers such as chat servers. *Access points* decouple spawning a thread from establishing a session. An access point has the unrestricted type $\text{AP}(S)$; we write $\text{un}(A)$ to mean that A is unrestricted and $\text{un}(\Gamma)$ if $\text{un}(A_i)$ for all $x_i : A_i \in \Gamma$. Figure 9 shows the syntax, typing rules, and reduction rules for EGV extended with access points.

Unrestricted Types. To support unrestricted types, we introduce a splitting judgement $(\Gamma = \Gamma_1 + \Gamma_2)$, which allows variables of unrestricted type to be shared across sub-environments, but requires linear variables to be used only in a single sub-environment. We relax rule T-VAR to allow the use of unrestricted environments, and adapt all rules containing multiple subterms to use the splitting judgement. We detail T-APP in the figure; the adaptations of other rules are similar. While unrestricted types are useful in general, we show the specific case of unrestricted access points.

Access points. The **spawn** M construct spawns M as a new thread, **new**_S creates a fresh access point, and **request** M and **accept** M generate fresh endpoints that are matched up nondeterministically to form channels. With access points we can macro-express **fork**:

$$\mathbf{fork} M \triangleq \mathbf{let} ap = \mathbf{new}_S \mathbf{in} \mathbf{spawn} (M (\mathbf{accept} ap)); \mathbf{request} ap$$

883

Syntax

884

Types $A ::= \dots \mid \text{AP}(S)$

885

Access Point Names z

886

Terms $M ::= \dots \mid z \mid \text{spawn } M \mid \text{new}_S \mid \text{request } M \mid \text{accept } M$

887

Configurations $C ::= \dots \mid (vz)C \mid z(\mathcal{X}, \mathcal{Y})$

888

Runtime typing environments $\Delta ::= \dots \mid \Delta, z : S$

889

Splitting

$$\boxed{\Gamma = \Gamma_1 + \Gamma_2}$$

890

891

892

893

894

Typing

$$\boxed{\Gamma \vdash M : A}$$

895

T-VAR

$$\frac{x : A \in \Gamma \quad \text{un}(\Gamma)}{\Gamma \vdash x : A}$$

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918

919

920

921

922

923

924

925

926

927

928

929

930

931

Reduction

$$\boxed{C \longrightarrow D}$$

E-SPAWN

$$\mathcal{F}[\text{spawn } M] \longrightarrow \mathcal{F}[\cdot] \parallel \circ M$$

E-NEW

$$\mathcal{F}[\text{new}_S] \longrightarrow (vz)(\mathcal{F}[z] \parallel z(\epsilon, \epsilon))$$

 z is fresh

E-ACCEPT

$$\mathcal{F}[\text{accept } z] \parallel z(\mathcal{X}, \mathcal{Y}) \longrightarrow (va)(\mathcal{F}[a] \parallel z(\{a\} \cup \mathcal{X}, \mathcal{Y}))$$

 a is fresh

E-REQUEST

$$\mathcal{F}[\text{request } z] \parallel z(\mathcal{X}, \mathcal{Y}) \longrightarrow (va)(\mathcal{F}[a] \parallel z(\mathcal{X}, \{a\} \cup \mathcal{Y}))$$

 a is fresh

E-MATCH

$$z(\{a\} \cup \mathcal{X}, \{b\} \cup \mathcal{Y}) \longrightarrow z(\mathcal{X}, \mathcal{Y}) \parallel a(\epsilon) \leftrightarrow b(\epsilon)$$

Configuration Typing

$$\boxed{\Gamma; \Delta \vdash^\phi C}$$

910

911

912

913

914

915

916

917

918

919

920

921

922

923

924

925

926

927

928

929

930

931

TA-APNAME

$$\frac{\Gamma, z : \text{AP}(S); \Delta, z : S \vdash^\phi C}{\Gamma; \Delta \vdash^\phi (vz)C}$$

TA-AP

$$\frac{\text{un}(\Gamma)}{\Gamma, z : \text{AP}(S); \mathcal{X} : \bar{S}, \mathcal{Y} : S, z : S \vdash^\circ z(\mathcal{X}, \mathcal{Y})}$$

TA-CONNECTN

$$\frac{\Gamma = \Gamma_1 + \Gamma_2 \quad \Gamma_1, a : \bar{S}; \Delta_1, b : \bar{T} \vdash^{\phi_1} C \quad \Gamma_2, b : \bar{T}; \Delta_2, a : \bar{S} \vdash^{\phi_2} D}{\Gamma; \Delta_1, \Delta_2, a : S^\sharp, b : T^\sharp \vdash C \parallel D}$$

Fig. 9. Access Points

Reduction rules. We let z range over access point names. Configuration $(vz)C$ denotes binding access point name z in C , and $z(\mathcal{X}, \mathcal{Y})$ is an access point with name z and two sets \mathcal{X} and \mathcal{Y} containing endpoints to be matched.

Rule E-SPAWN creates a new child thread but, unlike **fork**, returns the unit value instead of creating a channel and returning an endpoint. Rule E-NEW creates a new access point with fresh name z . Rules E-ACCEPT and E-REQUEST create a fresh name a , returning the newly-created name to the thread, and adding the name to sets \mathcal{X} and \mathcal{Y} respectively. Rule E-MATCH matches two endpoints a and b contained in \mathcal{X} and \mathcal{Y} , and creates an empty buffer $a(\epsilon) \leftrightarrow b(\epsilon)$.

Configuration typing. Configuration typing judgements again have the shape $\Gamma; \Delta \vdash^\phi C$. Whereas Γ may contain unrestricted variables, Δ remains entirely linear.

Read bottom-up, rule TA-APNAME adds an unrestricted reference $z : AP(S)$ to Γ , and a linear entry $z : S$ to Δ . Rule TA-AP types an access point configuration. We write $X : S$ for $a_1 : S, \dots, a_n : S$, where $X = \{a_1, \dots, a_n\}$. For an access point $z(X, \mathcal{Y})$ to be well-typed, Δ must contain $z : S$, along with the names in X having type \bar{S} and the names in \mathcal{Y} having type S . Rule T-CONNECTN generalises T-CONNECT₁ and T-CONNECT₂ to allow any number of channels to communicate across a buffer; this therefore introduces the possibility of deadlock.

Interaction with cancellation. We need no additional reduction rules to account for interaction between access points and channel cancellation. Should an endpoint waiting to be matched be cancelled, it is paired as usual, and interaction with its associated buffer raises an exception:

$$\begin{aligned} \downarrow a \parallel \mathcal{F}[\text{receive } b] \parallel z(\{a\}, \{b\}) &\Longrightarrow \downarrow a \parallel \mathcal{F}[\text{receive } b] \parallel z(\epsilon, \epsilon) \parallel a(\epsilon) \rightsquigarrow b(\epsilon) \\ &\Longrightarrow \downarrow a \parallel \mathcal{F}[\text{raise}] \parallel \downarrow b \parallel z(\epsilon, \epsilon) \parallel a(\epsilon) \rightsquigarrow b(\epsilon) \end{aligned}$$

Metatheory. By decoupling process and channel creation we lose the guarantee that the communication topology is acyclic, and therefore introduce the possibility of deadlock. Preservation continues to hold—in fact, we gain a stronger preservation result since the use of TA-CONNECTN allows typeability to be preserved by equivalences.

THEOREM 4.1 (PRESERVATION MODULO EQUIVALENCE (ACCESS POINTS)). *If $\Psi; \Delta \vdash^\phi C$ and $C \Longrightarrow \mathcal{D}$, then there exist Ψ', Δ' such that $\Psi; \Delta \longrightarrow \Psi'; \Delta'$ and $\Psi'; \Delta' \vdash^\phi \mathcal{D}$.*

PROOF. By induction on the derivation of $C \longrightarrow \mathcal{D}$ and preservation by \equiv ; see Appendix D. \square

Alas, the introduction of cyclic topologies and therefore the loss of deadlock-freedom necessarily violates global progress. Nevertheless, a weaker form of progress still holds: if a configuration does not reduce, then it is due to deadlock rather than cancellation.

THEOREM 4.2 (PROGRESS (ACCESS POINTS)). *Suppose $;\cdot \vdash^\phi C$ and $C \not\Longrightarrow$. Then each thread in C is either a value; a buffer; a zapper thread; an access point; requesting or accepting on an access point; or ready to perform a communication action.*

If C contains a thread ϕM and $\text{ready}(a, M)$ for some name a , then C contains some buffer $a(\epsilon) \rightsquigarrow b(\vec{W})$, and C does not contain a zapper thread $\downarrow b$.

PROOF. We can prove a similar property for open configurations by induction on the derivation of $\Psi; \Delta \vdash^\phi C$; the above result arises as a corollary and by inspection of the reduction rules. \square

In the presence of access points confluence and termination no longer hold: access points are nondeterministic and can encode higher-order state and hence fixpoints via Landin’s knot.

4.3 Recursive Session Types

Recursive session types support repeating protocols. The extension of EGV with recursive session types is standard [Lindley and Morris 2016, 2017] and orthogonal to the main ideas of this paper, so we do not spell out the details here. The implementation (§5) does provide recursive session types.

5 SESSION TYPES WITHOUT TIERS

In this section we describe our extensions to Links to support exception handling, as well as extensions to the Links concurrency runtimes to support distribution. Links [Cooper et al. 2007] is a statically-typed, ML-inspired, impure functional programming language designed for the web. Links is designed to allow code for all “tiers” of a web application—client, server, and database—to be written in a single language. Lindley and Morris [2017] extend Links with first-class session types, relying on lightweight linear typing [Mazurak et al. 2010] and row polymorphism [Rémy

981 1994]. We extend their work to account for distributed web applications, which amongst other
982 things necessitates handling failure.

983

984

5.1 The Links Model

985

986

987

988

989

990

991

992

993

994

Links provides a uniform language for web applications. Client code is compiled to JavaScript, server code is interpreted, and database queries are compiled to SQL. Each client and server has its own concurrency runtime, providing lightweight processes and message passing communication. Earlier versions of Links [Cooper et al. 2007] invoked a fresh copy of the server per server request and communication between client and server was via RPC calls. Advances such as WebSockets allow socket-like bidirectional asynchronous communication between client and server, in turn allowing richer applications where data (for example, comments on a GitHub pull request) flows more freely between client and server. Moving to a model based on lightweight threads and session-typed channels avoids the inversion of control inherent in RPC-style systems, and allows development to be driven by the communication protocol.

995

996

997

998

Links now adopts a persistent application server model, incorporating client-server communication using session-typed channels. Since channels are a location-transparent abstraction, we also optionally allow the abstraction of client-to-client communication, routed through the server.

999

5.2 Concurrency

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

Links provides typed actor-style concurrency where processes have a single incoming message queue and can send asynchronous messages. Lindley and Morris [2017] extend Links with session-typed channels, using Links' process-based model but replacing actor mailboxes with session-typed channels. We extend their implementation to support distribution and failure handling.

1014

1015

1016

1017

1018

1019

The client relies on continuation-passing style (CPS), trampolining, and co-operative threading. Client code is compiled to CPS, and explicit `yield` instructions are inserted at every function application. When a process has yielded a given number of times, the continuation is pushed to the back of a queue, and the next process is pulled from the front of the queue. While modern browsers are beginning to integrate tail-recursion, and we have updated the Links library to support it, adoption is not yet widespread. Thus, we periodically discard the call stack using a trampoline. Cooper [2009] discusses the Links client concurrency model in depth. The server implements concurrency on top of the OCaml `lwt` library [Vouillon 2008], which provides lightweight co-operative threading. At runtime, a channel is represented as a pair of endpoint identifiers:

(Peer endpoint, Local endpoint)

1016

1017

Endpoint identifiers are unique. If a channel (a, b) exists at a given location, then that location should contain a buffer for b .

1018

1019

5.3 Distributed Communication

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

To support bidirectional communication between client and server we use WebSockets [Fette and Melnikov 2011]. A WebSocket connection is established by a client. When a request is made and a web page is generated, each client is assigned a unique identifier, which it uses to establish a WebSocket connection. Any messages the server attempts to send prior to a WebSocket connection being established are buffered and delivered after the connection is established. We use a JSON protocol to communicate messages such as access point operations, remote session messages, and endpoint cancellation notifications.

It is possible that one client will hold one endpoint of a channel, and another client will hold the other endpoint. In order to provide the illusion of client-to-client communication, we route the

communication between the two clients via the server. The server maintains a map

$$\text{Endpoint ID} \mapsto \text{Location}$$

where `Location` is either `Server` or `Client (ID)`, where `ID` identifies a particular client. The map is updated if a new connection is established; an endpoint is sent as part of a message; or a client disconnects. The server also maintains a map

$$\text{Client ID} \mapsto [\text{Channel}]$$

associating each client with the publicly-facing channels residing on that client, where `Channel` is a pair of endpoints (a, b) such that b is the endpoint residing on the client. Much like TCP connections, WebSocket connections raise an event when a connection is disconnected. Upon receiving such an event, all channels associated with the client are cancelled, and exceptions are invoked as per the exception handling mechanism described in §2 and §5.4.

Distributed Delegation. It is possible to send endpoints as part of a message. Session delegation in the presence of distributed communication requires some care to ensure that messages are delivered to the correct participant; our implementation adapts the algorithms of Hu et al. [2008]. Further details can be found in Appendix E.

5.4 Session Typing with Failure Handling

Effect Handlers. Effect handlers [Plotkin and Pretnar 2013] provide a modular approach to programming with user-defined effects. Exception handlers are a special case of effect handlers. Consequently, we leverage the existing implementation of effect handlers in Links [Hillerström and Lindley 2016; Hillerström et al. 2017]. In §4 we generalise `try – as – in – otherwise` to accommodate user defined exceptions. Effect handlers generalise further to support what amounts to *resumable exceptions* in which the handler has access not only to a payload, but also the delimited continuation (i.e. evaluation context) from the point at which the exception was raised up to the handler, allowing effect handlers to implement arbitrary side-effects; not just exceptions. We translate exception handling as follows.

$$\begin{aligned} \llbracket \text{raise} \rrbracket &= \text{do raise} & \llbracket \text{try } L \text{ as } x \text{ in } M \text{ otherwise } N \rrbracket &= \text{handle } \llbracket L \rrbracket \text{ with} \\ & & & \text{return } x \mapsto \llbracket M \rrbracket \\ & & & \text{raise } r \mapsto \text{cancel } r; \llbracket N \rrbracket \end{aligned}$$

The introduction form `do op` invokes an operation `op` (which may represent raising an exception or some other effect). The elimination form `handle M with H` runs effect handler `H` on the computation `M`. In general an effect handler `H` consists of a *return clause* of the form `return x ↦ N`, which behaves just like the success continuation $(x \text{ in } N)$ of an exception handler, and a collection of *operation clauses*, each of the form `op \vec{p} r ↦ N`, specifying how to handle an operation analogously to how exception handler clauses specify how to handle an exception, except that as well as binding payload parameters \vec{p} , an operation clause also binds a *resumption* parameter r . The resumption r binds a closure representing the continuation up to the nearest enclosing effect handler, allowing control to pass back to the program after handling the effect. In the case of our translation, the `raise` operation has no payload, and rather than invoking the resumption r we cancel it, assuming the natural extension of cancellation to arbitrary linear values, whereby all free names in the value are cancelled (r being bound to the current evaluation context reified as a value). A formalisation of linear effect handlers for session typing is outside the scope of this paper and left as future work.

1079 *Raising exceptions.* An exception may be raised either explicitly through an invocation of **raise**
1080 (desugared to **do raise**), or through a blocked **receive** call where the peer endpoint has been
1081 cancelled. Thus, we know statically where any exceptions may be raised. To support cancella-
1082 tion of closures on the client, we adorn closures with an explicit environment field that can be
1083 directly inspected. Currently, Links does not closure-convert continuations on the client, so we
1084 use a workaround to simulate cancelling a resumption (as required by the translation $\llbracket - \rrbracket$). When
1085 compiling client code, for each occurrence of **do raise**, we compile a function that inspects all
1086 affected variables and cancels any affected endpoints in the continuation. For each occurrence of
1087 **receive**, we compile a continuation to cancel affected endpoints to be invoked by the runtime
1088 system if the receive operation fails.

1089

1090 5.5 Distributed Exceptions

1091 Our implementation fully supports the semantics described in §2. The concurrency runtime at each
1092 location maintains a set of cancelled endpoints.

1093

1094 *Cancellation.* Suppose endpoint a is connected to peer endpoint b . If a is cancelled, then all
1095 endpoints in the queue for a are also cancelled according to the E-ZAP rule. If a and b are at the
1096 same location, then a is added to the set of cancelled endpoints. If they are at different locations,
1097 then a cancellation notification for a is routed to b 's location. Zapper threads are modelled in the
1098 implementation by recording sets of cancelled endpoints and propagating cancellation messages.

1099

1100 *Failed communications.* Again, suppose endpoint a is connected to peer endpoint b . Should a
1101 process attempt to read from a when the buffer for a is empty, then the runtime will check to see
1102 whether b is in the set of cancelled endpoints. If so, then a is cancelled and an exception is raised in
1103 the blocked process; if not, the process is suspended until a message is ready. Should the runtime
1104 later add b to the set of cancelled endpoints, then again a is cancelled and an exception raised.
1105 These actions implement the E-RECEIVEZAP rule.

1106

1107 *Disconnection.* To handle disconnection, the server maintains a map from client IDs to the list of
1108 endpoints at the associated client. WebSockets—much like TCP sockets—raise a *closed* event on
1109 disconnection. Consequently, when a connection is closed, the runtime looks up the endpoints
1110 owned by the terminated client and notifies all other clients containing the peer endpoints.

1111

1112 6 EXAMPLE: A CHAT APPLICATION

1113 In this section we outline the design and implementation of a web-based chat application in Links
1114 making use of distributed session-typed channels. We write the following informal specification:

1115

- 1116 • To initialise, a client must:
 - 1117 – connect to the chat server; then
 - 1118 – send a nickname; then
 - 1119 – receive the current topic and list of nicknames.
- 1120 • After initialisation the client is connected and can:
 - 1121 – send a chat message to the room; or
 - 1122 – change the room's topic; or
 - 1123 – receive messages from other users; or
 - 1124 – receive changes of topic from other users.
- 1125 • Clients cannot connect with a nickname that is already in use in the room.
- 1126 • All participants should be notified whenever a participant joins or leaves the room.

1127

```

1128
1129 typename ChatClient = !Nickname.
1130   [&| Join:
1131     ?(Topic, [Nickname], ClientReceive).ClientSend,
1132     Nope:End |&];
1133
1134 typename ClientReceive =
1135   [&| Join      : ?Nickname      .ClientReceive,
1136     Chat       : ?(Nickname, Message).ClientReceive,
1137     NewTopic   : ?Topic         .ClientReceive,
1138     Leave      : ?Nickname      .ClientReceive
1139   |&];
1140
1141 typename ClientSend =
1142   [+| Chat     : ?Message.ClientSend,
1143     Topic     : ?Topic .ClientSend |+];
1144
1145 typename ChatServer = ~ChatClient;
1146 typename WorkerSend = ~ClientReceive;
1147 typename WorkerReceive = ~ClientSend;

```

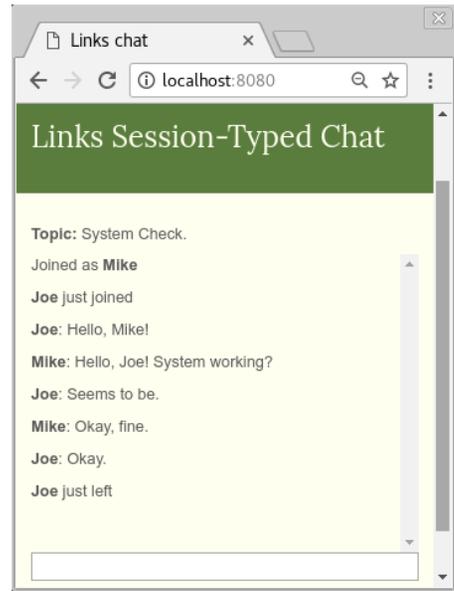


Fig. 10. Chat Application Session Types

Session Types. We can encode much of the specification more precisely as a session type, as shown in Figure 10. The client begins by sending a nickname, and then offers the server a choice of a `Join` message or a `Nope` message. In the former case, the client then receives a triple containing the current topic, a list of existing nicknames, and an endpoint (of type `ClientReceive`) for receiving further updates from the server; and may then continue to send messages to the server as a connected client endpoint (of type `ClientSend`). (Observe the essential use of session delegation.) In the latter case, communication is terminated. The intention is that the server will respond with `Nope` if a client with the supplied nickname is already in the chat room (the details of this check are part of the implementation, not part of the communication protocol).

The `ClientReceive` endpoint allows the client to offer a choice of four different messages: `Join`, `Chat`, `NewTopic`, or `Leave`. In each case the client then receives a payload (depending on the choice, a nickname, pair of nickname and chat message, or topic change) before offering another choice. The `ClientSend` endpoint allows the client to select between two different messages: `Chat` and `NewTopic`. In each case the client subsequently sends a payload (a chat message or a new topic) before selecting another choice. The chat server communicates with the client along endpoints with dual types.

How can session types help? The connect function (Fig. 11a) is run when a client enters a nickname. First, the client requests a fresh channel of type `ChatClient` from access point `wap` of type `AP(ChatServer)`. Next, the client obtains the content of the DOM input box for the nickname by calling `getInputContents(nameBoxId)`, where `nameBoxId` is the DOM ID for the nickname entry box. Next, the client sends the nickname to the server and waits for a response; in the case of a `Join` message, the client receives the room data and an incoming message channel, and calls the `beginChat` function. In the case of a `Nope` message, an error is printed and the session ends.

Now, suppose the developer forgets to write code to check the server response (Fig. 11b). This implementation is incorrect since there is a *communication mismatch*: the server is expecting to

```

1177 fun connect() {
1178   var s = request(wap);
1179   var nick = getInputContents(nameBoxId);
1180   var s = send(nick, s);
1181   offer(s) {
1182     case Join(s) ->
1183       var ((topic, nicks, incoming), s) =
1184         receive(s);
1185       beginChat(topic, nicks, incoming, s)
1186     case Nope(s) ->
1187       print("Nickname '" ^ nick ^ "' already taken")
1188   }
1189 }

```

(a) Correct connect function

```

1177 fun connect() {
1178   var s = request(wap);
1179   var nick = getInputContents(nameBoxId);
1180   var s = send(nick, s);
1181   offer(s) {
1182     case Join(s) ->
1183       var ((topic, nicks, incoming), s) =
1184         receive(s);
1185       beginChat(topic, nicks, incoming, s)
1186     case Nope(s) ->
1187       print("Nickname '" ^ nick ^ "' already taken")
1188   }
1189 }

```

(b) Incorrect connect function

Fig. 11. Implementations of connect function

accept or reject the request to join the room, whereas the client is expecting to receive data about the room. However, since s has type `ChatClient` but does not follow the protocol, Links catches the communication mismatch statically. Similarly, Links will statically detect an unused endpoint (e.g. the developer forgets to finish a protocol) or an endpoint being used more than once, as in §1.2.

Architecture. Figure 12a depicts the architecture of the chat application. Each client has a process which sends messages over a distributed session channel of type `ClientSend` to its own worker process on the server, which in turn sends internal messages to a supervisor process containing the state of the chat room. These messages trigger the supervisor process to broadcast a message to all chat clients over a channel of type `~ClientReceive`. As is evident from the figure, the communication topology is cyclic; in order to construct this topology the code makes essential use of access points.

Disconnection. Figure 12b shows the implementation of a worker process which receives messages from a client. The worker takes the nickname of the client, as well as a channel endpoint of type `WorkerReceive` (which is the dual of `ClientSend`). The server offers the client a choice of sending a message (`Chat`), or changing topic (`NewTopic`); in each case, the associated data is received and a message dispatched to the supervisor process by calling `chat` or `newTopic`. When a client closes its connection to the server, all associated endpoints are cancelled. Consequently, an exception will be raised when evaluating the `offer` or `receive` expressions. To handle disconnection, we wrap the function in an exception handler, which recursively calls `worker` if the interaction is successful, and notifies the supervisor that the user has left via a call to `leave` if an exception is raised.

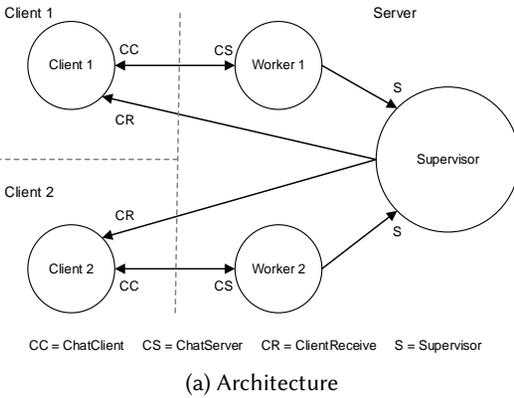
Additional examples. We have concentrated on the chat server example for exposition, but have also implemented an extended chat server and a multiplayer game. These can be found at <http://www.github.com/SimonJF/distributed-links-examples>.

7 RELATED WORK

7.1 Session Types with Failure Handling

Carbone et al. [2008] provide the first formal basis for exceptions in a session-typed process calculus. Our approach provides significant simplifications: zipper threads provide a simpler semantics and remove the need for their queue levels, meta-reduction relation, and liveness protocol.

1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274



```
sig worker : (Nickname, WorkerReceive) -> ()
fun worker(nick, c) {
  try {
    offer(c) {
      case Chat(c) ->
        var (msg, c) = receive(c);
        chat(nick, msg); c
      case NewTopic(c) ->
        var (topic, c) = receive(c);
        newTopic(topic); c
    }
  } as (c) in {
    worker(nick, c)
  } otherwise { leave(nick) }
}
```

(b) Worker Implementation

Fig. 12. Chat Application Architecture and Worker Implementation

Our work draws on that of [Mostrous and Vasconcelos \[2014\]](#), who introduce the idea of cancellation. Our work differs from theirs in several key ways. Their system is a process calculus; ours is a λ -calculus. Their channels are synchronous; ours are asynchronous. Their exception handling construct scopes over a single action; ours scopes over an arbitrary computation.

[Caires and Pérez \[2017\]](#) describe a core, logically-inspired process calculus supporting non-determinism and abortable behaviours encoded via a nondeterminism modality. Processes may either provide or not provide a prescribed behaviour; if a process attempts to consume a behaviour that is not provided, then its linear continuation is safely discarded by propagating the failure of sessions contained within the continuation. Their approach is similar in spirit to our zipper threads. Additionally, they give a core λ -calculus with abortable behaviours and exception handling, and define a type-preserving translation into their core process calculus.

Our approach differs in several important ways. First, our semantics is asynchronous, handling the intricacies involved with cancelling values contained in message queues. Second, we give a direct semantics to EGV, whereas [Caires and Pérez](#) rely on a translation into their underlying process calculus. Third, to handle the possibility of disconnection, our calculus allows *any* channel to be discarded, whereas they opt for an approach more closely resembling checked exceptions, aided by a monadic presentation.

The above works are all theoretical. Backed by our theoretical development, our implementation integrates session types and exceptions, extending Links.

Multiparty Session Types. [Fowler \[2016\]](#) describes an Erlang implementation of the Multiparty Session Actor framework proposed by [Neykova and Yoshida \[2014, 2017b\]](#) with a limited form of failure recovery; [Neykova and Yoshida \[2017a\]](#) present a more comprehensive approach, based on refining existing Erlang supervision strategies. [Chen et al. \[2016\]](#) introduce a formalism based on multiparty session types [[Honda et al. 2016](#)] that handles partial failures by transforming programs to detect possible failures at a set of statically determined synchronisation points. These approaches rely on a fixed communication topology, using mechanisms such as dependency graphs or synchronisation points to determine which participants are affected when one participant fails. Delegation implies location transparency, thus we must consider dynamic topologies.

7.2 Session Types and Distribution

Hu et al. [2008] introduce Session Java (SJ), which allows distributed session-based communication in the Java programming language, making use of the Polyglot framework [Nystrom et al. 2003] to statically check session types. Hu et al. are the first to present the challenges of distributed delegation along with distributed algorithms which address those challenges. We adapt their algorithms to web applications. SJ restricts communication to a fixed set of simple types; Links allows arbitrary values to be sent. SJ provides statically scoped exception handling, propagating exceptions to ensure liveness (but this feature is not formalised).

Scalas and Yoshida [2016] introduce `lchannels`, a lightweight implementation of session types in Scala. To maximise applicability of their approach and not require any modifications to Scala, their approach detects duplicate endpoint use at runtime. By virtue of the translation into the linear π -calculus introduced by Kobayashi [2003] and later expanded on by Dardha et al. [2017], `lchannels` is particularly amenable to distribution. Scalas et al. [2017] build upon this approach to translate a multiparty session calculus into the linear π -calculus, providing the first distributed implementation of multiparty session types to support delegation.

7.3 Session Types via Affine Types

Rust [Matsakis and Klock II 2014] provides *ownership types* [Clarke 2003], ensuring that an object has at most one owner. Jespersen et al. [2015] use Rust's ownership types to encode affine session types, but since affine endpoints can be discarded implicitly, their library does not guarantee progress. Although it is not possible to distinguish between dynamic failure and a developer forgetting to finish an implementation, our semantics can be implemented using Rust's destructor mechanism, enabling a progress property [Kokke 2018].

8 CONCLUSION AND FUTURE WORK

Session types allow protocol conformance to be checked statically. The prevailing consensus has hitherto been to require that endpoints be used linearly to enforce session fidelity and prevent premature discarding of open channels. We have argued that in order to write realistic applications in the presence of distribution and failure, linearity should be supplemented with an *explicit* cancellation operation. We show that, even in the presence of channel cancellation, our core calculus is well-behaved, being deadlock-free, type sound, confluent, and terminating.

In tandem with the formal development, we have developed an extension of the Links programming language to support distributed session-based communication for web applications, thus providing the first implementation of asynchronous session types with failure handling in a functional programming language. Our implementation leverages recent work on effect handlers.

Future work. Our implementation combines linearity and effect handlers. Linear effect handlers are new, and a ripe area of study in their own right; we plan to formalise session-typed concurrency and failure handling directly in terms of linear effect handlers. Multiparty session types [Honda et al. 2016] are yet to be included as a first-class construct of a core functional language. A natural starting point is to identify a λ -calculus into which we can translate the MCP calculus of Carbone et al. [2016] and then investigate how our approach adapts to the multiparty setting.

ACKNOWLEDGMENTS

Thanks to James McKinna and the anonymous reviewers for detailed comments and suggestions. This work was supported by EPSRC grants EP/L01503X/1 (EPSRC CDT in Pervasive Parallelism) and EP/K034413/1 (From Data Types to Session Types—A Basis for Concurrency and Distribution), and an LFCS internship.

REFERENCES

- 1324
1325 H. P. Barendregt. 1984. *The Lambda Calculus Its Syntax and Semantics* (revised ed.). Vol. 103. North Holland.
- 1326 Nick Benton and Andrew Kennedy. 2001. Exceptional Syntax. *Journal of Functional Programming* 11, 4 (2001), 395–410.
- 1327 Luís Caires and Jorge A Pérez. 2017. Linearity, control effects, and behavioral types. In *ESOP*. Springer, 229–259.
- 1328 Luís Caires and Frank Pfenning. 2010. Session types as intuitionistic linear propositions. In *CONCUR*, Vol. 10. Springer, 222–236.
- 1329 Marco Carbone, Ornela Dardha, and Fabrizio Montesi. 2014. Progress as compositional lock-freedom. In *COORDINATION*. Springer, 49–64.
- 1330 Marco Carbone, Kohei Honda, and Nobuko Yoshida. 2008. Structured interactional exceptions in session types. In *CONCUR*. Springer, 402–417.
- 1331 Marco Carbone, Sam Lindley, Fabrizio Montesi, Carsten Schürmann, and Philip Wadler. 2016. Coherence generalises duality: A logical explanation of multiparty session types. In *CONCUR (LIPICs)*, Vol. 59. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 33:1–33:15.
- 1332 Tzu-Chun Chen, Malte Viering, Andi Bejleri, Lukasz Ziarek, and Patrick Eugster. 2016. A type theory for robust failure handling in distributed systems. In *FORTE (Lecture Notes in Computer Science)*, Vol. 9688. Springer, 96–113.
- 1333 David Gerard Clarke. 2003. *Object Ownership and Containment*. Ph.D. Dissertation. New South Wales, Australia. AAI0806678.
- 1334 Ezra Cooper. 2009. *Programming Language Features for Web Application Development*. Ph.D. Dissertation. University of Edinburgh.
- 1335 Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. 2007. Links: Web programming without tiers. In *FMCO*. Springer, 266–296.
- 1336 Ornela Dardha, Elena Giachino, and Davide Sangiorgi. 2017. Session types revisited. *Inf. Comput.* 256 (2017), 253–286.
- 1337 Ian Fette and Alexey Melnikov. 2011. *The WebSocket Protocol*. RFC 6455. RFC Editor. 70 pages. <http://www.rfc-editor.org/rfc/rfc6455.txt>
- 1338 Simon Fowler. 2016. An Erlang implementation of multiparty session actors. In *ICE (EPTCS)*, Vol. 223. 36–50.
- 1339 Simon J Gay and Vasco T Vasconcelos. 2010. Linear type theory for asynchronous session types. *Journal of Functional Programming* 20, 1 (2010), 19–50.
- 1340 Daniel Hillerström and Sam Lindley. 2016. Liberating effects with rows and handlers. In *TyDe@ICFP*. ACM, 15–27.
- 1341 Daniel Hillerström, Sam Lindley, Robert Atkey, and K. C. Sivaramakrishnan. 2017. Continuation passing style for effect handlers. In *FSCD (LIPICs)*, Vol. 84. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 18:1–18:19.
- 1342 Kohei Honda. 1993. Types for dyadic interaction. In *CONCUR*. Springer, 509–523.
- 1343 Kohei Honda, Vasco T Vasconcelos, and Makoto Kubo. 1998. Language primitives and type discipline for structured communication-based programming. In *ESOP*. Springer, 122–138.
- 1344 Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2016. Multiparty asynchronous session types. *Journal of the ACM (JACM)* 63, 1 (2016), 9.
- 1345 Raymond Hu, Nobuko Yoshida, and Kohei Honda. 2008. Session-based distributed programming in java. In *ECOOP*. Springer, 516–541.
- 1346 Thomas Bracht Laumann Jespersen, Philip Munksgaard, and Ken Friis Larsen. 2015. Session types for Rust. In *WGP*. ACM, 13–22.
- 1347 Naoki Kobayashi. 2003. *Type Systems for Concurrent Programs*. Springer Berlin Heidelberg, Berlin, Heidelberg, 439–453.
- 1348 Wen Kokke. 2018. `rusty-variation`: a library for deadlock-free session-typed communication in Rust. <https://github.com/wenkokke/rusty-variation>. (2018).
- 1349 Sam Lindley and J. Garrett Morris. 2015. A semantics for propositions as sessions. In *ESOP (Lecture Notes in Computer Science)*, Vol. 9032. Springer, 560–584.
- 1350 Sam Lindley and J Garrett Morris. 2016. Talking bananas: structural recursion for session types. In *ICFP*. ACM, 434–447.
- 1351 Sam Lindley and J Garrett Morris. 2017. Lightweight functional session types. In *Behavioural Types: from Theory to Tools*. River Publishers, 265–286.
- 1352 Nicholas D. Matsakis and Felix S. Klock II. 2014. The Rust language. In *HILT*. ACM, 103–104.
- 1353 Karl Mazurak, Jianzhou Zhao, and Steve Zdancewic. 2010. Lightweight linear types in System F^{*}. In *TLDI*. ACM, 77–88.
- 1354 Robin Milner. 1999. *Communicating and mobile systems: the pi calculus*. Cambridge university press.
- 1355 Dimitris Mostros and Vasco Thudichum Vasconcelos. 2014. Affine Sessions. In *COORDINATION*. Springer, 115–130.
- 1356 Romyana Neykova and Nobuko Yoshida. 2014. Multiparty session actors. In *COORDINATION (Lecture Notes in Computer Science)*, Vol. 8459. Springer, 131–146.
- 1357 Romyana Neykova and Nobuko Yoshida. 2017a. Let it recover: multiparty protocol-induced recovery. In *CC*. ACM, 98–108.
- 1358 Romyana Neykova and Nobuko Yoshida. 2017b. Multiparty session actors. *Logical Methods in Computer Science* 13, 1 (2017).
- 1359 Nathaniel Nystrom, Michael Clarkson, and Andrew Myers. 2003. Polyglot: An extensible compiler framework for java. In *CC*. Springer, 138–152.
- 1360 Luca Padovani. 2017. A simple library implementation of binary sessions. *Journal of Functional Programming* 27 (2017), e4.

1372

- 1373 Jorge A Pérez, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2012. Linear logical relations for session-based
1374 concurrency. In *European Symposium on Programming*. Springer, 539–558.
- 1375 Gordon D. Plotkin and Matija Pretnar. 2013. Handling algebraic effects. *Logical Methods in Computer Science* 9, 4 (2013).
- 1376 Didier Rémy. 1994. Type inference for records in a natural extension of ML. In *Theoretical Aspects Of Object-Oriented
Programming*, Carl A. Gunter and John C. Mitchell (Eds.). MIT Press, Cambridge, MA, 67–95.
- 1377 Alceste Scalas, Ornela Dardha, Raymond Hu, and Nobuko Yoshida. 2017. A linear decomposition of multiparty sessions
1378 for safe distributed programming. In *ECOOP (LIPICs)*, Vol. 74. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik,
1379 24:1–24:31.
- 1380 Alceste Scalas and Nobuko Yoshida. 2016. Lightweight session programming in scala. In *ECOOP (LIPICs)*, Vol. 56. Schloss
1381 Dagstuhl - Leibniz-Zentrum fuer Informatik, 21:1–21:28.
- 1382 Jérôme Vouillon. 2008. Lwt: a cooperative thread library. In *ML*. ACM, 3–12.
- 1383 Philip Wadler. 2014. Propositions as sessions. *Journal of Functional Programming* 24, 2-3 (2014), 384–418.
- 1384
- 1385
- 1386
- 1387
- 1388
- 1389
- 1390
- 1391
- 1392
- 1393
- 1394
- 1395
- 1396
- 1397
- 1398
- 1399
- 1400
- 1401
- 1402
- 1403
- 1404
- 1405
- 1406
- 1407
- 1408
- 1409
- 1410
- 1411
- 1412
- 1413
- 1414
- 1415
- 1416
- 1417
- 1418
- 1419
- 1420
- 1421

1422 **APPENDIX CONTENTS**

1423			
1424	A	Example Runtime Typing Derivation	31
1425	B	Deadlock-freedom	32
1426	C	Supplement to Section 3 (Metatheory of EGV)	33
1427	C.1	Preservation	33
1428	C.2	Canonical Forms	46
1429	C.3	Progress	46
1430	C.4	Confluence	50
1431	D	Supplement to Section 4.1 (Metatheory of EGV with Access Points)	51
1432	E	Distributed Delegation	54
1433	E.1	Challenges of Distributed Delegation	54
1434	E.2	Approaches to Distributed Delegation	55
1435	E.3	Delegation in Distributed Session Links	55
1436	E.4	Correctness	56
1437			
1438			
1439			
1440			
1441			
1442			
1443			
1444			
1445			
1446			
1447			
1448			
1449			
1450			
1451			
1452			
1453			
1454			
1455			
1456			
1457			
1458			
1459			
1460			
1461			
1462			
1463			
1464			
1465			
1466			
1467			
1468			
1469			
1470			

1471 A EXAMPLE RUNTIME TYPING DERIVATION

1472 We give an example derivation to illustrate how channels are introduced by name restrictions
 1473 and then split into endpoints using the $T\text{-CONNECT}_i$ rules. We assume suitable encodings of linear
 1474 booleans and integers using linear sums and products.

1475 Let us assume we have derivations for:

1476 $\Gamma_1, a : !\text{Int.End} \vdash^\circ E[\text{send } 5 \ a] : 1$ $\Gamma_2, b : ?\text{Bool}.\text{?Int.End} \vdash E'[\text{receive } b] : A$ $\cdot \vdash \text{true} : \text{Bool}$
 1477

1478 We construct a derivation \mathbf{D} of $(\nu a)(\nu b)(\circ E[\text{send } 5 \ a] \parallel (a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b]))$.

1479 First let \mathbf{D}_1 be the following subderivation.

$$1480 \frac{1481 \frac{1482 \frac{1483 \text{?Int.End}/\epsilon = !\text{Bool}.\text{!Int.End}/\text{Bool} \quad \cdot \vdash \epsilon : \epsilon \quad \cdot \vdash \text{true} : \text{Bool}}{\cdot ; a : ?\text{Int.End}, b : !\text{Bool}.\text{!Int.End} \vdash^\circ a(\epsilon) \rightsquigarrow b(\text{true})} \text{T-BUFFER}}{\cdot ; a : ?\text{Int.End}, b : (?\text{Bool}.\text{?Int.End})^\sharp \vdash^\bullet a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b]} \text{T-MAIN}}{\Gamma_2; a : ?\text{Int.End}, b : (?\text{Bool}.\text{?Int.End})^\sharp \vdash^\bullet a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b]} \text{T-CONNECT}_2$$

1484 Then let \mathbf{D}_2 be the following subderivation.

$$1485 \frac{1486 \frac{1487 \frac{1488 \Gamma_2, b : ?\text{Bool}.\text{?Int.End} \vdash E'[\text{receive } b] : A}{\Gamma_2, b : ?\text{Bool}.\text{?Int.End}; \cdot \vdash^\bullet \bullet E'[\text{receive } b]} \text{T-MAIN}}{\Gamma_2; a : ?\text{Int.End}, b : (?\text{Bool}.\text{?Int.End})^\sharp \vdash^\bullet a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b]} \text{T-CONNECT}_2}{\Gamma_1, a : !\text{Int.End} \vdash^\circ E[\text{send } 5 \ a] : 1} \text{T-THREAD}$$

1489 The complete derivation \mathbf{D} is as follows.

$$1490 \frac{1491 \frac{1492 \frac{1493 \frac{1494 \frac{1495 \frac{1496 \Gamma_1, a : !\text{Int.End} \vdash^\circ E[\text{send } 5 \ a] : 1}{\Gamma_1, a : !\text{Int.End}; \cdot \vdash^\circ E[\text{send } 5 \ a]} \text{T-THREAD}}{\Gamma_1, \Gamma_2; a : (!\text{Int.End})^\sharp, b : (?\text{Bool}.\text{?Int.End})^\sharp \vdash^\circ \circ E[\text{send } 5 \ a] \parallel (a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b])} \text{T-CONNECT}_1}{\Gamma_1, \Gamma_2; a : (!\text{Int.End})^\sharp \vdash^\bullet (\nu a)(\nu b)(\circ E[\text{send } 5 \ a] \parallel (a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b]))} \text{T-NU}}{\Gamma_1, \Gamma_2; \cdot \vdash^\bullet \bullet (\nu a)(\nu b)(\circ E[\text{send } 5 \ a] \parallel (a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b]))} \text{T-NU}}{\Gamma_1, \Gamma_2; \cdot \vdash^\bullet \bullet (\nu a)(\nu b)(\circ E[\text{send } 5 \ a] \parallel (a(\epsilon) \rightsquigarrow b(\text{true}) \parallel \bullet E'[\text{receive } b]))} \text{T-NU}$$

1497 Let us read \mathbf{D} bottom-upwards. The two instances of the $T\text{-NU}$ rule introduce channels a and b
 1498 into the runtime environment. The $T\text{-CONNECT}_1$ rule splits channel a into dual endpoints: on the
 1499 left the endpoint a appears in the type environment and the sending thread; on the right the end
 1500 point a appears in the runtime environment and the buffer. The $T\text{-CONNECT}_2$ rule splits channel b
 1501 into dual endpoints: on the left the endpoint b appears in the runtime environment and the buffer;
 1502 on the right the endpoint b appears in the type environment and the receiving thread.

B DEADLOCK-FREEDOM

Here we give a graph-theoretic account of deadlock-freedom in EGV, independent of our notion of progress, following Lindley and Morris [2015].

Due to the asynchronous semantics of EGV, sending on an endpoint and cancelling an endpoint reduce immediately. Deadlocks may therefore only occur when cycles occur receiving or closing an endpoint. We begin by classifying the notion of a *blocked thread*: that is, a thread which is blocked performing an action on some channel endpoint.

Definition B.1. We say that term M is *blocked on name a* if M is about to receive on or close a . Formally:

$$\text{blocked}(a, M) \triangleq \exists E. (M = E[\mathbf{receive} \ a]) \vee (M = E[\mathbf{close} \ a])$$

Given the notion of a blocked thread, we may characterise the notion of a dependency between communication actions.

Definition B.2. Let C be a configuration such that a and b are not bound by C . We say that a *depends on b in C* , written $\text{depends}(a, b, C)$, if C is a buffer connecting a and b , or a appears in some thread blocked on b , or if a depends on some name c which depends on b . Formally:

- $\text{depends}(a, b, a(\vec{V}) \leftrightarrow b(\vec{W}))$
- $\text{depends}(a, b, b(\vec{W}) \leftrightarrow a(\vec{V}))$
- $\text{depends}(a, b, \phi M) \triangleq \text{blocked}(b, M) \wedge a \in \text{fn}(M)$
- $\text{depends}(a, b, C) \triangleq \exists \mathcal{G}, \mathcal{D}, \mathcal{E}, c. C \equiv \mathcal{G}[\mathcal{D} \parallel \mathcal{E}] \wedge \text{depends}(a, c, \mathcal{D}) \wedge \text{depends}(c, b, \mathcal{E})$

Remark. The above definition of dependency is an over-approximation to the intuitive notion, as a buffer need not have dependencies in both directions, but for our purposes this does not matter.

Definition B.3. We say that a configuration is *deadlocked* if it contains cyclic dependencies:

$$\text{deadlocked}(C) \triangleq \exists \mathcal{D}, \mathcal{E}, a, b. C \equiv \mathcal{G}[\mathcal{D} \parallel \mathcal{E}] \wedge \text{depends}(a, b, \mathcal{D}) \wedge \text{depends}(b, a, \mathcal{E})$$

With these definitions in place, we can show that EGV configurations are deadlock-free.

LEMMA B.4. *If $\text{depends}(a, b, C)$ then $a, b \in \text{fn}(C)$.*

PROOF. By induction on the definition of $\text{depends}(a, b, C)$. □

THEOREM B.5. *If $\Gamma; \Delta \vdash C$, then $\neg \text{deadlocked}(C)$.*

PROOF. By contradiction. Suppose $\text{deadlocked}(C)$, that is:

$$\exists \mathcal{D}, \mathcal{E}, a, b. C \equiv \mathcal{G}[\mathcal{D} \parallel \mathcal{E}] \wedge \text{depends}(a, b, \mathcal{D}) \wedge \text{depends}(b, a, \mathcal{E})$$

Thus by Lemma B.4, $a, b \in \text{fn}(\mathcal{D})$ and $b, a \in \text{fn}(\mathcal{E})$. Then by Lemma 3.6, C must be ill-typed. □

Remark. We regard blocked threads as deadlocked only if there is a cyclic dependency. It is perfectly possible for a configuration to include blocked threads without there being a deadlock.

- Deadlock-free open terms can block on external communication along a free endpoint.
- Deadlock-free closed terms can block on communication along an endpoint that appears in the return value of a program. This also amounts to being blocked on external communication.

All blocked threads can be ruled out by restricting the type of a program to be free of both session types and function types (the latter is necessary as closures can capture endpoints).

C SUPPLEMENT TO SECTION 3 (METATHEORY OF EGV)

C.1 Preservation

In this section, we present proofs that typeability is preserved by configuration reduction.

C.1.1 Equivalence. We begin by describing the properties of configuration equivalence. As described in §3, typeability of configurations is *not* preserved by equivalence. Nonetheless, Lemma C.1 shows that only the associativity of parallel composition may cause a configuration to be ill-typed.

LEMMA C.1. *If $\Gamma; \Delta \vdash^\phi C$ and $C \equiv \mathcal{D}$, where the derivation of $C \equiv \mathcal{D}$ does not contain a use of the axiom for associativity, then $\Gamma; \Delta \vdash^\phi \mathcal{D}$.*

PROOF. By induction on the derivation of $C \equiv \mathcal{D}$, examining the equivalence in both directions to account for symmetry. We show that a typing derivation of the left-hand side of an equivalence rule implies the existence of the right-hand side, and vice versa.

That reflexivity, transitivity, and symmetry of the equivalence relation respect typing follows immediately because equality of typing derivations is an equivalence relation.

We make implicit use of the induction hypothesis.

Congruence rules

Case Name restriction

$$\frac{C \equiv \mathcal{D}}{(va)C \equiv (va)\mathcal{D}}$$

$$\frac{\Gamma; \Delta, a : S^\# \vdash^\phi C}{\Gamma; \Delta \vdash^\phi (va)C} \iff \frac{\Gamma; \Delta, a : S^\# \vdash^\phi \mathcal{D}}{\Gamma; \Delta \vdash^\phi (va)\mathcal{D}}$$

Case Parallel Composition

$$\frac{C \equiv \mathcal{D}}{C \parallel \mathcal{E} \equiv \mathcal{D} \parallel \mathcal{E}}$$

There are three subcases, based on whether the parallel composition arises from T-CONNECT₁, T-CONNECT₂, or T-MIX.

Subcase T-MIX

$$\frac{\Gamma_1; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2 \vdash^{\phi_2} \mathcal{E}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{E}} \iff \frac{\Gamma_1; \Delta_1 \vdash^{\phi_1} \mathcal{D} \quad \Gamma_2; \Delta_2 \vdash^{\phi_2} \mathcal{E}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} \mathcal{D} \parallel \mathcal{E}}$$

Subcase T-CONNECT₁

$$\frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{E}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{E}} \iff \frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} \mathcal{D} \quad \Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{E}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} \mathcal{D} \parallel \mathcal{E}}$$

Subcase T-CONNECT₂

$$\frac{\Gamma_1; \Delta_1, a : \bar{S} \vdash^{\phi_1} C \quad \Gamma_2, a : S; \Delta_2 \vdash^{\phi_2} \mathcal{E}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{E}} \iff \frac{\Gamma_1; \Delta_1, a : \bar{S} \vdash^{\phi_1} \mathcal{D} \quad \Gamma_2, a : S; \Delta_2 \vdash^{\phi_2} \mathcal{E}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} \mathcal{D} \parallel \mathcal{E}}$$

Equivalence Axioms

Case $C \parallel \mathcal{D} \equiv \mathcal{D} \parallel C$

There are three subcases, based on which rule is used for parallel composition.

Subcase T-MIX

$$\frac{\Gamma_1; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2 \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}} \iff \frac{\Gamma_2; \Delta_2 \vdash^{\phi_2} \mathcal{D} \quad \Gamma_1; \Delta_1 \vdash^{\phi_1} C}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_2 + \phi_1} \mathcal{D} \parallel C}$$

Subcase T-CONNECT₁

$$\frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}} \iff \frac{\Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D} \quad \Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_2 + \phi_1} \mathcal{D} \parallel C}$$

Subcase T-CONNECT₂

$$\frac{\Gamma_1; \Delta_1, a : \bar{S} \vdash^{\phi_1} C \quad \Gamma_2, a : S; \Delta_2 \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}} \iff \frac{\Gamma_2, a : S; \Delta_2 \vdash^{\phi_2} \mathcal{D} \quad \Gamma_1; \Delta_1, a : \bar{S} \vdash^{\phi_1} C}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_2 + \phi_1} \mathcal{D} \parallel C}$$

Case $C \parallel (va)\mathcal{D} \equiv (va)(C \parallel \mathcal{D})$ if $a \notin \text{fn}(C)$

There are again three subcases based on which parallel composition rule is used. The exact rule does not affect the discussion, so without loss of generality we assume this is T-MIX.

$$\frac{\Gamma_1; \Delta_1 \vdash^{\phi_1} C \quad \frac{\Gamma_2; \Delta_2, a : S^\# \vdash^{\phi_2} \mathcal{D}}{\Gamma_2; \Delta_2 \vdash^{\phi_2} (va)\mathcal{D}}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} C \parallel (va)\mathcal{D}} \iff \frac{\Gamma_1; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2, a : S^\# \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}} \iff \frac{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} (va)(C \parallel \mathcal{D})}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} (va)(C \parallel \mathcal{D})}$$

In the left-to-right direction, that $\Gamma_1, \Gamma_2; \Delta_1, \Delta_2, a : S^\#$ is well-defined follows because $a \notin \text{fn}(C)$.

Case $(va)(vb)C \equiv (vb)(va)C$

$$\frac{\frac{\Gamma; \Delta, a : S^\#, b : T^\# \vdash^\phi C}{\Gamma; \Delta, a : S^\# \vdash^\phi (vb)C}}{\Gamma; \Delta \vdash^\phi (va)(vb)C} \iff \frac{\Gamma; \Delta, b : T^\#, a : S^\# \vdash^\phi C}{\Gamma; \Delta, b : T^\# \vdash^\phi (va)C}}{\Gamma; \Delta \vdash^\phi (vb)(va)C}$$

Case $a(\vec{V}) \iff b(\vec{W}) \equiv b(\vec{W}) \iff a(\vec{V})$

$$\frac{S/\vec{A} = \overline{T/\vec{B}} \quad \Gamma_1 \vdash \vec{V} : \vec{A} \quad \Gamma_2 \vdash \vec{W} : \vec{B}}{\Gamma_1, \Gamma_2; a : S, b : T \vdash^\circ a(\vec{V}) \iff b(\vec{W})} \iff \frac{T/\vec{B} = \overline{S/\vec{A}} \quad \Gamma_2 \vdash \vec{W} : \vec{B} \quad \Gamma_1 \vdash \vec{V} : \vec{A}}{\Gamma_1, \Gamma_2; a : S, b : T \vdash^\circ b(\vec{W}) \iff a(\vec{V})}$$

The above holds because $S/\overrightarrow{A} = \overline{\overrightarrow{T/\overrightarrow{B}}} \iff T/\overrightarrow{B} = \overline{\overrightarrow{S/\overrightarrow{A}}}$:

$$\begin{aligned}
& S/\overrightarrow{A} = \overline{\overrightarrow{T/\overrightarrow{B}}} \\
& \iff \text{(duality)} \\
& \overline{\overrightarrow{S/\overrightarrow{A}}} = \overline{\overrightarrow{T/\overrightarrow{B}}} \\
& \iff \text{(duality is involutive)} \\
& S/\overrightarrow{A} = T/\overrightarrow{B} \\
& \iff \text{(equality is symmetric)} \\
& T/\overrightarrow{B} = S/\overrightarrow{A}
\end{aligned}$$

Case $\circ() \parallel C \equiv C$

$$\frac{\overline{\cdot \vdash () : \mathbf{1}}}{\cdot; \cdot \vdash^\circ \circ() \quad \Gamma; \Delta \vdash^\phi C} \iff \Gamma; \Delta \vdash^\phi \circ() \parallel C$$

Case $(\nu a)(\nu b)(\not\downarrow a \parallel \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)) \parallel C \equiv C$

$$\frac{\frac{\frac{\frac{a : S; \cdot \vdash^\circ \not\downarrow a}{\cdot; \cdot \vdash^\circ \not\downarrow a} \quad \frac{\frac{\overline{S/\epsilon} = \overline{\overline{T/\epsilon}} \quad \cdot \vdash \epsilon : \epsilon \quad \cdot \vdash \epsilon : \epsilon}{\cdot; a : \overline{S}, b : \overline{T} \vdash^\circ a(\epsilon) \rightsquigarrow b(\epsilon)}}{\cdot; a : \overline{S}, b : T^\sharp \vdash^\circ \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}}{\cdot; a : S^\sharp, b : T^\sharp \vdash^\circ \not\downarrow a \parallel \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}}{\cdot; a : S^\sharp \vdash^\circ (\nu b)(\not\downarrow a \parallel \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\epsilon))}}{\cdot; \cdot \vdash^\circ (\nu a)(\nu b)(\not\downarrow a \parallel \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\epsilon))} \quad \Gamma; \Delta \vdash^\phi C} \iff \Gamma; \Delta \vdash^\phi C$$

While it is true that re-associating parallel composition may cause a configuration to be ill-typed, Lemma C.2 shows that it is always possible to re-associate parallel composition either directly, or by first commuting two subconfigurations.

LEMMA C.2 (ASSOCIATIVITY).

- If $\Gamma; \Delta \vdash^\phi C \parallel (\mathcal{D} \parallel \mathcal{E})$, then either $\Gamma; \Delta \vdash^\phi (C \parallel \mathcal{D}) \parallel \mathcal{E}$ or $\Gamma; \Delta \vdash^\phi (C \parallel \mathcal{E}) \parallel \mathcal{D}$.
- If $\Gamma; \Delta \vdash^\phi (C \parallel \mathcal{D}) \parallel \mathcal{E}$, then either $\Gamma; \Delta \vdash^\phi C \parallel (\mathcal{D} \parallel \mathcal{E})$ or $\Gamma; \Delta \vdash^\phi \mathcal{D} \parallel (C \parallel \mathcal{E})$.

PROOF. The cases where either parallel composition arises by T-MIX are unproblematic and can be re-associated without jeopardising typeability. Therefore, we concentrate on the cases where both compositions arise via T-CONNECT_i.

Case $C \parallel (\mathcal{D} \parallel \mathcal{E})$

1716 By the assumption that $\Gamma; \Delta \vdash^\phi C \parallel (\mathcal{D} \parallel \mathcal{E})$ we have that $\Gamma = \Gamma_1, \Gamma_2, \Gamma_3$, and $\Delta = \Delta_1, \Delta_2, \Delta_3$, $a :$
 1717 $S^\sharp, b : T^\sharp$, and $\phi = \phi_1 + \phi_2 + \phi_3$. There are 8 cases, based on whether $a, b \in \text{fn}(C)$ or $a, b \in \text{fn}(\mathcal{D})$
 1718 (it cannot be the case that $a, b \in \text{fn}(\mathcal{E})$, as \mathcal{E} only occurs under a single parallel composition), and
 1719 the exact dualisation (i.e., whether composition happens via T-CONNECT₁ or T-CONNECT₂).

1720 Of these, we are only interested in the cases where the sharing of the names differs, as opposed
 1721 to the dualisation. Thus, we consider the following two cases, where both compositions occur using
 1722 T-CONNECT₁:

1723 (1) $\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C$, and $\Gamma_2, b : T; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D}$, and $\Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E}$

1724 (2) $\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C$, and $\Gamma_2, b : T; \Delta_2 \vdash^{\phi_2} \mathcal{D}$, and $\Gamma_3; \Delta_3, a : \bar{S}, b : \bar{T} \vdash^{\phi_3} \mathcal{E}$

1725 **Subcase** $a \in \text{fn}(C), a, b \in \mathcal{D}, b \in \mathcal{E}$

$$\frac{\frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \frac{\Gamma_2, b : T; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D} \quad \Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E}}{\Gamma_2, \Gamma_3; \Delta_2, \Delta_3, a : \bar{S}, b : T^\sharp \vdash^{\phi_2 + \phi_3} \mathcal{D} \parallel \mathcal{E}}}{\Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp \vdash^{\phi_1 + \phi_2 + \phi_3} C \parallel (\mathcal{D} \parallel \mathcal{E})}}{\Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp \vdash^{\phi_1 + \phi_2 + \phi_3} C \parallel (\mathcal{D} \parallel \mathcal{E})}$$

1732 As \mathcal{D} contains both a and b , associativity does not alter the sharing of names and may be applied
 1733 safely.

$$\frac{\frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2, b : T; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D}}{\Gamma_1, \Gamma_2, b : T; \Delta_1, \Delta_2, a : S^\sharp \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D}} \quad \Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E}}{\Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp \vdash^{\phi_1 + \phi_2 + \phi_3} (C \parallel \mathcal{D}) \parallel \mathcal{E}}$$

1739 **Subcase** $a \in \text{fn}(C); b \in \mathcal{D}; a, b \in \mathcal{E}$

$$\frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \frac{\Gamma_2, b : T; \Delta_2 \vdash^{\phi_2} \mathcal{D} \quad \Gamma_3; \Delta_3, a : \bar{S}, b : \bar{T} \vdash^{\phi_3} \mathcal{E}}{\Gamma_2, \Gamma_3; \Delta_2, \Delta_3, a : \bar{S}, b : T^\sharp \vdash^{\phi_2 + \phi_3} \mathcal{D} \parallel \mathcal{E}}}{\Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp \vdash^{\phi_1 + \phi_2 + \phi_3} C \parallel (\mathcal{D} \parallel \mathcal{E})}$$

1742 Here, we may not apply associativity directly. But, we may first commute \mathcal{D} and \mathcal{E} :

$$\frac{\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \frac{\Gamma_3; \Delta_3, a : \bar{S}, b : \bar{T} \vdash^{\phi_3} \mathcal{E} \quad \Gamma_2, b : T; \Delta_2 \vdash^{\phi_2} \mathcal{D}}{\Gamma_2, \Gamma_3; \Delta_2, \Delta_3, a : \bar{S}, b : T^\sharp \vdash^{\phi_2 + \phi_3} \mathcal{E} \parallel \mathcal{D}}}{\Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp \vdash^{\phi_1 + \phi_2 + \phi_3} C \parallel (\mathcal{E} \parallel \mathcal{D})}$$

1752 and from here we may safely re-associate to the left:

$$\frac{\frac{\Gamma_2, a : S; \Delta_2 \vdash^{\phi_1} C \quad \Gamma_3; \Delta_3, a : \bar{S}, b : \bar{T} \vdash^{\phi_3} \mathcal{E}}{\Gamma_2, \Gamma_3; \Delta_2, \Delta_3, a : S^\sharp, b : \bar{T} \vdash^{\phi_1 + \phi_2} \mathcal{D} \parallel \mathcal{E}} \quad \Gamma_3, b : T; \Delta_3 \vdash^{\phi_3} \mathcal{D}}{\Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp \vdash^{\phi_1 + \phi_2 + \phi_3} (C \parallel \mathcal{E}) \parallel \mathcal{D}}$$

1759 **Case** $(C \parallel \mathcal{D}) \parallel \mathcal{E}$

1760 (1) $\Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C$, and $\Gamma_2, b : T; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D}$, and $\Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E}$

1761 (2) $\Gamma_1, a : S, b : T; \Delta_1 \vdash^{\phi_1} C$, and $\Gamma_2; \Delta_2, b : \bar{T} \vdash^{\phi_2} \mathcal{D}$, and $\Gamma_3; \Delta_3, a : \bar{S} \vdash^{\phi_3} \mathcal{E}$

1762 **Subcase** $a \in C; a, b \in \mathcal{D}; b \in \mathcal{E}$

1763

1764

1765 Assumption:

$$\begin{array}{c}
 1766 \quad \Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2, b : T; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D} \\
 1767 \quad \hline \\
 1768 \quad \Gamma_1, \Gamma_2, b : T; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D} \quad \Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E} \\
 1769 \quad \hline \\
 1770 \quad \Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\#, b : T^\# \vdash^{\phi_1 + \phi_2 + \phi_3} (C \parallel \mathcal{D}) \parallel \mathcal{E}
 \end{array}$$

1771 Applying associativity here does not make the configuration ill-typed, as \mathcal{D} contains both names:

$$\begin{array}{c}
 1772 \quad \Gamma_2, b : T; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D} \quad \Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E} \\
 1773 \quad \hline \\
 1774 \quad \Gamma_1, a : S; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2, \Gamma_3; \Delta_2, \Delta_3, a : \bar{S}, b : T^\# \vdash^{\phi_2 + \phi_3} \mathcal{D} \parallel \mathcal{E} \\
 1775 \quad \hline \\
 1776 \quad \Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\#, b : T^\# \vdash^{\phi_1 + \phi_2 + \phi_3} C \parallel (\mathcal{D} \parallel \mathcal{E})
 \end{array}$$

1776 **Subcase** $a, b \in C; a \in \mathcal{D}; b \in \mathcal{E}$

1777 Assumption:

$$\begin{array}{c}
 1778 \quad \Gamma_1, a : S, b : T; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D} \\
 1779 \quad \hline \\
 1780 \quad \Gamma_2, \Gamma_3, b : T; \Delta_2, \Delta_3, a : S^\# \vdash^{\phi_2 + \phi_3} C \parallel \mathcal{D} \quad \Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E} \\
 1781 \quad \hline \\
 1782 \quad \Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\#, b : T^\# \vdash^{\phi_1 + \phi_2 + \phi_3} (C \parallel \mathcal{D}) \parallel \mathcal{E}
 \end{array}$$

1783 By commutativity:

$$\begin{array}{c}
 1784 \quad \Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D} \quad \Gamma_1, a : S, b : T; \Delta_1 \vdash^{\phi_1} C \\
 1785 \quad \hline \\
 1786 \quad \Gamma_2, \Gamma_3, b : T; \Delta_2, \Delta_3, a : S^\# \vdash^{\phi_2 + \phi_1} \mathcal{D} \parallel C \quad \Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E} \\
 1787 \quad \hline \\
 1788 \quad \Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\#, b : T^\# \vdash^{\phi_1 + \phi_2 + \phi_3} (\mathcal{D} \parallel C) \parallel \mathcal{E}
 \end{array}$$

1789 By associativity:

$$\begin{array}{c}
 1790 \quad \Gamma_1, a : S, b : T; \Delta_1 \vdash^{\phi_1} C \quad \Gamma_3; \Delta_3, b : \bar{T} \vdash^{\phi_3} \mathcal{E} \\
 1791 \quad \hline \\
 1792 \quad \Gamma_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D} \quad \Gamma_1, \Gamma_3, a : S; \Delta_1, \Delta_3, b : T^\# \vdash^{\phi_1 + \phi_3} C \parallel \mathcal{E} \\
 1793 \quad \hline \\
 1794 \quad \Gamma_1, \Gamma_2, \Gamma_3; \Delta_1, \Delta_2, \Delta_3, a : S^\#, b : T^\# \vdash^{\phi_1 + \phi_2 + \phi_3} \mathcal{D} \parallel (C \parallel \mathcal{E})
 \end{array}$$

1794 as required. □

1797 **C.1.2 Configuration Reduction.** We may now show that configuration reduction preserves
 1798 typeability of configurations. We begin by stating some auxiliary results about substitution and
 1799 evaluation contexts.

1800 Typing of terms is preserved by substitution.

1801 LEMMA C.3 (SUBSTITUTION). *If:*

- 1803 (1) $\Gamma_1 \vdash M : B$
- 1804 (2) $\Gamma_2, x : B \vdash N : A$
- 1805 (3) Γ_1, Γ_2 is well-defined

1806 then $\Gamma_1, \Gamma_2 \vdash N\{M/x\} : A$.

1807 PROOF. By induction on the derivation of $\Gamma_2, x : B \vdash N : A$. □

1809 Lemma C.4 shows that a subterm of a well-typed evaluation context E (and therefore also a pure
 1810 evaluation context P) is typeable with a subset of the type environment. Lemma C.5 states that the
 1811 subterm of a well-typed evaluation context can be replaced. Both follow the formulation of Gay
 1812 and Vasconcelos [2010].

1814 LEMMA C.4 (TYPEABILITY OF SUBTERMS). *If \mathbf{D} is a derivation of $\Gamma \vdash E[M] : A$, then there exist Γ_1, Γ_2*
 1815 *and B such that $\Gamma = \Gamma_1, \Gamma_2$, that \mathbf{D} has a subderivation \mathbf{D}' that concludes $\Gamma_2 \vdash M : B$, and the position*
 1816 *of \mathbf{D}' in \mathbf{D} corresponds to the position of the hole in E .*

1817 PROOF. By induction on the structure of E . □

1819 LEMMA C.5 (REPLACEMENT (EVALUATION CONTEXTS)). *If:*

- 1820 • \mathbf{D} is a derivation of $\Gamma_1, \Gamma_2 \vdash E[M] : A$
- 1821 • \mathbf{D}' is a subderivation of \mathbf{D} concluding $\Gamma_2 \vdash M : B$
- 1822 • The position of \mathbf{D}' in \mathbf{D} corresponds to that of the hole in E
- 1823 • $\Gamma_3 \vdash N : B$
- 1824 • Γ_1, Γ_3 is well-defined

1825 then $\Gamma_1, \Gamma_3 \vdash E[N] : A$.

1827 PROOF. By induction on the structure of E . □

1829 To prove preservation on configurations, we must first establish some auxiliary results on
 1830 configuration contexts. Lemma C.6 states how we may type subconfigurations.

1831 LEMMA C.6 (TYPEABILITY OF SUBCONFIGURATIONS). *If \mathbf{D} is a derivation of $\Gamma; \Delta \vdash^\phi \mathcal{G}[C]$, then*
 1832 *there exist Γ', Δ', ϕ' such that \mathbf{D} has a subderivation \mathbf{D}' that concludes $\Gamma'; \Delta' \vdash^{\phi'} C$, and the position*
 1833 *of \mathbf{D}' in \mathbf{D} corresponds to the position of the hole in \mathcal{G} .*

1835 PROOF. By induction on the structure of \mathcal{G} . □

1837 Lemma C.7 states that we may replace a subconfiguration of a configuration context. The lemma
 1838 is slightly complicated by the fact that $(va)\mathcal{G}$ binds a variable a , but replacement is safe if the typing
 1839 environments are related by the environment reduction relation.

1841 LEMMA C.7 (REPLACEMENT (CONFIGURATIONS)). *If:*

- 1842 • \mathbf{D} is a derivation of $\Gamma; \Delta \vdash^\phi \mathcal{G}[C]$
- 1843 • \mathbf{D}' is a subderivation of \mathbf{D} concluding that $\Gamma'; \Delta' \vdash^{\phi'} C$ for some Γ', Δ', ϕ'
- 1844 • $\Gamma''; \Delta'' \vdash^{\phi'} C'$ for some Γ'', Δ'' such that $\Gamma'; \Delta' \longrightarrow^? \Gamma''; \Delta''$
- 1845 • The position of \mathbf{D} in \mathbf{D}' corresponds to that of the hole in \mathcal{G}

1846 then there exist some Γ''', Δ''' such that $\Gamma'''; \Delta''' \vdash^\phi \mathcal{G}[C']$ and $\Gamma; \Delta \longrightarrow^? \Gamma'''; \Delta'''$.

1848 PROOF. By induction on the structure of \mathcal{G} . □

1850 Theorem 3.2 (Preservation (Configurations))

1851 Assume Γ only contains entries of the form $a_i : S_i$.

1852 If $\Gamma; \Delta \vdash^\phi C$ and $C \longrightarrow \mathcal{D}$, then there exist Γ', Δ' such that $\Gamma; \Delta \longrightarrow^? \Gamma'; \Delta'$ and $\Gamma'; \Delta' \vdash^\phi \mathcal{D}$.

1854 PROOF. By induction on the derivation of $C \longrightarrow \mathcal{D}$. Where there is a choice of value for ϕ , we
 1855 consider the case where $\phi = \bullet$; the cases where $\phi = \circ$ are similar.

1856 Case E-Fork

1858 Assumption:

$$1860 \frac{\Gamma_1, \Gamma_2 \vdash \bullet E[\mathbf{fork} \lambda x.M] : A}{1861 \Gamma_1, \Gamma_2; \cdot \vdash \bullet \bullet E[\mathbf{fork} \lambda x.M]}$$

1862

By Lemma C.4:

$$\frac{\Gamma_2, x : S \vdash M : \mathbf{1}}{\Gamma_2 \vdash \lambda x.M : S \multimap \mathbf{1}} \\ \Gamma_2 \vdash \mathbf{fork} \lambda x.M : \bar{S}$$

By Lemma C.3, $\Gamma_2, b : S \vdash M\{b/x\} : \mathbf{1}$, and by Lemma C.5, $\Gamma_1, a : \bar{S} \vdash E[a] : A$. As duality is involutive, $\bar{\bar{S}} = S$.

Reconstructing:

$$\frac{\frac{\Gamma_1, a : \bar{S} \vdash E[a] : A}{\Gamma_1, a : \bar{S}; \cdot \vdash \bullet \bullet E[a]} \quad \frac{\Gamma_2, b : S \vdash \circ M\{b/x\} \quad \frac{S/\epsilon = \bar{\bar{S}}/\epsilon \quad \cdot \vdash \epsilon : \epsilon \quad \cdot \vdash \epsilon : \epsilon}{\cdot; a : S, b : \bar{S} \vdash \circ a(\epsilon) \rightsquigarrow b(\epsilon)}}{\Gamma_2; a : S, b : S^\# \vdash \circ M\{b/x\} \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}}{\Gamma_1, \Gamma_2; a : \bar{S}^\#, b : S^\# \vdash \bullet \bullet E[a] \parallel \circ M\{b/x\} \parallel a(\epsilon) \rightsquigarrow b(\epsilon)} \\ \Gamma_1, \Gamma_2; a : \bar{S}^\# \vdash \bullet (vb)(\bullet E[a] \parallel \circ M\{b/x\} \parallel a(\epsilon) \rightsquigarrow b(\epsilon)) \\ \Gamma_1, \Gamma_2; \cdot \vdash \bullet (va)(vb)(\bullet E[a] \parallel \circ M\{b/x\} \parallel a(\epsilon) \rightsquigarrow b(\epsilon))$$

Case E-Send

Assumption:

$$\frac{\Gamma_1, \Gamma_2 \vdash E[\mathbf{send} U a] : C \quad \frac{\bar{S}/\bar{A} = \overline{T/\bar{B}} \quad \Gamma_3 \vdash \vec{V} : \vec{A} \quad \Gamma_4 \vdash \vec{W} : \vec{B}}{\Gamma_3, \Gamma_4; a : \bar{S}, b : T \vdash \circ a(\vec{V}) \rightsquigarrow b(\vec{W})}}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : S^\#, b : T \vdash \bullet \bullet E[\mathbf{send} U a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})}$$

By Lemma C.4:

$$\frac{\Gamma_2 \vdash U : A \quad a : !A.S' \vdash a : !A.S'}{\Gamma_2, a : !A.S' \vdash \mathbf{send} U a : S'}$$

Thus, $S = !A.S'$, and $\bar{S} = ?A.\bar{S}'$. We may therefore refine our original derivation:

$$\frac{\Gamma_1, \Gamma_2, a : !A.S' \vdash E[\mathbf{send} U a] : C \quad \frac{?A.\bar{S}'/\bar{A} = \overline{T/\bar{B}} \quad \Gamma_3 \vdash \vec{V} : \vec{A} \quad \Gamma_4 \vdash \vec{W} : \vec{B}}{\Gamma_3, \Gamma_4; a : ?A.\bar{S}', b : T \vdash \circ a(\vec{V}) \rightsquigarrow b(\vec{W})}}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : !A.S'^\#, b : T \vdash \bullet \bullet E[\mathbf{send} U a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})}$$

Since $?A.\bar{S}'/\bar{A} = \overline{T/\bar{B}}$ is well-defined, we have that $\vec{A} = \epsilon$. By the definition of slicing, we have that $\bar{T} = !B_1 \cdot \dots \cdot !B_n \cdot !A.S'$, where $\vec{B} = B_1, \dots, B_n$. It follows that $\bar{S}'/\bar{A} = \overline{T/\bar{B}} \cdot A$.

By Lemma C.5, we have $\Gamma_1, \Gamma_2, a : S' \vdash E[a] : C$.

Reconstructing:

$$\frac{\Gamma_1, a : S' \vdash E[a] : C \quad \frac{\bar{S}'/\bar{A} = \overline{T/\bar{B}} \cdot A \quad \Gamma_3 \vdash \vec{V} : \vec{A} \quad \Gamma_2, \Gamma_4 \vdash \vec{W} \cdot U : \vec{B} \cdot A}{\Gamma_2, \Gamma_3, \Gamma_4; a : \bar{S}', b : T \vdash \circ a(\vec{V}) \rightsquigarrow b(\vec{W} \cdot U)}}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : S'^\#, b : T \vdash \bullet \bullet E[a] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W} \cdot U)}$$

Finally, we must show environment reduction:

$$\frac{!A.S' \longrightarrow S'}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : (!A.S')^\#, b : T \longrightarrow \Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : S'^\#, b : T}$$

as required.

Case E-Receive

Assumption:

$$\frac{\frac{\Gamma_1, a : S \vdash E[\mathbf{receive} \ a] : C}{\Gamma_1, a : S; \cdot \vdash^\bullet E[\mathbf{receive} \ a]} \quad \frac{\overline{S/\bar{A}} = \overline{T/\bar{B}} \quad \Gamma_2, \Gamma_3 \vdash U \cdot \vec{V} : \vec{A} \quad \Gamma_4 \vdash \vec{W} : \vec{B}}{\Gamma_2, \Gamma_3, \Gamma_4; a : \bar{S}, b : T \vdash^\circ a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})}}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : S^\#, b : T \vdash^\bullet \bullet E[\mathbf{receive} \ a] \parallel a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})}$$

By [Theorem C.4](#):

$$\frac{a : ?A.S' \vdash a : ?A.S'}{a : ?A.S' \vdash \mathbf{receive} \ a : (A \times S')}$$

Thus, we have that $S = ?A.S'$ and $\bar{S} = !A.\bar{S}'$, and we may therefore refine the original typing derivation:

$$\frac{\frac{\Gamma_1, a : ?A.S' \vdash E[\mathbf{receive} \ a] : C}{\Gamma_1, a : ?A.S'; \cdot \vdash^\bullet E[\mathbf{receive} \ a]} \quad \frac{\frac{!A.\bar{S}'/A \cdot \vec{A}' = \overline{T/\bar{B}} \quad \frac{\Gamma_1 \vdash U : A \quad \Gamma_3 \vdash \vec{V} : \vec{A}'}{\Gamma_2, \Gamma_3 \vdash U \cdot \vec{V} : A \cdot \vec{A}'}}{\Gamma_2, \Gamma_3, \Gamma_4; a : !A.\bar{S}', b : T \vdash^\circ a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})}}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : (?A.S')^\#, b : T \vdash^\bullet \bullet E[\mathbf{receive} \ a] \parallel a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})}}$$

By [Lemma C.5](#), we have $\Gamma_1, \Gamma_2, a : S' \vdash E[(U, a)] : C$ (that Γ_1, Γ_2 is defined follows from the fact that Γ_1 and Γ_2 are sub-environments of the original typing environment and are therefore necessarily disjoint).

By the definition of slicing, $!A.\bar{S}'/A \cdot \vec{A}' \iff \bar{S}'/\vec{A}'$.

Thus, recomposing:

$$\frac{\frac{\Gamma_1, \Gamma_2, a : S' \vdash E[(U, a)] : C}{\Gamma_1, \Gamma_2, a : S'; \cdot \vdash^\bullet E[(U, a)]} \quad \frac{\bar{S}'/\vec{A}' = \overline{T/\bar{B}} \quad \Gamma_3 \vdash \vec{V} : \vec{A}' \quad \Gamma_4 \vdash \vec{W} : \vec{B}}{\Gamma_3, \Gamma_4; a : \bar{S}', b : T \vdash^\circ a(\vec{V}) \rightsquigarrow b(\vec{W})}}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : S'^\#, b : T \vdash^\bullet \bullet E[(U, a)] \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})}$$

Finally, we must show environment reduction:

$$\frac{?A.S' \longrightarrow S'}{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : (?A.S')^\#; b : T \longrightarrow \Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4; a : S'^\#, b : T}$$

as required.

Case E-Close

1961 Assumption:

$$\begin{array}{c}
 1962 \\
 1963 \\
 1964 \\
 1965 \\
 1966 \\
 1967 \\
 1968 \\
 1969 \\
 1970 \\
 1971
 \end{array}
 \frac{
 \frac{
 \frac{
 \Gamma_1, a : S \vdash E[\mathbf{close} a] : C
 }{
 \Gamma_1, a : S; \cdot \vdash \bullet E[\mathbf{close} a]
 }
 \quad
 \frac{
 \Gamma_2, b : T \vdash E'[\mathbf{close} b] : \mathbf{1} \quad \overline{S/\epsilon} = \overline{T/\epsilon} \quad \overline{\cdot \vdash \epsilon : \epsilon} \quad \overline{\cdot \vdash \epsilon : \epsilon}
 }{
 \Gamma_2, b : T; \cdot \vdash^\circ \circ E'[\mathbf{close} b] \quad ; a : \overline{S}, b : \overline{T} \vdash^\circ a(\epsilon) \rightsquigarrow b(\epsilon)
 }
 }{
 \Gamma_2; a : \overline{S}, b : T^\# \vdash^\circ E'[\mathbf{close} b] \parallel a(\epsilon) \rightsquigarrow b(\epsilon)
 }
 }{
 \Gamma_1, \Gamma_2; a : S^\#, b : T^\# \vdash \bullet \bullet E[\mathbf{close} a] \parallel \circ E'[\mathbf{close} b] \parallel a(\epsilon) \rightsquigarrow b(\epsilon)
 }
 }{
 \Gamma_1, \Gamma_2; a : S^\# \vdash \bullet (vb)(\bullet E[\mathbf{close} a]) \parallel \circ E'[\mathbf{close} b] \parallel a(\epsilon) \rightsquigarrow b(\epsilon)
 }
 }{
 \Gamma_1, \Gamma_2; \cdot \vdash \bullet (va)(vb)(\bullet E[\mathbf{close} a]) \parallel \circ E'[\mathbf{close} b] \parallel a(\epsilon) \rightsquigarrow b(\epsilon)
 }$$

1972 By Lemma C.4:

$$\begin{array}{c}
 1973 \\
 1974 \\
 1975 \\
 1976
 \end{array}
 \frac{
 \frac{a : \text{End} \vdash a : \text{End}}{a : \text{End} \vdash \mathbf{close} a : \mathbf{1}} \quad \frac{b : \text{End} \vdash b : \text{End}}{b : \text{End} \vdash \mathbf{close} b : \mathbf{1}}
 }{
 }$$

1977 By Lemma C.5, we have that $\Gamma_1 \vdash E[()] : C$ and that $\Gamma_2 \vdash E'[()] : \mathbf{1}$. Thus by T-Mix, we may show:

$$\begin{array}{c}
 1978 \\
 1979 \\
 1980 \\
 1981
 \end{array}
 \frac{
 \frac{
 \frac{
 \Gamma_1 \vdash E[()] : C
 }{
 \Gamma_1; \cdot \vdash \bullet \bullet E[()]
 }
 \quad
 \frac{
 \Gamma_2 \vdash E'[()] : \mathbf{1}
 }{
 \Gamma_2; \cdot \vdash^\circ \circ E'[()]
 }
 }{
 \Gamma_1, \Gamma_2; \cdot \vdash \bullet \bullet E[()] \parallel \circ E'[()]
 }
 }{
 }$$

1982 as required.

1983 **Case E-Cancel**

$$\mathcal{F}[\mathbf{cancel} a] \longrightarrow \mathcal{F}[()] \parallel \not\downarrow a$$

1987 Assumption:

$$\begin{array}{c}
 1988 \\
 1989 \\
 1990 \\
 1991
 \end{array}
 \frac{
 \Gamma \vdash E[\mathbf{cancel} a] : C
 }{
 \Gamma; \cdot \vdash \bullet \bullet E[\mathbf{cancel} a]
 }$$

1992 By Lemma C.4, $\Gamma = \Gamma_1, \Gamma_2$, where

$$\begin{array}{c}
 1993 \\
 1994 \\
 1995
 \end{array}
 \frac{
 \Gamma_2 \vdash a : S
 }{
 \Gamma_2 \vdash \mathbf{cancel} a : \mathbf{1}
 }$$

1996 Thus $\Gamma_2 = a : S$. By Lemma C.5, $\Gamma_1 \vdash E[()] : C$. By T-ZAP, we have that $a : S \vdash^\circ \not\downarrow a$. Thus,

1997 recomposing:

$$\begin{array}{c}
 1998 \\
 1999 \\
 2000 \\
 2001 \\
 2002
 \end{array}
 \frac{
 \frac{
 \Gamma \vdash E[()] : C
 }{
 \Gamma_1; \cdot \vdash \bullet \bullet E[()]
 }
 \quad
 \frac{
 }{
 a : S; \cdot \vdash^\circ \not\downarrow a
 }
 }{
 \Gamma_1, a : S; \cdot \vdash \bullet \bullet E[()] \parallel \not\downarrow a
 }$$

2003 as required.

2004 **Case E-Zap**

$$\not\downarrow a \parallel a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W}) \longrightarrow \not\downarrow a \parallel \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})$$

2008 where $\text{fn}(U) = \{c_i\}_i$.

Assumption:

$$\frac{\frac{a : S; \cdot \vdash^\circ \not\leq a}{\Gamma_1, \Gamma_2, \Gamma_3; a : \bar{S}, b : T \vdash^\circ a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})} \quad \frac{\bar{S}/\vec{A} = \overline{T/\vec{B}} \quad \Gamma_1, \Gamma_2 \vdash U \cdot \vec{V} : \vec{A} \quad \Gamma_3 \vdash \vec{W} : \vec{B}}{\Gamma_1, \Gamma_2, \Gamma_3; a : S^\sharp, b : T \vdash^\circ \not\leq a \parallel a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})}}$$

By the definition of slicing, we have that there exist some A and S' such that $\bar{S} = !A.\bar{S}'$. Thus, we may refine our judgement:

$$\frac{\frac{a : ?A.S'; \cdot \vdash^\circ \not\leq a}{\Gamma_1, \Gamma_2, \Gamma_3; a : \bar{S}, b : T \vdash^\circ a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})} \quad \frac{!A.\bar{S}'/A \cdot \vec{A}' = \overline{T/\vec{B}} \quad \Gamma_1, \Gamma_2 \vdash U \cdot \vec{V} : A \cdot \vec{A}' \quad \Gamma_3 \vdash \vec{W} : \vec{B}}{\Gamma_1, \Gamma_2, \Gamma_3; a : S^\sharp, b : T \vdash^\circ \not\leq a \parallel a(U \cdot \vec{V}) \rightsquigarrow b(\vec{W})}}$$

By the definition of buffer typing, we have that $\Gamma_1 \vdash U : A$. By the definition of the reduction rule, $\text{fn}(U) = \{c_i\}_i$, and by assumption, Γ contains only runtime names. Thus, we may conclude that U is closed and therefore that $\Gamma_1 = c_1 : S_1, \dots, c_n : S_n$ for some session types S_1, \dots, S_n .

By the definition of slicing, we have that $!A.\bar{S}'/A \cdot \vec{A}' \iff \bar{S}'/\vec{A}'$. Correspondingly, by T-BUFFER, we may show

$$\frac{\bar{S}'/\vec{A}' = \overline{T/\vec{B}} \quad \Gamma_2 \vdash \vec{V} : \vec{A}' \quad \Gamma_3 \vdash \vec{W} : \vec{B}}{\Gamma_2, \Gamma_3; a : \bar{S}', b : T \vdash^\circ a(\vec{V}) \rightsquigarrow b(\vec{W})}$$

By repeated applications of T-ZAP and T-MIX, we have that

$$\Gamma_2, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; a : \bar{S}', b : T \vdash^\circ \not\leq c_1 \parallel \dots \parallel \not\leq c_n \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})$$

Recomposing:

$$\frac{\frac{a : S'; \cdot \vdash^\circ \not\leq a}{\Gamma_2, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; a : \bar{S}', b : T \vdash^\circ \not\leq c_1 \parallel \dots \parallel \not\leq c_n \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})} \quad \frac{\frac{c_1 : S_n; \cdot \vdash^\circ \not\leq c_n}{\Gamma_2, \Gamma_3; a : \bar{S}', b : T \vdash^\circ a(\vec{V}) \rightsquigarrow b(\vec{W})} \quad \frac{\bar{S}'/\vec{A}' = \overline{T/\vec{B}} \quad \Gamma_2 \vdash \vec{V} : \vec{A}' \quad \Gamma_3 \vdash \vec{W} : \vec{B}}{\Gamma_2, \Gamma_3; a : \bar{S}', b : T \vdash^\circ a(\vec{V}) \rightsquigarrow b(\vec{W})}}{\Gamma_2, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; a : S'^\sharp, b : T \vdash^\circ \not\leq c_1 \parallel \dots \parallel \not\leq c_n \parallel a(\vec{V}) \rightsquigarrow b(\vec{W})} \quad \frac{c_1 : S_1; \cdot \vdash^\circ \not\leq c_1}{\vdots}}$$

Finally, we must show environment reduction:

$$\frac{?A.S' \longrightarrow S'}{\Gamma_2, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; a : (?A.S'^\sharp), b : T \longrightarrow \Gamma_2, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; a : S'^\sharp, b : T}$$

as required.

Case E-CloseZap

$$\mathcal{F}[\mathbf{close} \ a] \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon) \longrightarrow \mathcal{F}[\mathbf{raise}] \parallel \not\leq a \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)$$

2059 Assumption:

$$\begin{array}{c}
 2060 \\
 2061 \\
 2062 \\
 2063 \\
 2064 \\
 2065 \\
 2066 \\
 2067 \\
 2068 \\
 2069 \\
 2070 \\
 2071 \\
 2072 \\
 2073 \\
 2074 \\
 2075 \\
 2076 \\
 2077 \\
 2078 \\
 2079
 \end{array}$$

$$\frac{\Gamma, a : S \vdash E[\mathbf{close} a] : C \quad \frac{b : T; \cdot \vdash^\circ \not\leq b \quad \frac{\overline{\overline{S} = \overline{\overline{T}}} \quad \overline{\cdot \vdash \epsilon : \epsilon} \quad \overline{\cdot \vdash \epsilon : \epsilon}}{\cdot; a : \overline{S}, b : \overline{T} \vdash^\circ a(\epsilon) \rightsquigarrow b(\epsilon)}}{\cdot; a : \overline{S}, b : T^\# \vdash^\circ \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}}{\Gamma; a : S^\#, b : T^\# \vdash^\bullet \bullet E[\mathbf{close} a] \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}$$

By Lemma C.4:

$$\frac{a : \mathbf{End} \vdash a : \mathbf{End}}{a : S \vdash \mathbf{close} a : 1}$$

We may therefore refine our original derivation:

$$\begin{array}{c}
 2080 \\
 2081 \\
 2082 \\
 2083 \\
 2084 \\
 2085 \\
 2086 \\
 2087 \\
 2088 \\
 2089 \\
 2090 \\
 2091 \\
 2092 \\
 2093 \\
 2094 \\
 2095 \\
 2096 \\
 2097 \\
 2098 \\
 2099 \\
 2100 \\
 2101 \\
 2102 \\
 2103 \\
 2104 \\
 2105 \\
 2106 \\
 2107
 \end{array}$$

$$\frac{\Gamma, a : \mathbf{End} \vdash E[\mathbf{close} a] : C \quad \frac{b : \mathbf{End}; \cdot \vdash^\circ \not\leq b \quad \frac{\mathbf{End} = \mathbf{End} \quad \overline{\cdot \vdash \epsilon : \epsilon} \quad \overline{\cdot \vdash \epsilon : \epsilon}}{\cdot; a : \mathbf{End}, b : \mathbf{End} \vdash^\circ a(\epsilon) \rightsquigarrow b(\epsilon)}}{\cdot; a : \mathbf{End}, b : \mathbf{End}^\# \vdash^\circ \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}}{\Gamma; a : \mathbf{End}^\#, b : \mathbf{End}^\# \vdash^\bullet \bullet E[\mathbf{close} a] \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}$$

By Lemma C.5, $\Gamma \vdash E[\mathbf{raise}] : C$.

Thus, recomposing:

$$\begin{array}{c}
 2080 \\
 2081 \\
 2082 \\
 2083 \\
 2084 \\
 2085 \\
 2086 \\
 2087 \\
 2088 \\
 2089 \\
 2090 \\
 2091 \\
 2092 \\
 2093 \\
 2094 \\
 2095 \\
 2096 \\
 2097 \\
 2098 \\
 2099 \\
 2100 \\
 2101 \\
 2102 \\
 2103 \\
 2104 \\
 2105 \\
 2106 \\
 2107
 \end{array}$$

$$\frac{\Gamma \vdash E[\mathbf{raise}] : C \quad \frac{a : \mathbf{End}; \cdot \vdash^\circ \not\leq a \quad \frac{b : \mathbf{End}; \cdot \vdash^\circ \not\leq b \quad \frac{\mathbf{End} = \mathbf{End} \quad \overline{\cdot \vdash \epsilon : \epsilon} \quad \overline{\cdot \vdash \epsilon : \epsilon}}{\cdot; a : \mathbf{End}, b : \mathbf{End} \vdash^\circ a(\epsilon) \rightsquigarrow b(\epsilon)}}{\cdot; a : \mathbf{End}^\#, b : \mathbf{End}^\# \vdash^\circ \not\leq a \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}}{\Gamma; a : \mathbf{End}^\#, b : \mathbf{End}^\# \vdash^\bullet \bullet E[\mathbf{close} a] \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}$$

as required.

2091 **Case E-ReceiveZap**

$$\bullet E[\mathbf{receive} a] \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\vec{W}) \longrightarrow \bullet E[\mathbf{raise}] \parallel \not\leq a \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})$$

Assumption:

$$\begin{array}{c}
 2096 \\
 2097 \\
 2098 \\
 2099 \\
 2100 \\
 2101 \\
 2102 \\
 2103 \\
 2104 \\
 2105 \\
 2106 \\
 2107
 \end{array}$$

$$\frac{\Gamma_1, a : S \vdash E[\mathbf{receive} a] : C \quad \frac{b : T; \cdot \vdash^\circ \not\leq b \quad \frac{\overline{\overline{S/\epsilon} = \overline{\overline{T/\vec{B}}}} \quad \overline{\cdot \vdash \epsilon : \epsilon} \quad \Gamma_2 \vdash \vec{W} : \vec{B}}{\Gamma_2; a : \overline{S}, b : \overline{T} \vdash^\circ a(\epsilon) \rightsquigarrow b(\vec{W})}}{\Gamma_2; a : \overline{S}, b : T^\# \vdash^\circ \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})}}{\Gamma_1, \Gamma_2; a : S^\#, b : T^\# \vdash^\bullet \bullet E[\mathbf{receive} a] \parallel \not\leq b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})}$$

By Lemma C.4:

$$\frac{a : ?A.S' \vdash a : ?A.S'}{a : ?A.S' \vdash \mathbf{receive} a : (A \times S')}$$

By Lemma C.5, $\Gamma_1 \vdash E[\mathbf{raise}] : S'$. Thus, recomposing:

$$\frac{\frac{\Gamma_1 \vdash E[\mathbf{raise}] : C}{\Gamma_1; \cdot \vdash^\bullet \bullet E[\mathbf{raise}]}}{\Gamma_1, \Gamma_2; a : S^\sharp, b : T^\sharp \vdash^\bullet \bullet E[\mathbf{raise}] \parallel \not\downarrow a \parallel \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})}}{\frac{\frac{\frac{\frac{\Gamma_2; \cdot \vdash^\circ \not\downarrow b}{\Gamma_2; a : \bar{S}, b : \bar{T} \vdash^\circ a(\epsilon) \rightsquigarrow b(\vec{W})}}{\bar{S}/\epsilon = \bar{T}/\bar{B}} \quad \cdot \vdash \epsilon : \epsilon \quad \Gamma_2 \vdash \vec{W} : \vec{B}}{\Gamma_2; a : \bar{S}, b : T^\sharp \vdash^\circ \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})}}{\Gamma_2; a : S^\sharp, b : T^\sharp \vdash^\circ \not\downarrow a \parallel \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})}}{\Gamma_1, \Gamma_2; a : S^\sharp, b : T^\sharp \vdash^\bullet \bullet E[\mathbf{raise}] \parallel \not\downarrow a \parallel \not\downarrow b \parallel a(\epsilon) \rightsquigarrow b(\vec{W})}}$$

as required.

Case E-Raise

$$\bullet E[\mathbf{try } P[\mathbf{raise}] \mathbf{ as } x \mathbf{ in } M \mathbf{ otherwise } N] \longrightarrow E[N] \parallel \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n$$

and $\text{fn}(P) = \{c_i\}_i$.

Assumption:

$$\frac{\Gamma \vdash E[\mathbf{try } P[\mathbf{raise}] \mathbf{ as } x \mathbf{ in } M \mathbf{ otherwise } N] : A'}{\Gamma; \cdot \vdash^\bullet \bullet E[\mathbf{try } P[\mathbf{raise}] \mathbf{ as } x \mathbf{ in } M \mathbf{ otherwise } N]}$$

By Lemma C.4, there exist $\Gamma_1, \Gamma_2, A, B, C$ such that $\Gamma = \Gamma_1, \Gamma_2, \Gamma_3$ and

$$\frac{\Gamma_2 \vdash P[\mathbf{raise}] : A \quad \Gamma_3, x : B \vdash M : C \quad \Gamma_3 \vdash N : C}{\Gamma_2, \Gamma_3 \vdash \mathbf{try } P[\mathbf{raise}] \mathbf{ as } x \mathbf{ in } M \mathbf{ otherwise } N : C}$$

Since Γ contains only runtime names and $\text{fn}(P) = \{c_i\}_i$, we know that $\Gamma_2 = c_1 : S_1, \dots, c_n : S_n$ for some S_i .

By Lemma C.5, we have that:

$$\frac{}{\Gamma_1, \Gamma_3 \vdash E[N] : A'}$$

By repeated applications of T-ZAP and T-MIX, we have that $\Gamma_2 \vdash \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n$.

Therefore, recomposing:

$$\frac{\frac{\Gamma_1, \Gamma_3 \vdash E[N] : C}{\Gamma_1, \Gamma_3; \cdot \vdash^\bullet \bullet E[N]}}{\Gamma_1, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; \cdot \vdash^\bullet \bullet E[N] \parallel \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n}}{\frac{\frac{\frac{\frac{\frac{\Gamma_1, \Gamma_3 \vdash E[N] : C}{\Gamma_1, \Gamma_3; \cdot \vdash^\bullet \bullet E[N]}}{\Gamma_1, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; \cdot \vdash^\bullet \bullet E[N] \parallel \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n}}{\Gamma_1, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; \cdot \vdash^\bullet \bullet E[N] \parallel \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n}}{\Gamma_1, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; \cdot \vdash^\bullet \bullet E[N] \parallel \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n}}{\Gamma_1, \Gamma_3, c_1 : S_1, \dots, c_n : S_n; \cdot \vdash^\bullet \bullet E[N] \parallel \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n}}}$$

as required.

Case E-RaiseChild

$$\circ P[\mathbf{raise}] \longrightarrow \not\downarrow c_1 \parallel \cdots \parallel \not\downarrow c_n$$

Assumption:

$$\frac{\Gamma \vdash P[\mathbf{raise}] : \mathbf{1}}{\Gamma; \cdot \vdash^\circ \circ P[\mathbf{raise}]}$$

and $M \longrightarrow_M N$. By Lemma 3.1, we have that $\Gamma \vdash N : A$. Recomposing:

$$\frac{\Gamma \vdash N : A}{\Gamma; \cdot \vdash \bullet N}$$

as required. \square

C.2 Canonical Forms

Theorem 3.7: Canonical Forms *Given C such that $\Gamma; \Delta \vdash^\bullet C$, there exists some $C' \equiv C$ such that $\Gamma; \Delta \vdash^\bullet C'$ and C' is in canonical form.*

PROOF. The proof is by induction on the count of ν -bound variables, following Lindley and Morris [2015]. Without loss of generality, assume that the ν -bound variables of C are distinct. Let $\{a_i \mid 1 \leq i \leq n\}$ be the set of ν -bound variables in C and let $\{\mathcal{D}_j \mid 1 \leq j \leq m\}$ be the set of threads in C .

In the case that $n = 0$, by Lemma C.1 we can safely commute the main thread such that it is the rightmost configuration, and associate parallel composition to the right using Lemma C.2 to derive a well-typed canonical form.

In the case that $n \geq 1$, pick some a_i and \mathcal{D}_j such that a_i is the only ν -bound variable in $\text{fn}(\mathcal{D}_j)$; Lemma 3.6 and a standard counting argument ensure that such a name and configuration exist. By the equivalence rules, there exists \mathcal{E} such that $\Gamma; \Delta \vdash^\bullet C \equiv (\nu a_i)(\mathcal{D}_j \parallel \mathcal{E})$ (that a_i is the only ν -bound variable in $\text{fn}(\mathcal{D}_j)$ ensures well-typing). Moreover, we have that there exist $\Gamma' \subseteq \Gamma$, $\Delta' \subseteq \Delta$, and S , such that either $\Gamma', a_i : S; \Delta' \vdash^\bullet \mathcal{E}$ or $\Gamma'; \Delta', a_i : S \vdash^\bullet \mathcal{E}$. By the induction hypothesis, there exists \mathcal{E}' in canonical form such that either $\Gamma', a_i : S; \Delta' \vdash^\bullet \mathcal{E} \equiv \mathcal{E}'$ or $\Gamma'; \Delta', a_i : S \vdash^\bullet \mathcal{E} \equiv \mathcal{E}'$. Let $C' = (\nu a_i)(\mathcal{D}_j \parallel \mathcal{E}')$. By construction it holds that $\Gamma; \Delta \vdash^\bullet C \equiv C'$ and that C' is in canonical form. \square

C.3 Progress

To prove Theorem 3.9, we prove a similar property in which canonical configurations are decomposed step-by-step rather than in one go.

Definition C.8 (Open Progress). Suppose $\Psi; \Delta \vdash^\bullet C$, where C is in canonical form and $C \rightleftharpoons$.

We say that C satisfies open progress if:

- (1) $C = (\nu a)(\mathcal{A} \parallel \mathcal{D})$, where $\Psi = \Psi_1, \Psi_2$ and $\Delta = \Delta_1, \Delta_2$ such that either:
 - (a) $\Psi_1, a : S; \Delta_1 \vdash^\circ \mathcal{A}$ and $\Psi_2; \Delta_2, a : \bar{S} \vdash^\bullet \mathcal{D}$ where \mathcal{D} satisfies open progress, and \mathcal{A} is either:
 - (i) A thread $\circ M$ where $\text{ready}(b, M)$ for some $b \in \text{fn}(\Psi_1, a : S)$; or
 - (ii) A zipper thread $\not\downarrow a$; or
 - (iii) A buffer $b(\vec{V}) \rightleftharpoons c(\vec{W})$ where $b, c \neq a$ and either $a \in \vec{V}$ or $a \in \vec{W}$
 - (b) $\Psi_1; \Delta_1, a : \bar{S} \vdash^\circ \mathcal{A}$ and $\Psi_2, a : S; \Delta_2 \vdash^\bullet \mathcal{D}$, where \mathcal{D} satisfies open progress, and \mathcal{A} is either $a(\vec{V}) \rightleftharpoons b(\vec{W})$ or $b(\vec{V}) \rightleftharpoons a(\vec{W})$ for some $b \in \text{fn}(\Delta_1)$
- (2) $C = \mathcal{A} \parallel \mathcal{M}$, where $\Psi = \Psi_1, \Psi_2$ and either:
 - (a) $\Delta = \Delta_1, \Delta_2, a : S^\sharp$, where $\Psi_1, a : S; \Delta_1 \vdash^\circ \mathcal{A}$ and $\Psi_2; \Delta_2, a : \bar{S} \vdash^\bullet \mathcal{M}$, where \mathcal{M} satisfies open progress, and \mathcal{A} is either:
 - (i) A thread $\circ M$ where $\text{ready}(b, M)$ for some $b \in \text{fn}(\Psi_1, a : S)$; or
 - (ii) A zipper thread $\not\downarrow a$; or
 - (iii) A buffer $b(\vec{V}) \rightleftharpoons c(\vec{W})$ where $b, c \neq a$ and either $a \in \text{fn}(\vec{V})$ or $a \in \text{fn}(\vec{W})$
 - (b) $\Delta = \Delta_1, \Delta_2, a : S^\sharp$, where $\Psi_1; \Delta_1, a : \bar{S} \vdash^\circ \mathcal{A}$ and $\Psi_2, a : S; \Delta_2 \vdash^\bullet \mathcal{M}$, where \mathcal{M} satisfies open progress, and \mathcal{A} is either $a(\vec{V}) \rightleftharpoons b(\vec{W})$ or $b(\vec{V}) \rightleftharpoons a(\vec{W})$ for some $b \in \text{fn}(\Delta_1)$

- 2255 (c) $\Delta = \Delta_1, \Delta_2$, where $\Psi_1; \Delta_1 \vdash^\circ \mathcal{A}$ and $\Psi_2; \Delta_2 \vdash^\bullet \mathcal{M}$, where \mathcal{M} satisfies open progress, and
 2256 \mathcal{A} is either:
 2257 (i) A thread $\circ M$ where either $M = ()$, or $\text{ready}(a, M)$ for some $a \in \text{fn}(\Psi_1)$; or
 2258 (ii) A zipper thread $\downarrow a$ for some $a \in \text{fn}(\Psi_1)$; or
 2259 (iii) A buffer $a(\vec{V}) \leftrightarrow c(\vec{W})$ for some $a, b \in \text{fn}(\Delta_1)$
 2260 (3) $C = \mathcal{T}$, where either:
 2261 (a) $\mathcal{T} = \bullet N$, where N is either a value or $\text{ready}(b, N)$ for some $b \in \text{fn}(\Psi)$
 2262 (b) $\mathcal{T} = \mathbf{halt}$
 2263

2264 **LEMMA C.9.** *Suppose $\Psi; \Delta \vdash^\bullet C$, where C is in canonical form and $C \Rightarrow$. Then C satisfies open*
 2265 *progress.*

2266 **PROOF.** By induction on the derivation of $\Psi; \Delta \vdash^\bullet C$. We have three cases, based on the structure
 2267 of the given canonical form.
 2268

2269 **Case** $C = (va)(\mathcal{A} \parallel \mathcal{D})$, with $a \in \text{fn}(\mathcal{A})$, and where \mathcal{D} is in canonical form
 2270

2271 By assumption, we know that $\Psi; \Delta \vdash^\phi (va)(\mathcal{A} \parallel \mathcal{D})$.

2272 This configuration is typeable by T-NU, followed by either T-CONNECT₁ or T-CONNECT₂. As
 2273 the definition of canonical forms requires that $a \in \text{fn}(\mathcal{A})$, it cannot be the case that the parallel
 2274 composition arises as a result of T-MIX.

2275 We consider these two subcases to show that \mathcal{A} satisfies the properties required by open progress.

2276 **Subcase** T-CONNECT₁
 2277

$$\frac{\frac{\Psi_1, a : S; \Delta_1 \vdash^{\phi_1} \mathcal{A} \quad \Psi_2; \Delta_2, a : \bar{S} \vdash^{\phi_2} \mathcal{D}}{\Psi_1, \Psi_2; \Delta_1, \Delta_2, a : S^\# \vdash^{\phi_1 + \phi_2} \mathcal{A} \parallel \mathcal{D}}}{\Psi_1, \Psi_2; \Delta_1, \Delta_2 \vdash^{\phi_1 + \phi_2} (va)(\mathcal{A} \parallel \mathcal{D})}$$

2282 By the definition of auxiliary threads and inversion on the typing relation, we know that \mathcal{A} is of
 2283 the following forms:

- 2284 • $\circ M$, where $a \in \text{fn}(M)$, and $\Psi_1, a : S \vdash M : 1$
- 2285 • $\downarrow a$
- 2286 • $b(\vec{V}) \leftrightarrow c(\vec{W})$, where $b, c \in \text{fn}(\Delta_1)$ and $a \in \text{fn}(V)$
- 2287 • $b(\vec{V}) \leftrightarrow c(\vec{W})$, where $b, c \in \text{fn}(\Delta_1)$ and $a \in \text{fn}(W)$

2288 (since $a \notin \Delta_1$, it cannot be the case that a appears as a buffer endpoint).
 2289

2290 Lemma 3.4 tells us that either there exists some M' such that $M \rightarrow_M M'$; that M is a value; or
 2291 that M is a communication and concurrency construct. Since $C \Rightarrow$, we have that M is unable to
 2292 reduce (as otherwise C could reduce by E-LIFTM). Since $a \in \text{fn}(M)$ and a does not have type 1,
 2293 it cannot be the case that M is a value.

2294 Therefore, we have that M has the form $E[N]$, where N is a communication / concurrency
 2295 construct. This cannot be **fork**, since **fork** may always reduce by E-FORK, so there must exist
 2296 some $b \in \text{fn}(\Psi, a : S)$ such that $\text{ready}(b, M)$.
 2297

2298 **Subcase** T-CONNECT₂
 2299

$$\frac{\frac{\Psi_1; \Delta_1, a : \bar{S} \vdash^\circ \mathcal{A} \quad \Psi_2, a : S; \Delta_2 \vdash^\bullet \mathcal{D}}{\Psi_1, \Psi_2; \Delta_1, \Delta_2, a : S^\# \vdash^\bullet \mathcal{A} \parallel \mathcal{D}}}{\Psi_1, \Psi_2; \Delta_1, \Delta_2 \vdash^\bullet (va)(\mathcal{A} \parallel \mathcal{D})}$$

2304 By the definition of auxiliary threads and inversion on the typing relation, we know that \mathcal{A} is of
 2305 the following forms:

2306 • $a(\vec{V}) \rightsquigarrow b(\vec{W})$, where $b \in \Delta_1$

2307 • $b(\vec{V}) \rightsquigarrow a(\vec{W})$, where $b \in \Delta_1$

2308 (as $a \in \text{fn}(\mathcal{A})$ and $a \in \Delta_1$, it cannot be the case that \mathcal{A} is a child thread or a zipper thread, as
 2309 these require empty runtime typing environments).

2310 By the induction hypothesis, we know that \mathcal{D} satisfies open progress; hence $(\nu a)(\mathcal{A} \parallel \mathcal{D})$ satisfies
 2311 open progress.

2312 **Case** $C = \mathcal{A} \parallel \mathcal{M}$

2313 There are three subcases, based on whether the parallel composition arises as a result of T-CONNECT₁,
 2314 T-CONNECT₂, or T-MIX.

2315 **Subcase** T-CONNECT₁

$$\frac{\Psi_1, a : S; \Delta_1 \vdash^\circ \mathcal{A} \quad \Psi_2; \Delta_2, a : \bar{S} \vdash^\bullet \mathcal{M}}{\Psi_1, \Psi_2; \Delta_1, \Delta_2, a : S^\# \vdash^\bullet \mathcal{A} \parallel \mathcal{M}}$$

2316 By inversion on the typing rules, we have that \mathcal{A} may be:

2317 • A child thread $\circ M$, where $a \in \text{fn}(M)$

2318 • A zipper thread $\not\downarrow a$

2319 • A buffer $b(\vec{V}) \rightsquigarrow c(\vec{W})$, where $b, c \neq a$ and either $a \in \text{fn}(\vec{V})$ or $a \in \text{fn}(\vec{W})$

2320 In the case of (1), by Lemma 3.4, we have that either M is a value; there exists N such that
 2321 $M \rightarrow_M N$; or $M = E[N]$ for some E, N , where N is a communication / concurrency construct.

2322 By T-CHILD, $\Psi_1, a : S \vdash M : 1$. Since $a \in \text{fn}(M)$ and the only value with type 1 is the unit
 2323 value $()$ it therefore cannot be the case that M is a value. Since $C \not\Rightarrow$, it cannot be the case that
 2324 $M \rightarrow_M N$, since otherwise C could reduce. Thus, it must be the case that $M = E[N]$ where N is
 2325 a communication and concurrency construct; by similar reasoning as above cases, it therefore
 2326 must be the case that $\text{ready}(b, M)$ for some $b \in \text{fn}(\Psi_1, a : S)$.

2327 (2) and (3) satisfy the required conditions by definition.

2328 **Subcase** T-CONNECT₂

$$\frac{\Psi_1; \Delta_1, a : \bar{S} \vdash^\circ \mathcal{A}; \Psi_2, a : S; \Delta_2 \vdash^\bullet \mathcal{M}}{\Psi_1, \Psi_2; \Delta_1, \Delta_2, a : S^\# \vdash^\bullet \mathcal{A} \parallel \mathcal{M}}$$

2329 Since the runtime typing environment $\Delta_1, a : \bar{S}$ is non-empty, it cannot be the case that \mathcal{A} is a
 2330 child thread or zipper thread. Thus, \mathcal{A} must either be of the form:

2331 (1) $a(\vec{V}) \rightsquigarrow b(\vec{W})$, where $a, b \in \Delta_1$; or

2332 (2) $b(\vec{V}) \rightsquigarrow a(\vec{W})$, where $a, b \in \Delta_1$

2333 which satisfy the required conditions by definition.

2334 **Subcase** T-MIX

$$\frac{\Psi_1; \Delta_1 \vdash^\circ \mathcal{A} \quad \Psi_2; \Delta_2 \vdash^\bullet \mathcal{M}}{\Psi_1, \Psi_2; \Delta_1, \Delta_2 \vdash^\bullet \mathcal{A} \parallel \mathcal{M}}$$

2335 By inversion on the typing rules, we have that \mathcal{A} may either be:

2336 (1) A child thread $\circ M$

2337 (2) A zipper thread $\not\downarrow a$ for some $a \in \text{fn}(\Psi_1)$

2338

(3) A buffer thread $a(\vec{V}) \leftrightarrow b(\vec{W})$ for some $a, b \in \text{fn}(\Delta_1)$

By Lemma 3.4, we have that M is either a value V ; there exists some N such that $M \longrightarrow_M N$; or $M = E[N]$ for some E, N such that N is a communication and concurrency primitive. It cannot be the case that $M \longrightarrow_M N$ since otherwise the configuration could reduce.

By T-CHILD, it must be the case that $\Psi_1; \Delta_1 \vdash M : \mathbf{1}$; if M is a value then by inversion on the term typing rules, it must be the case that $M = ()$.

Following the same reasoning as previous cases, if $M = E[N]$ for some communication / concurrency primitive N , it must be that $\text{ready}(a, M)$ for some $a \in \Psi_1$.

By the induction hypothesis, we know that \mathcal{M} satisfies open progress; hence $\mathcal{A} \parallel \mathcal{M}$ satisfies open progress.

Case $C = \mathcal{T}$

Assumption: $\Psi; \Delta \vdash^\bullet \mathcal{T}$. By the definition of \mathcal{T} , we have two subcases:

Subcase $\mathcal{T} = \bullet M$

$$\frac{\Psi \vdash M : A}{\Psi; \cdot \vdash^\bullet \bullet M}$$

By Lemma 3.4, we have that either M is a value; that there exists some N such that $M \longrightarrow_M N$; or that there exist some E, N such that $M = E[N]$ where N is a communication / concurrency primitive.

Again, as $C \Rightarrow$, it cannot be the case that $M \longrightarrow_M N$, since otherwise C could reduce. If M is a value, then \mathcal{T} satisfies open progress.

Finally, if $M = E[N]$ where N is a communication / concurrency primitive, it cannot be the case that $N = \text{fork } M'$ since it could reduce by T-FORK, and so it must be the case that $\text{ready}(a, M)$ for some $a \in \Psi$, satisfying open progress, as required.

Subcase $\mathcal{T} = \text{halt}$

Immediate by the definition of open progress. □

Theorem 3.9 provides a more global and concise view of the properties exhibited by a non-reducing process in canonical form, and arises as an immediate corollary.

Theorem 3.9 *Suppose $\Psi; \Delta \vdash^\bullet C$ where C is in canonical form and $C \Rightarrow$.*

Let $C = (va_1)(\mathcal{A}_1 \parallel (va_2)(\mathcal{A}_2 \parallel \dots \parallel (va_n)(\mathcal{A}_n \parallel \mathcal{M})) \dots)$.

Either there exists some C' such that $C \Rightarrow C'$, or:

(1) *For $1 \leq i \leq n$, each thread in \mathcal{A}_i is either:*

- (a) *a child thread $\circ M$ for which there exists $a \in \{a_j \mid 1 \leq j \leq i\} \cup \text{fn}(\Psi)$ such that $\text{ready}(a, M)$;*
- (b) *a zipper thread ζa_i ; or*
- (c) *a buffer.*

(2) *$\mathcal{M} = \mathcal{A}'_1 \parallel \dots \parallel \mathcal{A}'_m \parallel \mathcal{T}$ such that for $1 \leq j \leq m$:*

(a) *\mathcal{A}'_j is either:*

- (i) *a child thread $\circ N$ such that $N = ()$ or $\text{ready}(a, N)$ for some $a \in \{a_i \mid 1 \leq i \leq n\} \cup \text{fn}(\Psi) \cup \text{fn}(\Delta)$;*
- (ii) *a zipper thread ζa for some $a \in \{a_i \mid 1 \leq i \leq n\} \cup \text{fn}(\Psi) \cup \text{fn}(\Delta)$; or*
- (iii) *a buffer.*

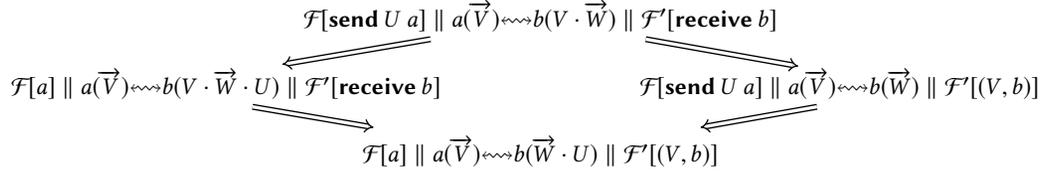
(b) *Either $\mathcal{T} = \bullet N$, where N is either a value or $\text{ready}(a, N)$ for some $a \in \{a_i \mid 1 \leq i \leq n\} \cup \text{fn}(\Psi) \cup \text{fn}(\Delta)$; or $\mathcal{T} = \text{halt}$.*

C.4 Confluence

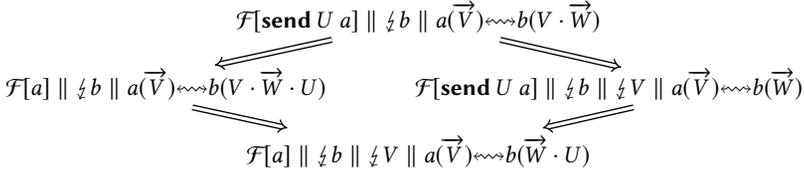
Theorem 3.12 (Diamond Property) *If $\Psi; \Delta \vdash^\phi C$, and $C \Longrightarrow \mathcal{D}_1$, and $C \Longrightarrow \mathcal{D}_2$, then either $\mathcal{D}_1 \equiv \mathcal{D}_2$, or there exists some \mathcal{D}_3 such that $\mathcal{D}_2 \Longrightarrow \mathcal{D}_3$ and $\mathcal{D}_1 \Longrightarrow \mathcal{D}_3$.*

PROOF. As noted in Section 3.4, \longrightarrow_M is deterministic and hence confluent due to the setup of term evaluation contexts, and linearity ensures that endpoints to a buffer may not be shared. Consequently, communication actions on different channels may be performed in any order.

Nevertheless, two critical pairs arise due to asynchrony. The first arises when it is possible to send to or receive from a buffer; there is a choice of whether the send or the receive happens first. Both cases reduce to the same configuration after a single further step.



The second critical pair arises when sending to a buffer where the peer endpoint has a non-empty buffer and has been cancelled. There is a choice as to whether the value at the head of the queue is cancelled before or after the send takes place. Again, both cases reduce to the same configuration after a single further step.



□

D SUPPLEMENT TO SECTION 4.1 (METATHEORY OF EGV WITH ACCESS POINTS)

In this section, we prove that the extension of EGV with access points satisfies preservation.

LEMMA D.1 (PRESERVATION, ACCESS POINTS (EQUIVALENCE)). *If $\Gamma; \Delta \vdash^\phi C$ and $C \equiv \mathcal{D}$, then $\Gamma; \Delta \vdash^\phi \mathcal{D}$*

PROOF. By induction on the derivation of $C \equiv \mathcal{D}$. Rule T-CONNECTN subsumes T-CONNECT₁ and T-CONNECT₂, so the majority of cases are similar to those we have proven in Lemma C.1. We consider the case for associativity in detail.

Case $C \parallel (\mathcal{D} \parallel \mathcal{E}) \equiv (C \parallel \mathcal{D}) \parallel \mathcal{E}$

$$\begin{array}{c}
 \Gamma = \Gamma_1 + \Gamma' \\
 \Gamma' = \Gamma_2 + \Gamma_3 \\
 \Gamma_2, b_1 : T_1, \dots, b_{m'} : T_{m'}, c : S'; \Delta_2, a_1 : \overline{S}_1, \dots, a_m : \overline{S}_m, d : \overline{T'} \vdash^{\phi_2} \mathcal{D} \\
 \Gamma_3, b_{m'+1}, \dots, b_{n'} : T_{n'}, d : \overline{T'}; \Delta_3, a_{m+1} : \overline{S}_{m+1}, \dots, a_n : \overline{S}_n, c : S' \vdash^{\phi_3} \mathcal{E} \\
 \hline
 \Gamma_1, a : \overline{S}; \Delta_1, b : \overline{T} \vdash^{\phi_1} C \quad \Gamma', b : \overline{T}; \Delta_2, \Delta_3, a : \overline{S}, c : S', d : \overline{T} \vdash^{\phi_2 + \phi_3} \mathcal{D} \parallel \mathcal{E} \\
 \hline
 \Gamma; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp, c : S', d : \overline{T} \vdash^{\phi_1 + \phi_2 + \phi_3} C \parallel (\mathcal{D} \parallel \mathcal{E})
 \end{array}$$

where

$$\begin{array}{c}
 \overline{a : \overline{S}} = a_1 : \overline{S}_1, \dots, a_m : \overline{S}_m, \dots, a_n : \overline{S}_n \\
 \overline{b : \overline{T}} = b_1 : T_1, \dots, b_{m'} : T_{m'}, \dots, b_{n'} : T_{n'}
 \end{array}$$

$$\iff$$

$$\Gamma = \Gamma'' + \Gamma_3$$

$$\begin{array}{c}
 \Gamma'' = \Gamma_1 + \Gamma_2 \\
 \Gamma_1, a : \overline{S}; \Delta_1, b : \overline{T} \vdash^{\phi_1} C \\
 \hline
 \Gamma_2, b_1 : T_1, \dots, b_m : T_m, c : S'; \Delta_2, d : \overline{T} \vdash^{\phi_2} \mathcal{D} \\
 \hline
 \Gamma'', a_{m+1} : S_{m+1}, \dots, a_n : S_n \\
 \Delta_1, \Delta_2, a_1 : S_1^\sharp, \dots, a_m : S_m^\sharp, \\
 b_1 : T_1^\sharp, \dots, b_{m'} : T_{m'}^\sharp; \\
 \hline
 \Gamma_3, b_{m'+1}, \dots, b_n : T_n, d : \overline{T'}; \\
 \hline
 \overline{b_{m'+1} : T_{m'+1}, \dots, b_{n'} : T_{n'}} \vdash^{\phi_1 + \phi_2} C \parallel \mathcal{D} \quad \Delta_3, a_{m+1} : \overline{S}_{m+1}, \dots, a_n : \overline{S}_n, c : S' \vdash^{\phi_3} \mathcal{E} \\
 \hline
 \Gamma; \Delta_1, \Delta_2, \Delta_3, a : S^\sharp, b : T^\sharp, c : S', d : \overline{T} \vdash^{\phi_1 + \phi_2 + \phi_3} C \parallel (\mathcal{D} \parallel \mathcal{E})
 \end{array}$$

□

The lemmas for subterm typeability and replacement are slightly different as we must consider unrestricted environments.

LEMMA D.2 (TYPEABILITY OF SUBTERMS (ACCESS POINTS)). *If \mathbf{D} is a derivation of $\Gamma \vdash E[M] : A$, then there exist Γ_1, Γ_2 and B such that $\Gamma = \Gamma_1 + \Gamma_2$, that \mathbf{D} has a subderivation \mathbf{D}' that concludes $\Gamma_2 \vdash M : B$, and the position of \mathbf{D}' in \mathbf{D} corresponds to the position of the hole in E .*

PROOF. By induction on the structure of E . □

LEMMA D.3 (REPLACEMENT (ACCESS POINTS)). *If:*

- \mathbf{D} is a derivation of $\Gamma \vdash E[M] : A$, such that $\Gamma = \Gamma_1 + \Gamma_2$
- \mathbf{D}' is a subderivation of \mathbf{D} concluding $\Gamma_2 \vdash M : B$

- 2500 • The position of D' in D corresponds to that of the hole in E
- 2501 • $\Gamma_3 \vdash N : B$
- 2502 • $\Gamma' = \Gamma_1 + \Gamma_3$ is defined

2503 then $\Gamma' \vdash E[N] : A$.

2504 PROOF. By induction on the structure of E . □

2505 THEOREM D.4 (PRESERVATION, ACCESS POINTS). *If $\Gamma; \Delta \vdash^\phi C$ and $C \implies \mathcal{D}$, then $\Gamma; \Delta \vdash^\phi \mathcal{D}$.*

2506 PROOF. Recall that \implies is defined as $\equiv \longrightarrow \equiv$. Therefore, the result arises by induction on the
2509 derivation of $C \longrightarrow \mathcal{D}$ and as a corollary of Lemma D.1.

2510 Again, since T-CONNECTN subsumes T-CONNECT₁ and T-CONNECT₂, it suffices only to prove the
2511 new cases required for access point reduction.

2513 Case E-Spawn

2514 Assumption:

$$2516 \frac{\Gamma \vdash E[\mathbf{spawn} M] : C}{2517 \Gamma; \cdot \vdash^\bullet \bullet E[\mathbf{spawn} M]}$$

2518 By Lemma D.2, we have that $\Gamma = \Gamma_1 + \Gamma_2$, and

$$2520 \frac{\Gamma_2 \vdash M : 1}{2521 \Gamma_2 \vdash \mathbf{spawn} M : 1}$$

2522 By Lemma D.3, we have that $\Gamma_1 \vdash E[()] : C$.

2523 Recomposing:

$$2526 \frac{\Gamma = \Gamma_1 + \Gamma_2 \quad \Gamma_1; \cdot \vdash^\bullet E[()] \quad \Gamma_2; \cdot \vdash^\circ \circ M}{2527 \Gamma; \cdot \vdash^\bullet E[()] \parallel \circ M}$$

2528 as required.

2530 Case E-New

2531 Assumption:

$$2533 \frac{\Gamma \vdash E[\mathbf{new}_S] : C}{2534 \Gamma; \cdot \vdash^\bullet \bullet E[\mathbf{new}_S]}$$

2535 By Lemma D.2 and TA-NEW, we have that $\cdot \vdash \mathbf{new}_S : \text{AP}(S)$.

2536 By Lemma D.3, we have that $\Gamma, z : \text{AP}(S) \vdash E[z] : C$.

2537 Thus, we can show:

$$2540 \frac{\Gamma, z : \text{AP}(S) \vdash E[z] : C}{2541 \Gamma, z : \text{AP}(S); \cdot \vdash^\bullet \bullet E[z] \quad \cdot; z : \text{AP}(S) \vdash^\circ z(\epsilon, \epsilon)}$$

$$2542 \frac{\Gamma, z : \text{AP}(S); z : S \vdash^\bullet \bullet E[z] \parallel z(\epsilon, \epsilon)}{2543 \Gamma; \cdot \vdash^\bullet (vz)(\bullet E[z]) \parallel z(\epsilon, \epsilon)}$$

2544 as required.

2547 Case E-Accept

2548

2549 Assumption:

$$\frac{\Gamma \vdash E[\mathbf{accept} z] : C}{\Gamma; \cdot \vdash^\bullet \bullet E[\mathbf{accept} z]} \quad ; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S \vdash^\circ z(\mathcal{X}, \mathcal{Y})$$

$$\Gamma; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S \vdash^\bullet \bullet E[\mathbf{accept} z] \parallel z(\mathcal{X}, \mathcal{Y})$$

2554 By Lemma D.2, we have that $\Gamma = \Gamma_1 + \Gamma_2$ and that $\Gamma_2 \vdash \mathbf{accept} z : S$. Thus by TA-ACCEPT we have
2555 that $z : \text{AP}(S) \in \Gamma$.

2556 By Lemma D.3, we have that $\Gamma, a : S \vdash E[a] : C$.

2557 Recomposing, we have that:

$$\frac{\Gamma, a : S \vdash E[a] : C}{\Gamma, a : S; \cdot \vdash^\bullet \bullet E[a]} \quad ; z : S, \mathcal{X} : \bar{S}, a : \bar{S}, \mathcal{Y} : S \vdash^\circ z(\{a\} \cup \mathcal{X}, \mathcal{Y})$$

$$\frac{\Gamma; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S, a : S^\# \vdash^\bullet \bullet E[a] \parallel z(\{a\} \cup \mathcal{X}, \mathcal{Y})}{\Gamma; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S \vdash^\bullet (va)(\bullet E[a]) \parallel z(\{a\} \cup \mathcal{X}, \mathcal{Y})}$$

2564 **Case E-Request**

2565 Assumption:

$$\frac{\Gamma \vdash E[\mathbf{request} z] : C}{\Gamma; \cdot \vdash^\bullet \bullet E[\mathbf{request} z]} \quad ; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S \vdash^\circ z(\mathcal{X}, \mathcal{Y})$$

$$\Gamma; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S \vdash^\bullet \bullet E[\mathbf{request} z] \parallel z(\mathcal{X}, \mathcal{Y})$$

2571 By Lemma D.2, we have that $\Gamma = \Gamma_1 + \Gamma_2$ and that $\Gamma_2 \vdash \mathbf{request} z : \bar{S}$. Thus by TA-REQUEST we
2572 have that $z : \text{AP}(S) \in \Gamma$.

2573 By Lemma D.3, we have that $\Gamma, a : \bar{S} \vdash E[a] : C$. As duality is involutive, we have that $\bar{\bar{S}} = S$.

2574 Recomposing, we have that:

$$\frac{\Gamma, a : \bar{S} \vdash E[a] : C}{\Gamma, a : \bar{S}; \cdot \vdash^\bullet \bullet E[a]} \quad ; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S, a : S \vdash^\circ z(\mathcal{X}, \{a\} \cup \mathcal{Y})$$

$$\frac{\Gamma; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S, a : S^\# \vdash^\bullet \bullet E[a] \parallel z(\mathcal{X}, \{a\} \cup \mathcal{Y})}{\Gamma; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S \vdash^\bullet (va)(\bullet E[a]) \parallel z(\mathcal{X}, \{a\} \cup \mathcal{Y})}$$

2582 as required.

2583 **Case E-Match**

2584 Assumption:

$$\frac{}{; z : S, a : \bar{S}, \mathcal{X} : \bar{S}, b : S, \mathcal{Y} : S \vdash^\circ z(\{a\} \cup \mathcal{X}, \{b\} \cup \mathcal{Y})}$$

2588 Recomposing:

$$\frac{}{; z : S, \mathcal{X} : \bar{S}, \mathcal{Y} : S \vdash^\circ z(\mathcal{X}, \mathcal{Y})} \quad \frac{\bar{S}/\epsilon = \bar{S}/\epsilon \quad \cdot \vdash \epsilon : \epsilon \quad \cdot \vdash \epsilon : \epsilon}{; a : \bar{S}, b : S \vdash^\circ a(\epsilon) \rightsquigarrow b(\epsilon)}$$

$$\frac{}{; z : S, a : \bar{S}, \mathcal{X} : \bar{S}, b : S, \mathcal{Y} : S \vdash^\circ z(\mathcal{X}, \mathcal{Y}) \parallel a(\epsilon) \rightsquigarrow b(\epsilon)}$$

2595 \square

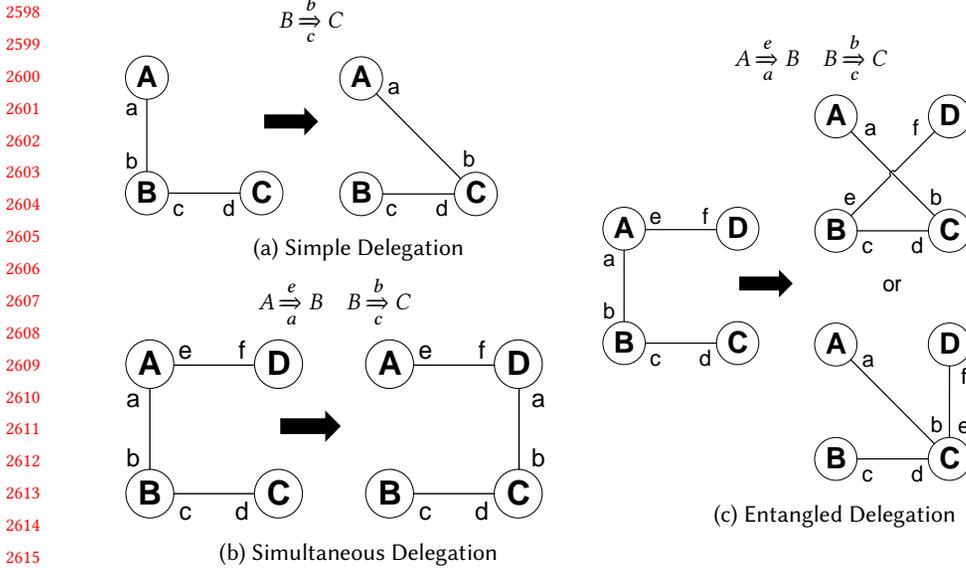


Fig. 13. Cases of Distributed Delegation

E DISTRIBUTED DELEGATION

A key feature of π -calculus is *mobility*, that is, sending channel names as values. In session-based languages and calculi, mobility is realised as *session delegation*, allowing session-typed channel endpoints to be sent over other session-typed channels. We saw an example of session delegation in §6, in the ChatClient type:

```

typename ChatClient = !Nickname.
  [&|Join:?(Topic, [Nickname], ClientReceive).ClientSend,
   Nope:End|&];

```

An endpoint of type `ClientReceive` is passed as a message.

E.1 Challenges of Distributed Delegation

Session delegation is a vital abstraction in session-based programming. However, its integration with both asynchrony *and* distribution brings several challenges. The seminal work on distributed delegation is Session Java [Hu et al. 2008].

Fig. 13 shows three scenarios of distributed delegation, as described by Hu et al. [2008]. We write $X \xrightarrow[x]{y} Y$ to indicate that X wishes to send x to Y over y on the basis that X 's last known

location of the corresponding endpoint for y is Y . Now suppose $B \xrightarrow[b]{c} C$. Following Hu et al. [2008], we refer to B as the *session-sender*, C as the *session-receiver*, and A as a *passive party*. There is no happens-before relation between A sending a message to B along a , and B delegating b to C along c . Thus, a message could be sent to A *after* A has given up control of a . Following Hu et al. [2008], we call such messages *lost messages*.

1. $A \rightarrow S : \text{Send}(t, v, [b \mapsto \vec{V}])$
2. $A : \text{start recording lost messages } \vec{W} \text{ for } b$
3. $S : \sigma = \sigma[b \mapsto B]; \delta = \delta \cup \{t\}$
4. $S \rightarrow B : \text{Deliver}(t, v, [b \mapsto \vec{V}])$
5. $S \rightarrow A : \text{GetLostMessages}([b])$
6. $A : \text{stop recording lost messages for } b$
7. $A \rightarrow S : \text{LostMessageResponse}([b \mapsto \vec{W}])$
8. $S \rightarrow B : \text{Commit}(t, [b \mapsto \vec{W}])$
9. $S : \delta = \delta \setminus \{t\}$
10. $B : \text{buffers}[b] = \vec{V} \uplus \vec{W} \uplus \vec{U}$
 where $\vec{U} = \text{messages received for } b \text{ between (3) and (8)}$

Fig. 14. Operation of Distributed Delegation Protocol

E.2 Approaches to Distributed Delegation

The simplest safe way to implement distributed delegation is to store all buffers on the server, but this requires a blocking remote call for every receive operation. A second naïve method is *indefinite redirection*, where the session-sender indefinitely forwards all messages to the session-receiver. This retains buffer locality, but requires the session-sender to remain online for the duration of the delegated session.

Hu et al. [2008] describe two more realistic distributed delegation algorithms: a *resending* protocol, which re-sends lost messages *after* a connection for the delegated session is established, and a *forwarding* protocol, which forwards lost messages *before* the delegated session is established. The key idea behind both algorithms is to establish a connection between the passive party and the session-receiver, ensure that the lost messages are received by the session-receiver, and to continue the session only once lost messages are received.

E.3 Delegation in Distributed Session Links

Alas, we cannot directly re-use the resending and forwarding protocols of Hu et al. [2008] because of two fundamental differences in our setting: Links clients do not connect to each other directly, and in Links multiple sessions may be sent at once. Thus, we describe the high-level details of a modified algorithm which addresses these two constraints. We utilise two key ideas:

- Much like the resending protocol, lost messages are retrieved and relayed to the session-receiver once the new session has been established.
- We ensure the session-receiver endpoint is not delegated until the delegation has completed, by queuing messages that include the session-receiver endpoint, and resending them once delegation has completed.

We now consider the case where session-sender and session-receiver are different clients; the case where session-sender is a client and session-receiver the server is similar. Let client A be session-sender and client B be session-receiver.

Example. Suppose client A sends a value v containing a session endpoint d along channel (s, t) , recalling that s is the peer endpoint and t is the local endpoint. The initial endpoint location table is:

$$\sigma \triangleq [s \mapsto A, t \mapsto B, b \mapsto A, c \mapsto A]$$

2696 Fig. 14 shows the operation of the delegation protocol on this example. In Step 1, A sends a message
 2697 to the server S , containing the peer endpoint t , value to send v , and the buffer \vec{V} for b , before
 2698 beginning to record lost messages for b . Upon receiving this message, the server updates its internal
 2699 mapping for the location of b to be B , adds t to the set of delegation carriers δ , and sends a Deliver
 2700 message containing t , v , and \vec{V} , before sending a GetLostMessages request to A . Upon receiving
 2701 this message, A will stop recording lost messages for b , and relay the lost messages \vec{W} for b to
 2702 S . The server then sends a Commit message containing t and the lost messages for all delegated
 2703 endpoints, and removes t from the set of delegation carriers.
 2704

2705 The final buffer for b is the concatenation of the initial buffer \vec{V} , the lost messages \vec{W} , and all
 2706 messages \vec{U} received for b before the Commit message.
 2707

2708 E.4 Correctness

2709 We argue correctness of the algorithm in a similar manner to Hu et al. [2008]. Due to co-operative
 2710 threading, we can treat each sequence of actions happening at a single participant (for example,
 2711 steps 3–8) as atomic. Since (as per step 3) the endpoint location table is updated prior to the lost
 2712 message request, we can safely split the buffer of the delegated session into three parts: the initial
 2713 buffer being delegated (\vec{V}); the lost messages (\vec{W}); and the messages received after the change in
 2714 the lookup table but before the Commit message is received (\vec{U}) and reassemble them, retaining
 2715 ordering.
 2716

2717 In our setting, since session channels are not associated with sockets, simultaneous delegation
 2718 (Fig. 13b) can be handled in the same way as simple delegation. In the case of entangled delegation
 2719 (Fig. 13c), since delegation carriers may not be delegated themselves until the lost messages have
 2720 been received, we can be sure that the lost message requests are sent to the correct participant.
 2721 Hence, the case devolves to simple delegation.
 2722
 2723
 2724
 2725
 2726
 2727
 2728
 2729
 2730
 2731
 2732
 2733
 2734
 2735
 2736
 2737
 2738
 2739
 2740
 2741
 2742
 2743
 2744