Connectivity Graphs: A Method for Proving Deadlock Freedom Based on Separation Logic

JULES JACOBS, Radboud University Nijmegen, The Netherlands

STEPHANIE BALZER, Carnegie Mellon University, USA

ROBBERT KREBBERS, Radboud University Nijmegen, The Netherlands

We introduce the notion of a *connectivity graph*—an abstract representation of the topology of concurrently interacting entities, which allows us to encapsulate generic principles of reasoning about *deadlock freedom*. Connectivity graphs are *parametric* in their vertices (representing entities like threads and channels) and their edges (representing references between entities) with labels (representing interaction protocols). We prove deadlock and memory leak freedom in the style of progress and preservation and use *separation logic* as a meta theoretic tool to treat connectivity graph edges and labels substructurally. To prove preservation locally, we distill generic separation logic rules for *local graph transformations* that preserve acyclicity of the connectivity graph. To prove global progress locally, we introduce a *waiting induction* principle for acyclic connectivity graphs. We mechanize our results in Coq, and instantiate our method with a higher-order binary session-typed language to obtain the first mechanized proof of deadlock and leak freedom.

CCS Concepts: • **Theory of computation** → **Separation logic**; • **Software and its engineering** → **Concurrent programming languages**.

Additional Key Words and Phrases: Message passing, concurrency, session types, separation logic, deadlock freedom, graphs, Coq

ACM Reference Format:

Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. 2022. Connectivity Graphs: A Method for Proving Deadlock Freedom Based on Separation Logic. *Proc. ACM Program. Lang.* 6, POPL, Article 1 (January 2022), 33 pages. <https://doi.org/10.1145/3498662>

1 INTRODUCTION

Binary session types [Honda 1993; Honda et al. 1998] are a type discipline for specifying protocols of interactions in message-passing concurrent programs. Session types have turned into an active area of research that enjoys strong theoretical and practical foundations. The theoretical foundations include a Curry-Howard correspondence between session-typed π -calculi and linear logic [Caires and Pfenning 2010; Wadler 2012; Caires et al. 2013; Pérez et al. 2014; Toninho et al. 2013; Lindley and Morris 2015; Toninho 2015] and session-typed λ -calculi with mainstream programming language features [Lindley and Morris 2016b, 2017; Igarashi et al. 2017; Fowler et al. 2019]. The practical foundations include libraries for session types in mainstream programming languages [Dezani-Ciancaglini et al. 2006; Pucella and Tov 2008; Imai et al. 2010; Jespersen et al. 2015; Lindley and Morris 2016a; Scalas and Yoshida 2016; Padovani 2017; Imai et al. 2019; Kokke 2019; Chen and Balzer 2020].

Authors' addresses: Jules Jacobs, Radboud University Nijmegen, The Netherlands, j.jacobs@science.ru.nl; Stephanie Balzer, Carnegie Mellon University, USA, balzers@cs.cmu.edu; Robbert Krebbers, Radboud University Nijmegen, The Netherlands, mail@robbertkrebbers.nl.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/1-ART1

<https://doi.org/10.1145/3498662>

Fig. 1. Connectivity graphs with run-time information for our example program (the End markers have been elided from session types). Boxes depict threads (red boxes are blocked threads, green dotted boxes are running threads). Blue circles depict channels. Black edges indicate references to channel endpoints, labeled with their session type. Red triangles reveal the waiting dependency for each reference to a channel endpoint: either the owner of the endpoint is waiting to receive a message from the channel, or the channel is waiting for the owner of the endpoint to initiate the next action (send or receive or close).

The main thread (T) uses the **fork** construct to spawn two threads (T1 and T2) with bidirectional channels (C1 and C2) connecting them to the main thread. The endpoints c_2 and c_2' of channel C2 (created on [Line 5](#)) have session types $(!N. \text{End})$ and $(?N. \text{End})$, respectively. These dual session types express that a number should be sent (!) over c_2 and received (?) over c_2' . The session types of channel C1 (created on [Line 1](#)) are more interesting—they are higher-order. Endpoint c_1 has session type $(!(?N. \text{End}). ?N. \text{End})$, which expresses that first a channel of type $(?N. \text{End})$ should be sent, after which a number can be received. The λ -expression of thread T2 captures endpoint c_1 , resulting in the ownership of c_1 being transferred from thread T to thread T2.

The first picture in [Figure 1](#) displays the connectivity graph after both **forks** have been executed, but no other steps have been performed yet. The solid red boxes correspond to threads that are blocked on a receive, while the dotted green boxes correspond to threads that can make progress. The small black arrowheads on the edges indicate the direction of channel ownership: an edge from a thread to a channel indicates that the thread owns an endpoint of that channel, and an edge from a channel to a channel indicates that an endpoint of the latter channel is stored in one of the buffers of the former channel (an edge between two threads is not possible—all references are *to* channels). The red triangles denote the waiting dependency. A crucial property of the connectivity graph is that the waiting dependency remains acyclic. Acyclicity enables us to find a thread that can make progress by starting at any vertex and repeatedly following the red triangles. For example, when starting at thread T1, which is blocked, we find that thread T2 can perform a step.

When we continue by letting thread T2 perform the **send** operation on [Line 6](#), the **send** will move the endpoint c_2' into the buffer of C1. In general, the effect of the **send** operation on the connectivity graph is as follows:

On the left, thread T has ownership of the transmitted value v and the endpoint of carrier channel C with session type $!\tau.s$. The value v could in general contain any number of channel references (depicted as C_1, C_2, C_3), for instance when it is a pair of channels or a closure that has captured more than one channel. On the right, we have the resulting connectivity graph that we obtain after the **send** operation has been performed. The session type changes to s (*i.e.*, the remainder of the

protocol) and the value v gets transferred to the buffer of C , so the incoming edges of the channels in v are changed from T to C . Note that the red waiting triangles and the information about whether a thread is blocked is not part of the connectivity graph because it can be derived from the run-time configuration. We therefore depict the general transformation rule without waiting triangles and use a neutral color for threads. In our running example, the thread is T_2 , the channel is C_1 , and the value v is the single channel C_2 . The second picture in Figure 1 displays the resulting connectivity graph: the session type of T_2 has advanced to $(?N.\text{End})$ and the incoming edge from C_2 to T_2 has turned into an incoming edge from C_2 to C_1 .

Next, we let thread T_1 perform the `recv` operation on Line 2, which will move the endpoint out of the buffer of channel C_1 and bind it to variable c . The rule to transform the connectivity graph for a `receive` operation is similar to `send` (the exact rule can be found in Figure 8 in §3.5). The third picture in Figure 1 displays the resulting connectivity graph, where we see that the session type of T_1 advanced, and that the outgoing edge from C_1 to C_2 has turned into an outgoing edge from T_1 to C_2 . Observe that the connectivity graph has a non-trivial structure—to find a thread that can unblock T_2 , we need to follow multiple edges to end up in thread T .

Progress and preservation. As shown by the examples in Figure 1, connectivity graphs describe the types and abstract reference topology of a program’s execution configuration, but not the concrete expressions and values that constitute the threads and channels. To prove a property about the operational semantics, we need to define a relation that expresses that a configuration ρ is well-formed w.r.t. a connectivity graph G . With that relation at hand, we can carry out a proof in the style of progress and preservation [Wright and Felleisen 1994; Harper 2016; Pierce 2002].

- **Progress:** If ρ is well-formed w.r.t. G , then either ρ is final (all threads have terminated and all channels have been deallocated), or ρ can step (deadlock freedom).
- **Preservation:** If ρ is well-formed w.r.t. G , and can take a step $\rho \rightsquigarrow \rho'$ in the operational semantics, then we can transform G into G' so that ρ' is well-formed w.r.t. G' .

It is important to point out that connectivity graphs generalize *heap typings* from traditional progress and preservation proofs for type systems with mutable references [Pierce 2002; Harper 2016]. Whereas heap typings are flat (they merely give the types of channels, which suffices for type safety), connectivity graphs additionally describe the reference topology and ensure its acyclicity (needed for deadlock and memory leak freedom).

Connectivity graphs as a parametric proof principle. When trying to formalize the above reasoning, we encounter two problems:

- (1) Due to linear types and concurrency, it is non-trivial to formalize the well-formedness relation of configurations w.r.t. connectivity graphs. Definitions easily end up cluttered with details about *resource separation*, which burdens mechanization in a proof assistant.
- (2) Proving preservation and progress involves non-trivial reasoning about graphs. For preservation we need to *transform* graphs (to type a post-configuration), and for progress we need to *traverse* graphs (to find a thread that can step). Reasoning about graphs is difficult in a proof assistant because graphs are not inductively defined.

To address these problems, we use separation logic [O’Hearn et al. 2001] as a meta theoretic tool to reason about graphs. Traditionally, separation logic is used as a specification language to write pre- and postconditions for individual programs in Hoare-style logics. Inspired by recent work that uses separation logic to establish type safety using logical relations [Krebbers et al. 2017; Jung et al. 2018a; Hinrichsen et al. 2021b] and intrinsically-typed interpreters and compilers [Rouvoet et al. 2020, 2021], we also use separation logic but in the context of a progress and preservation proof.

Our version of separation logic makes it possible to define the well-formedness relation in a way that is local (*i.e.*, talks about threads in isolation) and that hides resources. To adopt separation logic for our connectivity graph, we must decide on what to consider as a resource. The scenarios in Figure 1 suggest that we should consider a vertex's *outgoing edges* as resources, because then a graph transformation, such as the one induced by moving endpoint $c2'$ into $C1$'s buffer, simply amounts to an ownership transfer. To prove preservation, we distill a set of separation logic rules for reasoning about *graph transformations* by simply transferring ownership of resources. To prove progress, we distill a form of *waiting induction* to perform induction on the connectivity graph to find a vertex that can perform a step.

All ingredients of our method (the definition of connectivity graph, the separation logic, the graph transformations, and waiting induction) are parametric in the vertices, edges and labels of the connectivity graph. This is crucial for mechanization: we can encapsulate our proof method as a library that is independent of the concrete programming language. We use our library in combination with the Iris Proof Mode [Krebbers et al. 2017, 2018]—which provides tactics for separation-logic based reasoning—to effectively hide reasoning about graphs and resources in Coq.

Contributions. We present a parametric method for proving deadlock and memory leak freedom of binary linear session-typed languages. Concretely:

- We introduce *connectivity graphs* as a generalization of heap typings in progress and preservation proofs. In addition to typing, connectivity graphs track the reference topology.
- We show how to use separation logic in a non-standard way as a language for linking our abstract connectivity graphs to a concrete language's operational semantics and type system.
- We implement connectivity graphs as a library in the Coq proof assistant that is parametric in the type of vertices and edges. Our library includes *graph transformations* as separation logic rules to aid proving preservation, and a principle of *waiting induction* over connectivity graphs to aid proving progress.
- We use our connectivity graph library to obtain the first mechanized proof of deadlock and leak freedom for a binary session-typed λ -calculus with higher-order channels, recursive types, and unrestricted types.

We start by introducing our language (§2), and explain our key ideas by proving deadlock freedom for it (§3). Next, we present the parametric aspects of our proof method (§4 and 5). We then add extensions to our language, and prove a stronger deadlock and memory leak freedom property than the conventional formulation in terms of global progress (§6), and describe our Coq mechanization (§7). We finish with related and future work (§8 and 9). An archive of the Coq mechanization can be found at Jacobs et al. [2021], and the most recent version at <https://github.com/julesjacobs/cgraphs>.

2 LANGUAGE AND OPERATIONAL SEMANTICS

We present the core of our session-typed λ -calculus with concurrency and asynchronous bidirectional channels (extensions with more features are described in §6). This language is inspired by GV [Wadler 2012; Lindley and Morris 2015], but there are a couple of differences. First, we are more liberal and allow both channel endpoints to be closed anytime, rather than only when a thread terminates. For our proofs this extension poses no problem—it just means that our connectivity graphs might become disconnected. Second, our operational semantics uses a flat thread pool and heap rather than binders and structural congruences, resembling more closely a realistic implementation of message passing. The syntax of expressions of our core language is:

$$e \in \text{Expr} ::= x \mid () \mid n \mid (e, e) \mid \lambda x. e \mid c \mid e e \mid \text{let } x = e \text{ in } e \mid \text{let } () = e \text{ in } e \mid \text{let } (x_1, x_2) = e \text{ in } e \mid \text{if } e \text{ then } e \text{ else } e \mid \text{fork}(e) \mid \text{send}(e, e) \mid \text{receive}(e) \mid \text{close}(e)$$

Pure reduction relation:

$$\begin{aligned}
& (\lambda x. e) v \rightsquigarrow_{\text{pure}} e[v/x] \\
& \text{let } x = v \text{ in } e \rightsquigarrow_{\text{pure}} e[v/x] \\
& \text{let } () = () \text{ in } e \rightsquigarrow_{\text{pure}} e \\
& \text{let } (x_1, x_2) = (v_1, v_2) \text{ in } e \rightsquigarrow_{\text{pure}} e[v_1/x_1][v_2/x_2] \\
& \text{if } n \text{ then } e_1 \text{ else } e_2 \rightsquigarrow_{\text{pure}} e_1 \quad (\text{if } n \neq 0) \\
& \text{if } n \text{ then } e_1 \text{ else } e_2 \rightsquigarrow_{\text{pure}} e_2 \quad (\text{if } n = 0)
\end{aligned}$$

Head reduction relation:

$$\begin{aligned}
& (e_1, h) \rightsquigarrow_{\text{head}} (e_2, h, \epsilon) \quad (\text{if } e_1 \rightsquigarrow_{\text{pure}} e_2) \\
& (\text{fork}(v), h) \rightsquigarrow_{\text{head}} (\#(a, 1), h \uplus \{(a, 0) \mapsto \epsilon, (a, 1) \mapsto \epsilon\}, [v \#(a, 0)]) \\
& \quad (\text{if } (a, 0), (a, 1) \notin \text{dom}(h)) \\
& (\text{send}(c, v), h \uplus \{\bar{c} \mapsto \vec{v}\}) \rightsquigarrow_{\text{head}} (c, h \uplus \{\bar{c} \mapsto \vec{v} \uparrow [v]\}, \epsilon) \\
& (\text{receive}(c), h \uplus \{c \mapsto [v] \uparrow \vec{v}\}) \rightsquigarrow_{\text{head}} ((c, v), h \uplus \{c \mapsto \vec{v}\}, \epsilon) \\
& (\text{close}(c), h \uplus \{c \mapsto \epsilon\}) \rightsquigarrow_{\text{head}} ((), h, \epsilon)
\end{aligned}$$

Global reduction relation:

$$(\vec{e}_a \uparrow [K[e]] \uparrow \vec{e}_b, h) \rightsquigarrow_{\text{global}} (\vec{e}_a' \uparrow [K[e']] \uparrow \vec{e}_b' \uparrow \vec{e}, h') \quad (\text{if } (e, h) \rightsquigarrow_{\text{head}} (e', h', \vec{e}))$$

Evaluation contexts:

$$\begin{aligned}
K \in \text{Ctx} ::= & \square \mid (K, e) \mid (v, K) \mid K e \mid v K \mid \text{let } x = K \text{ in } e \mid \text{let } () = K \text{ in } e \mid \text{let } (x_1, x_2) = K \text{ in } e \mid \\
& \text{if } K \text{ then } e_1 \text{ else } e_2 \mid \text{fork}(K) \mid \text{send}(K, e) \mid \text{send}(v, K) \mid \text{receive}(K) \mid \text{close}(K)
\end{aligned}$$

Fig. 2. The operational semantics of our language.

The literals include the unit value $()$, numbers $n \in \mathbb{N}$, and channel endpoint references $c \in \text{Chan}$ (these enter expressions at run time, see the operational semantics below). As usual in a linearly-typed language, we consider let-binding constructs $\text{let } () = e \text{ in } e$ and $\text{let } (x_1, x_2) = e \text{ in } e$ for pattern matching on the unit value $()$ and pairs (e, e) , respectively.

Operational semantics. We use an asynchronous semantics with two buffers per channel to guarantee that sends in either direction are non-blocking.¹ This is formally modeled as:

$$\begin{aligned}
c \in \text{Chan} ::= & \#(a, t) & h \in \text{Heap} \triangleq \text{Chan} \xrightarrow{\text{fin}} \text{List Val} \\
v \in \text{Val} ::= & () \mid n \mid (v, v) \mid \lambda x. e \mid c & \rho \in \text{Cfg} \triangleq \text{List Expr} \times \text{Heap}
\end{aligned}$$

A heap h is a finite map from channel endpoint references to buffers (modeled as lists of values). Channel endpoint references $\#(a, t)$ consist of an address $a \in \text{Addr}$ and a tag $t \in \{0, 1\}$ denoting the endpoint. The operation $\#(n, t) \triangleq \#(n, 1 - t)$ gives the opposite endpoint. Configurations $(\vec{e}, h)$ consist of a list of expressions $\vec{e}$, modeling the threads, and a heap h that is shared by these threads. The semantics of most constructs is standard, so we focus on the message passing constructs:

¹Due to the session typing discipline, only one of the buffers is expected to be populated at any given time. The two buffers are important to distinguish the origin of the messages, because otherwise an asynchronous send followed by a receive creates a risk that the thread receives back its own message that it just sent.

Fig. 3. Examples of configurations that are deadlocked or have leaked.

`fork( $v$ )` Allocates a new channel with endpoints $c_{\text{left}} \triangleq \#(a, 0)$ and $c_{\text{right}} \triangleq \#(a, 1)$, where a is a fresh address. It starts a new thread running v at c_{left} (v should be a function) and returns c_{right} .
`send( $c, v$ )` Places the message v into the buffer of the opposite endpoint $\bar{c}$ of c and returns c .²

This construct does not block.

`receive( $c$ )` Takes a message v out of the buffer of endpoint c and returns the pair (c, v) . If the buffer is empty, it blocks until a message is available.

`close( $c$ )` Closes the endpoint c and returns the unit value $()$. This construct does not block.

The formal definition of the semantics is given in Figure 2. It involves three reduction relations: (1) pure reductions $e \rightsquigarrow_{\text{pure}} e'$, (2) head-reductions $(e, h) \rightsquigarrow_{\text{head}} (e', h', \vec{e})$, where $\vec{e}$ is a list of spawned threads (empty for all constructs but `fork`), and (3) global reductions $\rho \rightsquigarrow_{\text{global}} \rho'$. Global reductions make use of standard call-by-value evaluation contexts $K \in \text{Ctx}$.

Deadlocks and memory leaks. Untyped programs in our language can deadlock or have memory leaks. A configuration $(\vec{e}, h)$ is *deadlocked* if each expression $e \in \vec{e}$ is a `receive( $c$ )` that is waiting on an empty buffer c in the heap h . A configuration $(\vec{e}, h)$ has *leaked* if each expression $e \in \vec{e}$ is a value, but the heap h is not empty, meaning not all channels have been closed. In Figure 3 we show examples of both. On the left we show the code, and on the right we show a graphical representation of the resulting configuration. As in §1, boxes denote threads (i.e., expressions), circles denote channels (i.e., buffer pairs), black arrows denote channel references, and red triangles denote the waiting dependency. Concretely, a thread with a red triangle pointing to a channel is waiting to receive a message from that channel, and a channel with a red triangle pointing to a thread is waiting for the thread to send a message along that channel.

The simplest form of deadlock is a thread attempting to receive a message from a channel that nobody else has a reference to (first program). If threads violate the usual protocol that one side

²The reason why `send` returns the endpoint c is the session type system, which gives the endpoint a new type, prescribing the remainder of the protocol. The same applies to the `receive` operation.

$$\begin{array}{c}
\frac{\Gamma = \{x \mapsto \tau\}}{\Gamma \vdash x : \tau} \quad \frac{}{\emptyset \vdash () : \mathbf{1}} \quad \frac{n \in \mathbb{N}}{\emptyset \vdash n : \mathbb{N}} \quad \frac{\Gamma_1 \vdash e_1 : \tau_1 \quad \Gamma_2 \vdash e_2 : \tau_2}{\Gamma_1 \uplus \Gamma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2} \quad \frac{\Gamma \uplus \{x \mapsto \tau_1\} \vdash e : \tau_2}{\Gamma \vdash \lambda x. e : \tau_1 \multimap \tau_2} \\[10pt]
\frac{\Gamma_1 \vdash e_1 : \tau_1 \multimap \tau_2 \quad \Gamma_2 \vdash e_2 : \tau_1}{\Gamma_1 \uplus \Gamma_2 \vdash e_1 \ e_2 : \tau_2} \quad \frac{\Gamma_1 \vdash e_1 : \tau_1 \quad \Gamma_2 \uplus \{x \mapsto \tau_1\} \vdash e_2 : \tau_2}{\Gamma_1 \uplus \Gamma_2 \vdash \text{let } x = e_1 \text{ in } e_2 : \tau_2} \\[10pt]
\frac{\Gamma_1 \vdash e_1 : \mathbf{1} \quad \Gamma_2 \vdash e_2 : \tau}{\Gamma_1 \uplus \Gamma_2 \vdash \text{let } () = e_1 \text{ in } e_2 : \tau} \quad \frac{\Gamma_1 \vdash e_1 : \tau_1 \times \tau_2 \quad \Gamma_2 \uplus \{x_1 \mapsto \tau_1\} \uplus \{x_2 \mapsto \tau_2\} \vdash e_2 : \tau}{\Gamma_1 \uplus \Gamma_2 \vdash \text{let } (x_1, x_2) = e_1 \text{ in } e_2 : \tau} \\[10pt]
\frac{\Gamma_1 \vdash e_1 : \mathbb{N} \quad \Gamma_2 \vdash e_2 : \tau \quad \Gamma_3 \vdash e_3 : \tau}{\Gamma_1 \uplus \Gamma_2 \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau} \quad \frac{\Gamma \vdash e : \bar{s} \multimap \mathbf{1}}{\Gamma \vdash \text{fork}(e) : s} \quad \frac{\Gamma_1 \vdash e_1 : !\tau. s \quad \Gamma_2 \vdash e_2 : \tau}{\Gamma_1 \uplus \Gamma_2 \vdash \text{send}(e_1, e_2) : s} \\[10pt]
\frac{\Gamma \vdash e : ?\tau. s}{\Gamma \vdash \text{receive}(e) : s \times \tau} \quad \frac{\Gamma \vdash e : \text{End}}{\Gamma \vdash \text{close}(e) : \mathbf{1}}
\end{array}$$

Fig. 4. The static type system of our language.

receives and the other side sends a message, then a deadlock can occur if both try to receive (second program). A deadlock can occur even if all parties are locally well behaved, but cause a cyclic dependency (third program). Note that even though the reference structure (black arrows) of this example forms a directed acyclic graph, a deadlock occurs because the waiting direction (red triangles) can be opposite of the reference direction (black arrows). Finally, memory leaks can occur if channels are not properly closed (fourth program).

Session Typing. A linear type system with session types can be used to rule out deadlocks:³

$$\begin{aligned}
\tau \in \text{Type} &::= \mathbf{1} \mid \mathbb{N} \mid \tau \times \tau \mid \tau \multimap \tau \mid s \\
s \in \text{Session} &::= \text{End} \mid ?\tau. s \mid !\tau. s
\end{aligned}$$

A session type s denotes a sequence of actions, with $?\tau$ indicating a receive, $!\tau$ a send, and End termination, where τ denotes the type of the message. The *dual* $\bar{s}$ of a session type s is defined by flipping all sends (!) and receives (?):

$$\overline{\text{End}} \triangleq \text{End} \quad \overline{!\tau. s} \triangleq ?\tau. \bar{s} \quad \overline{?\tau. s} \triangleq !\tau. \bar{s}$$

The rules of the type system are shown in Figure 4. Note that the type system is *higher-order* because it allows sending any value over a channel, including functions and channel endpoints.

Session types ensure deadlock and leak freedom by combining channel and thread creation through the fork construct.⁴ Together with linear channel typing, this ensures that the resulting reference structure is *acyclic*, even when viewed as an *undirected graph*, in which edges may be traversed in either direction. Let us consider the deadlocked programs in Figure 3. The first one is ruled out by ensuring that there always exists a counter party (due to the absence of weakening). The second one is ruled out by ensuring that all threads adhere to protocols (due to session duality). The third one is ruled out by ensuring that the reference structure is acyclic (due to the absence of contraction). The memory leak in the last example is ruled out by a combination of these rules.

³For simplicity we let even integers be linear; in §6 we extend the type system with support for unrestricted types.

⁴If we had a `new_chan : 1  $\multimap$  (s  $\times$   $\bar{s}$ )` construct, then `let (c1, c2) = new_chan () in let x = receive(c1) in ...` would deadlock.

3 KEY IDEAS

Before we detail the abstractions that make our proof method parametric (§4 and §5), we describe a concrete instantiation of our method to our session-typed language (§2). To do so, we first discuss the well-known method of progress and preservation and the challenges in applying it to prove deadlock and resource leak freedom (§3.1). To address these challenges, we introduce connectivity graphs (§3.2) and describe how we use separation logic to formalize run-time typing judgments for individual expressions (§3.3) and a well-formedness predicate for configurations (§3.4). We finally show how to use our proof method to prove preservation (§3.5) and progress (§3.6).

3.1 Generalizing The Progress and Preservation Method

Traditionally, a language is said to be *type safe* if well typed programs do not get stuck. For purely functional languages, like the Simply Typed Lambda Calculus (STLC), this is stated as:

Theorem 3.1 (Type safety). *If $\emptyset \vdash e : \tau$ and $e \rightsquigarrow^* e'$, then either e' is a value, or e' can step further (i.e., $\exists e''. e' \rightsquigarrow e''$).*

Type safety is often proved using the method of *progress and preservation* [Wright and Felleisen 1994; Harper 2016; Pierce 2002], which decomposes type safety into two properties that imply it:

- **Preservation:** If $\emptyset \vdash e : \tau$ and $e \rightsquigarrow e'$, then $\emptyset \vdash e' : \tau$.
- **Progress:** If $\emptyset \vdash e : \tau$, then either e is a value, or e can step further (i.e., $\exists e''. e \rightsquigarrow e''$).

For pure languages like STLC, the proofs of these properties are straightforward: both properties are proved by induction on the structure of the typing judgment and/or the reduction relation.

For languages with mutable state or concurrency, the above properties must be generalized to account for a program's run-time configurations. In general, for a language with expressions *Expr* we have a set $\rho \in \text{Cfg}$ of configurations, an initial configuration $\text{init} \in \text{Expr} \rightarrow \text{Cfg}$, and a predicate $\text{final} \in \text{Cfg} \rightarrow \text{Prop}$ of configurations that are considered to be safely terminated.

Theorem 3.2 (Generalized type safety). *If $\emptyset \vdash e : ()$ and $\text{init}(e) \rightsquigarrow^* \rho$, then either $\text{final}(\rho)$ or ρ can step further (i.e., $\exists \rho'. \rho \rightsquigarrow \rho'$).*

Recall that for our session-typed language we have $\text{Cfg} \triangleq \text{List Expr} \times \text{Heap}$. We let final select configurations where all threads have terminated with a unit value, and the heap is empty:

$$\text{init}(e) \triangleq ([e], \emptyset) \qquad \text{final}(\vec{e}, h) \triangleq h = \emptyset \wedge \forall i. e_i = ()$$

By defining final this way, the type safety theorem expresses *deadlock and memory leak freedom*.⁵ To see why, consider a configuration $\rho = (\vec{e}, h)$ that does not satisfy $\text{final}(\rho)$ and cannot step any further. It must either consist of threads $\vec{e}$ that have not terminated but cannot step (indicating a deadlock), or of terminated threads $\vec{e}$ but a non-empty heap h (indicating a memory leak).

For deadlock and resource leak freedom, we need to restrict expressions to have a ground type. For example, an expression like $\text{fork}(\lambda x. \text{close}(x)) : \text{End}$ exhibits a trivial memory leak because the channel endpoint returned by $\text{fork}()$ is still active. For simplicity, we use the unit type $()$ in Theorem 3.2, but of course, other ground types like $\mathbf{N}$ would suffice too.

The method of progress and preservation can be generalized to prove our generalized type safety theorem (Theorem 3.2) by choosing a well-formedness predicate $\text{wf} \in \text{Cfg} \rightarrow \text{Prop}$ that satisfies:

- **Initialization:** If $\emptyset \vdash e : ()$, then $\text{wf}(\text{init}(e))$.
- **Generalized preservation:** If $\text{wf}(\rho)$ and $\rho \rightsquigarrow \rho'$, then $\text{wf}(\rho')$.
- **Generalized progress:** If $\text{wf}(\rho)$, then either $\text{final}(\rho)$ or ρ can step further (i.e., $\exists \rho'. \rho \rightsquigarrow \rho'$).

⁵This kind of deadlock freedom is also known as *global progress* in the session type literature. In §6.3 we prove a stronger property that also rules out partial deadlocks.

Fig. 5. An example of a connectivity graph. Brown boxes depict threads, and blue circles depict channels.

The primary challenge is to define a well-formedness predicate wf in such a way that these properties can be proved. A naive definition of $wf(\vec{e}, h)$ would simply demand each expression in the thread pool $\vec{e}$ and each buffer in the heap h to be well typed. Unfortunately, this naive definition does not quite work:

- (1) Channel references $\#(a, t)$ enter the expressions $\vec{e}$ and heap h throughout the execution of the program. The typing judgment $\Gamma \vdash e : \tau$ of our type system (Figure 4) does not account for channel references $\#(a, t)$ because they cannot be written in source programs.
- (2) Whenever a channel reference $\#(a, t)$ appears in a thread or channel buffer, the type of $\#(a, t)$ should match up with the values in the buffers at address a in the heap h , and with the type $\#(a, 1 - t)$ of the other endpoint.

For the simpler case of proving type safety for the STLC with references, Harper [2016] and Pierce [2002] remedy issue (1) by introducing a *run-time typing judgment* $\Gamma; \Sigma \vdash e : \tau$. This judgment extends the static typing judgment with a *heap typing* Σ , which assigns types to heap addresses. Issue (2) is addressed because the heap typing makes sure that the typing of each reference is consistent with the corresponding value in the heap.

Unfortunately, adapting this approach to prove deadlock and resource freedom is not as simple. Conventional heap typings only capture the typing of addresses, not the acyclicity of the reference topology. The latter is crucial to prove “generalized progress”, which states that the well-formedness predicate wf indeed implies deadlock and resource leak freedom.

3.2 Generalizing Heap Typings to Connectivity Graphs

Our notion of *connectivity graphs* generalizes the notion of heap typings by simultaneously keeping track of the types of channels in the heap, and providing an abstract representation of the reference topology. In their full generality, connectivity graphs are represented as finite maps from pairs of vertices V to the labels L on the edges between them:

$$G \in Cgraph(V, L) \triangleq \{G \in V \times V \xrightarrow{\text{fin}} L \mid G \text{ has no undirected cycles}\}$$

To reason about our language defined in §2, we instantiate the vertices V and edge labels L of a connectivity graph $Cgraph(V, L)$ as follows:

$$v \in V ::= \text{Thread}(i) \mid \text{Chan}(a) \qquad l \in L \triangleq \{0, 1\} \times \text{Session}$$

The vertices V are threads $\text{Thread}(i)$ (with position i in the thread pool) or channels $\text{Chan}(a)$ (with address a in the heap). The edges are channel references and have a label $(t, s) \in L$ that consists of a tag $t \in \{0, 1\}$, indicating the channel endpoint pointed to, and a session type $s \in \text{Session}$, indicating the endpoint’s session type. An example of a connectivity graph is depicted in Figure 5. Note that

the red triangles that we used to depict the waiting directions in Figures 1 and 3 in §1 are not part of the connectivity graph itself, because these can be derived from the run-time configuration. The session types on the edges *are* part of the connectivity graph, because they cannot be derived from the run-time configuration. In Definition 3.2 we formalize how the red triangles are derived.

Connectivity graphs corresponding to configurations in our language have a number of important properties. First, vertices $\text{Thread}(i)$ can have an arbitrary number of outgoing edges, but have no incoming edges. That is because threads can own channel endpoints, but threads can never be owned. Second, vertices $\text{Chan}(a)$ can also have an arbitrary number of outgoing edges, but at most two incoming edges. Outgoing edges are due to higher-order channels—a channel c_1 can be sent over another channel c_2 , resulting in an edge from c_2 to c_1 that models that c_2 owns c_1 . Ingoing edges correspond to a channel's endpoints, which can be at most two. The number of incoming edges is one in case one channel endpoint has been deallocated, but the other is still in use.

A key ingredient of connectivity graphs is the acyclicity restriction. They should be *acyclic* in the *undirected* sense: there must be no cycles even if we disregard the direction of the edges. In other words, a connectivity graph must be an unrooted undirected forest if we erase the direction of the edges. The third example in Figure 3 shows why acyclicity in the undirected sense is important.

To formally reason about ownership, we introduce the following functions:

$$\text{out}(G, v) \in V \xrightarrow{\text{fin}} L \qquad \text{in}(G, v) \in \text{Multiset } L$$

The outgoing edges $\text{out}(G, v)$ determine which *resources vertex v owns*, whereas the incoming edges $\text{in}(G, v)$ determine at which labels (*i.e.*, types) the *vertex v is owned*. We use the above functions in the definitions of the run-time typing judgment (§3.3) and the configuration well-formedness predicate (§3.4). We represent the outgoing edges $\text{out}(G, v)$ of a vertex v as a finite map from vertices to labels to track *which* resources a vertex v owns and at *which type*. The incoming edges $\text{in}(G, v)$ of a vertex v , however, we represent as a multiset of labels because it only matters at *which type* a vertex v is owned, but not by whom (note that only channel endpoints can be owned).

What is parametric in this section. Connectivity graphs $\text{Cgraph}(V, L)$ are parametric over the type of vertices V and labels L . All theory about connectivity graphs (including the separation logic) that we present throughout the rest of this section is parametric. Connectivity graphs and their theory are thus modularly separated from the operational semantics and type system of the language, which we found to be essential for keeping the complexity of the mechanization manageable.

3.3 Run-Time Typing Judgment Using Separation Logic

In the previous section, we developed the notion of a connectivity graph as a generalization of the heap typing, known from type safety proofs of the STLC with references [Harper 2016; Pierce 2002]. We now make this generalization precise, develop a run-time typing judgment for our language, and show how we can use separation logic to hide reasoning about linearity.

We start with the run-time judgment $\Gamma; \Sigma \vdash e : \tau$, where $\Sigma \in V \xrightarrow{\text{fin}} L$ provides the session types of the free channel references in e . Channel references amount to edges in our connectivity graph, and thus Σ becomes the set of *outgoing edges* $\text{out}(G, v)$ associated with the threads and channels v occurring in e . Let us consider the typing rule for channel references and function application:

$$\frac{}{\emptyset; \{\text{Chan}(a) \mapsto (t, s)\} \vdash \#(a, t) : s} \qquad \frac{\Gamma_1; \Sigma_1 \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma_2; \Sigma_2 \vdash e_2 : \tau_1}{\Gamma_1 \uplus \Gamma_2; \Sigma_1 \uplus \Sigma_2 \vdash e_1 e_2 : \tau_2}$$

Because our language is linear, we insist that the Σ -context is a singleton in the rule for channel references. In the rule for application, both contexts are split into disjoint parts for the subexpressions. Unfortunately, this leads to a multiplication of contexts and disjointness side conditions,

$$\begin{aligned}
(\text{Emp})(\Sigma) &\triangleq (\Sigma = \emptyset) & (P, Q \in oProp &\triangleq (V \xrightarrow{\text{fin}} L) \rightarrow Prop) \\
(\ulcorner \phi \urcorner)(\Sigma) &\triangleq \phi \wedge (\Sigma = \emptyset) & (\Sigma \in V \xrightarrow{\text{fin}} L) \\
(P \wedge Q)(\Sigma) &\triangleq P(\Sigma) \wedge Q(\Sigma) \\
(\exists x. P(x))(\Sigma) &\triangleq \exists x. P(x)(\Sigma) \\
(\text{own}(\Sigma'))(\Sigma) &\triangleq (\Sigma = \Sigma') \\
(P * Q)(\Sigma) &\triangleq \exists \Sigma_1 \Sigma_2. \text{dom}(\Sigma_1) \cap \text{dom}(\Sigma_2) = \emptyset \wedge \Sigma = \Sigma_1 \uplus \Sigma_2 \wedge P(\Sigma_1) \wedge Q(\Sigma_2) \\
(P \multimap Q)(\Sigma) &\triangleq \forall \Sigma'. (\text{dom}(\Sigma) \cap \text{dom}(\Sigma') = \emptyset \wedge P(\Sigma')) \Rightarrow Q(\Sigma \uplus \Sigma')
\end{aligned}$$

Fig. 6. Semantics of separation logic.

because of the disjoint unions in the conclusion. These side conditions cannot be ignored because we want to mechanize our results. This is not a big issue for the variable context Γ since we mostly deal with *closed* expressions (*i.e.*, with $\Gamma = \emptyset$) because the operational semantics operates on closed expressions. The Σ -context, however, is in general non-empty for run-time expressions.

We use separation logic [O'Hearn et al. 2001] to hide the Σ -context and its disjointness conditions. We work with separation logic propositions $oProp \triangleq (V \xrightarrow{\text{fin}} L) \rightarrow Prop$, which are predicates over a context $\Sigma \in V \xrightarrow{\text{fin}} L$ of outgoing edges. Our use of separation logic as an *internal*, meta theoretical tool is inspired by Rouvoet et al. [2020] and contrasts with traditional uses which employ separation logic as an *external*, user-visible tool when specifying programs in Hoare-style logics. The separation logic connectives are defined in Figure 6. To assert that a separation logic proposition P is true, is to assert that $P(\emptyset)$ is true. An important special case is that $P \multimap Q$ is true, if $\forall \Sigma. P(\Sigma) \Rightarrow Q(\Sigma)$.

Instead of the ordinary typing judgment $(\Gamma; \Sigma \vdash e : \tau) \in Prop$ we define a separation-logic based judgment $(\Gamma \vdash e : \tau) \in oProp$, so that $(\Gamma; \Sigma \vdash e : \tau)$ iff $(\Gamma \vdash e : \tau)(\Sigma)$. The preceding two rules are then reformulated as follows:

$$\frac{\text{own}(\text{Chan}(a) \mapsto (t, s))}{\emptyset \vdash \#(a, t) : s} * \quad \frac{\Gamma_1 \vdash e_1 : \tau_1 \multimap \tau_2 \quad \Gamma_2 \vdash e_2 : \tau_1}{\Gamma_1 \uplus \Gamma_2 \vdash e_1 e_2 : \tau_2} *$$

The Σ -contexts are hidden by the separation logic connectives, and the disjointness conditions on Σ are taken care of by the separating conjunction ($*$). At a channel reference, we use the $\text{own}(\Sigma)$ connective, which asserts that the separation logic resource is precisely Σ .

An exception to the rule that contexts are split up disjointly (with $*$) is the if e_1 then e_2 else e_3 expression. Although the channel references occurring in e_1 must be disjoint from those occurring in e_2 and e_3 , the same endpoint is allowed to occur in both e_2 and e_3 , because only one of the branches will be executed. This pattern too fits neatly in the separation logic methodology; we use separating conjunction ($*$) between e_1 and e_2, e_3 , but ordinary conjunction ($\wedge$) between e_2 and e_3 :

$$\frac{\Gamma_1 \vdash e_1 : N \quad (\Gamma_2 \vdash e_2 : \tau \quad \wedge \quad \Gamma_2 \vdash e_3 : \tau)}{\Gamma_1 \uplus \Gamma_2 \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau} *$$

Figure 7 contains the full definition of our run-time type system using separation logic. Although it is possible to define the meaning of general inductive separation logic inference rules via Tarski's fixed point theorem, this generality is not necessary here: the expressions in the premises of each rule are strictly smaller than the expression in the conclusion, so the rules can be interpreted as being defined by recursion on the expression. We use this approach in the Coq formalization, because it has the additional benefit that we get the inversion rules for free.

$$\begin{array}{c}
\frac{\ulcorner \Gamma = \{x \mapsto \tau\} \urcorner}{\Gamma \vdash x : \tau}^* \quad \frac{\text{Emp}}{\emptyset \vdash () : \mathbf{1}}^* \quad \frac{\ulcorner n \in \mathbb{N} \urcorner}{\emptyset \vdash n : \mathbb{N}}^* \quad \frac{\Gamma_1 \vdash e_1 : \tau_1 \quad * \quad \Gamma_2 \vdash e_2 : \tau_2}{\Gamma_1 \uplus \Gamma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2}^* \\
\\
\frac{\Gamma_1 \vdash e_1 : \tau_1 \multimap \tau_2 \quad * \quad \Gamma_2 \vdash e_2 : \tau_1}{\Gamma_1 \uplus \Gamma_2 \vdash e_1 e_2 : \tau_2}^* \quad \frac{\Gamma \uplus \{x \mapsto \tau_1\} \vdash e : \tau_2}{\Gamma \vdash \lambda x. e : \tau_1 \multimap \tau_2}^* \\
\\
\frac{\Gamma_1 \vdash e_1 : \tau_1 \quad * \quad \Gamma_2 \uplus \{x \mapsto \tau_1\} \vdash e_2 : \tau_2}{\Gamma_1 \uplus \Gamma_2 \vdash \text{let } x = e_1 \text{ in } e_2 : \tau_2}^* \quad \frac{\Gamma_1 \vdash e_1 : \mathbf{1} \quad * \quad \Gamma_2 \vdash e_2 : \tau}{\Gamma_1 \uplus \Gamma_2 \vdash \text{let } () = e_1 \text{ in } e_2 : \tau}^* \\
\\
\frac{\Gamma_1 \vdash e_1 : \tau_1 \times \tau_2 \quad * \quad \Gamma_2 \uplus \{x_1 \mapsto \tau_1\} \uplus \{x_2 \mapsto \tau_2\} \vdash e_2 : \tau}{\Gamma_1 \uplus \Gamma_2 \vdash \text{let } (x_1, x_2) = e_1 \text{ in } e_2 : \tau}^* \\
\\
\frac{\Gamma_1 \vdash e_1 : \mathbb{N} \quad * \quad (\Gamma_2 \vdash e_2 : \tau \quad \wedge \quad \Gamma_2 \vdash e_3 : \tau)}{\Gamma_1 \uplus \Gamma_2 \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau}^* \quad \frac{\Gamma_1 \vdash e_1 : !\tau.s \quad * \quad \Gamma_2 \vdash e_2 : \tau}{\Gamma_1 \uplus \Gamma_2 \vdash \text{send}(e_1, e_2) : s}^* \\
\\
\frac{\Gamma \vdash e : ?\tau.s}{\Gamma \vdash \text{receive}(e) : s \times \tau}^* \quad \frac{\Gamma \vdash e : \bar{s} \multimap \mathbf{1}}{\Gamma \vdash \text{fork}(e) : s}^* \quad \frac{\Gamma \vdash e : \text{End}}{\Gamma \vdash \text{close}(e) : \mathbf{1}}^* \quad \frac{\text{own}(\text{Chan}(a) \mapsto (t, s))}{\emptyset \vdash \#(a, t) : s}^*
\end{array}$$

Fig. 7. The separation-logic based run-time type system of our language.

A key strength of separation logic is that we can prove assertions using the proof rules of the logic of Bunched Implications (BI) [O’Hearn and Pym 1999]. For example, separating conjunction ($*$) is associative and commutative, separating conjunction ($*$) has Emp as identity element, and magic wand ($\multimap$) is the adjoint of separating conjunction ($*$). We use the Iris Proof Mode [Krebbers et al. 2017, 2018] to reason abstractly using the rules of separation logic in Coq (see §7 for details).

What is parametric in this section. The definition of the separation logic connectives and the proof rules for the separation logic are parametric in the types of vertices V and labels L .

3.4 Well-Formedness of Configurations Using Connectivity Graphs

Now that we have a run-time typing judgment for a single expression, we are in a position to define which configurations are well-formed. Recall that a configuration is a pair $(\vec{e}, h)$ where $\vec{e} : \text{List Expr}$ is the thread pool and where $h : \text{Chan} \xrightarrow{\text{fin}} \text{ListVal}$ is the heap of channel buffers. We must certainly insist that all threads $\vec{e}$ are well-typed expressions (of unit type), and that all the values inside the heap h are well-typed. For the latter we have to ensure that a channel’s endpoints are of dual types, modulo the messages queued up in the buffer. This requires us to consider the *incoming edges* $\text{in}(G, v)$ of a vertex v in addition to its *outgoing edges* $\text{out}(G, v)$. We can thus state well-formedness of a configuration in terms of its connectivity graph.

A configuration is well-formed if there exists a connectivity graph such that each thread and channel is locally well-formed with respect to its vertex in the graph:

$$\text{wf}(\vec{e}, h) \triangleq \exists G : \text{Cgraph}(V, L). \forall v \in V. \text{wf}_{(\vec{e}, h)}^{\text{local}}(v, \text{in}(G, v))(\text{out}(G, v))$$

Here, $\text{wf}_{(\vec{e}, h)}^{\text{local}} : V \times \text{Multiset } L \rightarrow \text{oProp}$ gives the *local well-formedness condition* for each vertex. It has two explicit arguments (the vertex $v \in V$ and its incoming edges $\text{in}(G, v) \in \text{Multiset } L$), and an

extra argument $\text{out}(G, v) \in V \xrightarrow{\text{fin}} L$ that will form the vertex' local Σ -context, which is implicit in the type signature of wf^{local} because $\text{oProp} \triangleq (V \xrightarrow{\text{fin}} L) \rightarrow \text{Prop}$.

A thread is locally well-formed if it is well-typed (with the implicit Σ -context given by its outgoing edges), and has no incoming edges (because threads cannot be owned):

$$\text{wf}_{(\vec{e}, h)}^{\text{local}}(\text{Thread}(i), \Delta) \triangleq \begin{cases} \ulcorner \Delta = \emptyset^\top * \emptyset \vdash e_i : 1 \urcorner & \text{if } i < |\vec{e}| \\ \ulcorner \Delta = \emptyset^\top \urcorner & \text{otherwise} \end{cases}$$

Note that wf quantifies over any vertex $v \in V$, and we thus have to consider any thread index i , including those $i \geq |\vec{e}|$ that are not yet in use. For such indexes, we use the separation logic proposition $\ulcorner \Delta = \emptyset^\top \urcorner$ to assert that both the incoming and outgoing edges are empty. The latter is implicit by the semantics of $\ulcorner \Delta = \emptyset^\top \urcorner$ (see Figure 6).

A channel is locally well-formed if the buffers are well-typed (with the implicit Σ -context given by its outgoing edges), and have matching incoming edges match the types of the endpoints:

$$\text{wf}_{(\vec{e}, h)}^{\text{local}}(\text{Chan}(a), \Delta) \triangleq \begin{cases} \exists s_0, s_1, s. \ulcorner \Delta = \{(0, s_0), (1, s_1)\}^\top * \urcorner * & \text{if } \#(a, 0), \#(a, 1) \in \text{dom}(h) \\ \quad \vdash_{\text{buf}} h(a, 0) : (s_0, s) * \\ \quad \vdash_{\text{buf}} h(a, 1) : (s_1, \bar{s}) \\ \exists b, s. \ulcorner \Delta = \{(t, s)\}^\top * \urcorner & \text{if } \#(a, t) \in \text{dom}(h) \\ \quad \vdash_{\text{buf}} h(a, t) : (s, \text{End}) & \text{and } \#(a, 1 - t) \notin \text{dom}(h) \\ \ulcorner \Delta = \emptyset^\top \urcorner & \text{otherwise} \end{cases}$$

In this definition we have to consider three cases. The first case corresponds to the situation in which both buffers are still in use. In that case, there must be two incoming edges in the connectivity graph, labeled with session types that are dual modulo the values in the buffers. For instance, if the left endpoint has session type $? \tau_1. ? \tau_2. s$ and the right endpoint has session type $\bar{s}$, then the buffer of the left endpoint must be $[v_1, v_2]$ with $\vdash v_1 : \tau_1$ and $\vdash v_2 : \tau_2$. The second case corresponds to the situation in which one buffer has been deallocated. The third case corresponds to the situation in which the channel is not allocated (or both buffers have been deallocated).

The *buffer typing judgment* $\vdash_{\text{buf}} \vec{v} : (s_1, s_2)$ is inductively defined by the following rules:

$$\frac{\text{Emp}}{\vdash_{\text{buf}} \epsilon : (s, s)} * \quad \frac{\emptyset \vdash v : \tau \quad * \quad \vdash_{\text{buf}} \vec{v} : (s_1, s_2)}{\vdash_{\text{buf}} ([v] ++ \vec{v}) : (? \tau. s_1, s_2)} *$$

These rules express that $\vdash_{\text{buf}} \vec{v} : (s_1, s_2)$ holds if s_1 is equal to prefixing s_2 with the types of the values $\vec{v}$ in the buffer. Note that similar to the other run-time judgments, the buffer typing judgment is defined in separation logic, which implicitly ensures that the Σ -environment is distributed disjointly over the values in the buffer.

What is parametric in this section. The definition of wf is parametric in the type of vertices V and labels L , but also a *local well-formedness predicate* wf^{local} that captures the language-specific information by linking the incoming and outgoing edges of each vertex to their run-time counterpart.

3.5 Proving Preservation Using Local Connectivity Graph Transformations

Now that we have defined the well-formedness predicate $\text{wf}(\vec{e}, h)$, we must prove that it is preserved by the operational semantics: if $(\vec{e}, h) \rightsquigarrow_{\text{global}} (\vec{e}', h')$, then $\text{wf}(\vec{e}, h)$ implies $\text{wf}(\vec{e}', h')$. Recall that the well-formedness predicate $\text{wf}(\vec{e}, h)$ intuitively means “there exists a connectivity graph G describing the configuration $(\vec{e}, h)$ ”, so when the configuration steps to a new configuration $(\vec{e}', h')$, we must show that there exists a new connectivity graph G' that describes $(\vec{e}', h')$.

Fig. 8. The operational steps and the corresponding connectivity graph transformations.

If the head step is a pure step, then the new connectivity graph is exactly the same as the old one, and the preservation of the well-formedness predicate follows by a standard case analysis of the possible pure steps, because the heap does not change and no new threads are spawned.

Operational steps that involve channel operations are the interesting cases because they may alter the connectivity graph. Figure 8 provides a schematic overview. We focus on the $\text{send}(c, v)$ operation, which moves the value v from the thread into the buffer of channel c . The session type in the label on the edge corresponding to c itself must change from $!\tau.s$ to s . Additionally, if the value v contains channel references, the connectivity graph must change to reflect this. The changes to the connectivity graph for send and the other channel operations are depicted in Figure 8.

Once we have chosen the appropriate new connectivity graph G' , we have to prove that this graph indeed describes the new configuration $(\vec{e}', h')$. This amounts to showing that the local well-formedness condition $\text{wf}_{(\vec{e}', h')}^{\text{local}}(v, \text{in}(G', v))(\text{out}(G', v))$ is re-established for every vertex v .

For $\text{send}(c, v)$ the relevant parts of the $(\vec{e}, h)$ -configuration and $(\vec{e}', h')$ -configuration are:

$$\begin{aligned} e_i &= K[\text{send}(c, v)] & h(\vec{c}) &= \vec{v} \\ e'_i &= K[c] & h'(\vec{c}) &= \vec{v} \uparrow [v] \end{aligned}$$

The thread pool and heap do not change at other locations. After this change to the configuration and the corresponding change to the connectivity graph (as depicted in Figure 8), we classify the vertices into three types and explain how the local well-formedness $\text{wf}_{(\vec{e}', h')}^{\text{local}}$ is restored.

- (1) For the vertices v' where neither the corresponding part of the configuration nor the incoming and outgoing edges change, $\text{wf}_{(\vec{e}, h)}^{\text{local}}(v', \text{in}(G, v'))(\text{out}(G, v'))$ remains valid.
- (2) For the vertices v' that correspond to channels referenced *inside the message* v , the owner changes from Thread(i) to Chan($c.1$) (corresponding to T1 and C in the figure). These vertices are not affected either, because $\text{in}(G, v') = \text{in}(G', v')$. Since $\text{in}(G, v')$ and $\text{in}(G', v')$ are multisets of *labels*, they are thus not affected by the change of owner.

- (3) The vertices $v_1 = \text{Thread}(i)$ and $v_2 = \text{Chan}(c.1)$ (corresponding to T1 and C in the figure) are the vertices that are truly affected. Re-establishing their $\text{wf}_{(\vec{e}, h')}^{\text{local}}(v_{12}, \text{in}(G', v_{12}))(\text{out}(G', v_{12}))$ requires some language-specific reasoning, because both their part of the configuration and their incoming and outgoing edges change.

There is another proof obligation that we need to meet: the connectivity graph has to remain acyclic when we do these local transformations.

Even though Figure 8 looks hopelessly language specific, we show that we can use our separation logic to distill abstract rules for *local graph transformations* (§5). These rules involve the transfer of resources between the old local well-formedness predicates $\text{wf}_{(\vec{e}, h)}^{\text{local}}(v, \text{in}(G, v))$ and the new local well-formedness predicates $\text{wf}_{(\vec{e}', h')}^{\text{local}}(v, \text{in}(G', v))$ (for the affected vertices v in question). To distill these rules, it is crucial that the local well-formedness predicate is a separation logic proposition, which enables reasoning using the abstract proof rules of separation logic, without explicitly having to reference the graph, nor having to explicitly establish acyclicity, nor having to deal with disjointness conditions. The reasoning left to the user of the rule is purely local and precisely the language-specific reasoning that cannot be done generically. The result is that the preservation proof appears to perform no graph reasoning at all: at no point in the preservation proof is there any value of type $G, G' : \text{Cgraph}(V, L)$ in the proof context.

What is parametric in this section. The separation-logic based rules for *local graph transformations* (§5) are parametric in the type of vertices V and labels L , and the local well-formedness predicate.

3.6 Proving Progress Using Waiting Induction

To prove progress, we have to show that if $\text{wf}(\vec{e}, h)$ holds, then either $\text{final}(\vec{e}, h)$ holds (i.e., $e_i = ()$ for all i and $h = \emptyset$), or the configuration can step. This is equivalent to saying that:

$$\text{wf}(\vec{e}, h) \text{ and } \text{active}(\vec{e}, h) \neq \emptyset \text{ implies that } (\vec{e}, h) \text{ can step}$$

Here, $\text{active}(\vec{e}, h)$ is the set of threads and channels that have not yet terminated and not yet been fully deallocated, respectively.

Definition 3.1 (Active). The set of *active vertices* in configuration $(\vec{e}, h)$ is formally defined as

$$\text{active}(\vec{e}, h) \triangleq \{\text{Thread}(i) \mid e_i \neq ()\} \cup \{\text{Chan}(a) \mid h(a, 0) \neq \perp \vee h(a, 1) \neq \perp\}$$

If $\text{active}(\vec{e}, h) \neq \emptyset$, then there exists a vertex $v \in \text{active}(\vec{e}, h)$ for which we must find a thread that can step. If the vertex v is a thread that can step, we are done. The difficulty is that v may be a thread that is blocked on a $\text{receive}(c)$, where the corresponding buffer of c in heap h is empty. If the configuration is well-formed, then we will presumably be able to find a non-blocked thread connected to the other endpoint of c , since that side will eventually be responsible for sending a message to c . However, the thread holding the other endpoint of c may be blocked itself, waiting on a receive on a different channel. Also, the endpoint c may not even be held by another thread; it could be stored in the buffer of some other channel c' .

Our way out is to use the connectivity graph: starting from vertex v , we search for another thread that can step. To organize this search process, we annotate edges of the connectivity graph with a *waiting direction* (as also done in §1), depicted as red triangles in Figure 9. The waiting direction is formalized using the notion of being blocked.

Definition 3.2 (Blocked). A vertex v_1 is *blocked* on vertex v_2 in configuration $(\vec{e}, h)$ if v_1 is a thread that is trying to receive from channel v_2 whose buffer is empty. Formally:

$$\text{blocked}_{(\vec{e}, h)}(v_1, v_2) \triangleq \exists i, a, t, K. v_1 = \text{Thread}(i) \wedge v_2 = \text{Chan}(a) \wedge e_i = K[\text{receive}(\#(a, t))] \wedge h(a, t) = \epsilon$$

Fig. 9. The connectivity graph from Figure 5 annotated with red triangles for the waiting direction.

The waiting direction (red triangle) $v_1 \rightarrow_G^l v_2$ of an edge coincides with its ownership direction (black arrowhead) if v_1 is blocked on v_2 . Otherwise, it is opposite to the ownership direction.

To find a thread that can step from $v \in \text{active}(\vec{e}, h)$, we follow edges in the waiting direction until we arrive at a vertex that has no outgoing waiting edges. As one can see in Figure 9, if we follow the waiting direction (red triangles) from *any* start vertex v , we always end up in a thread that can step (green dotted square). To prove that we can always find a thread that can step by simply following the waiting edges from any starting vertex, we have to show that:

- (1) If the current vertex v is a thread, it can either step, or it has an outgoing waiting edge.
- (2) If the current vertex v is a channel, it always has an outgoing waiting edge.
- (3) The search process terminates, because the graph is acyclic as an undirected graph.

To show (1): We show that active threads v can step or have an outgoing waiting arrow by induction on typing. The interesting cases are the channel operations, and receive in particular, so suppose that the thread's expression is $K[\text{receive}(v)]$. By run-time typing, we know that v is a channel reference $\#(a, t)$, and the typing rule for $\text{receive}(v)$ gives us the separation logic resource $\text{own}(\{ \text{Chan}(a) \mapsto (t, ?\tau.s) \})$. From this it follows that the thread has an outgoing edge to $\text{Chan}(a)$, and hence $\text{Chan}(a)$ has an incoming edge with the label $(t, ?\tau.s)$. From the channel's local well-formedness predicate, it follows that the required buffer exists in the heap. If the buffer is non-empty, then the receive can proceed, so the configuration can step. If the buffer is empty, we have an outgoing waiting arrow from the thread to the channel, so the search process can continue.

To show (2): The channel v is active, so it has a buffer, so it must have a corresponding incoming edge in the graph by the definition of the local well-formedness predicate for channels. If that incoming edge comes from a vertex v' , and that vertex v' is not blocked on v , we are done. That is because then the waiting direction is pointing from v to v' , and we can continue the search process from v' . If v' is a thread currently blocked on v , then the session type on that edge must be a receive. It follows from the channel's local well-formedness predicate that the other endpoint has not yet been closed, and thus there is another incoming edge. It cannot be the case that both buffers are empty and the other edge is also a receive, because that would violate duality. Thus, the other edge is coming from a vertex v'' that is *not* a thread currently blocked on us. So there is a waiting edge from v to v'' , and we can continue the search process from v'' .

To show (3): Although (3) is intuitively obvious if one looks at a picture such as Figure 9, one has two difficulties in a formal setting. Firstly, showing that such a search process actually terminates

requires formally reasoning about the (undirected) acyclicity of graphs. We refer the interested reader to our appendix and Coq mechanization for details [Jacobs et al. 2021]. Secondly, one has to *restructure the argument* in order to even formally state what it means that “the search process terminates”. Our key idea is that the progress proof can be proved with an *inductive argument*, with a non-standard graph induction principle.

We call this induction principle for connectivity graphs *waiting induction*. The induction principle says that in order to prove $P(v)$ for all vertices $v \in V$, we can assume that $P(v')$ already holds for all vertices v' that v is waiting for. Note the similarity with strong induction on natural numbers: in order to prove $P(n)$ for all $n \in \mathbb{N}$, we can assume that $P(n')$ already holds for all $n' < n$.

We restructure the progress proof by applying our waiting induction principle at the start. Whenever we said “continue the search process” in the argument above, we can apply the inductive hypothesis. The induction principle is formally stated and discussed in more detail in §4.

What is parametric in this section. The waiting induction principle is parametric in the types of vertices V and labels L . This induction principle encapsulates the acyclicity reasoning, so that the progress proof can focus on the language-specific reasoning.

4 CONNECTIVITY GRAPHS AND WAITING INDUCTION IN DETAIL

Reasoning about graphs in a progress and preservation proof is non-standard, and reasoning about graphs in a proof assistant is more involved than reasoning about inductively-defined types like lists or maps that are normally used to define heap typings. We therefore factor graph reasoning out into a connectivity graph *library* that is *parametric* in vertices V and labels L . In this section we explain the foundations of this library by presenting the formal definition of acyclicity, a selected set of primitive rules (which are used to prove soundness of our separation-logic based graph transformations in §5 for proving preservation), as well as our principle of waiting induction (for proving progress). We hope to convince the reader that our graph-based approach is feasible—even in a mechanized setting in a proof assistant.

Recall the informal definition of connectivity graphs *Cgraph* from §3.2:

$$\text{Cgraph}(V, L) \triangleq \{G \in V \times V \xrightarrow{\text{fin}} L \mid G \text{ has no undirected cycles}\}$$

To define “ G has no undirected cycles” formally, we need to introduce some basic notions about graphs. We let $\text{graph}(V, L) \triangleq V \times V \xrightarrow{\text{fin}} L$ be graphs without the acyclicity restriction. The notation $v_1 \xrightarrow{l}_G v_2$ expresses that there is an edge from vertex v_1 to v_2 with label l (i.e., we have $G(v_1, v_2) = l$). The notation $v_1 \leftrightarrow_G v_2$ expresses that there is an edge from v_1 to v_2 or from v_2 to v_1 . The notation $v_1 \leftrightarrow_G^* v_2$ expresses that vertices v_1 and v_2 are *connected* by a (possibly empty) path from v_1 to v_2 where we may follow edges in either direction, and $v_1 \nleftrightarrow_G^* v_2$ expresses that there is no such path.

Definition 4.1 (Undirected acyclicity). A graph $G \in \text{graph}(V, L)$ has no undirected cycles if:

- (1) The undirected erasure $\bar{G} = \{\{v_1, v_2\} \mid v_1 \leftrightarrow_G v_2\}$, where we forget the labels and directions of the edges, is acyclic. See Jacobs et al. [2021] for details about the formalization of acyclicity of undirected graphs and the undirected erasure.
- (2) There are no short loops, i.e., we do not both have $v_1 \xrightarrow{l}_G v_2$ and $v_2 \xrightarrow{l'}_G v_1$.

Our reasoning about the acyclicity of graphs rests on two primitive lemmas:

Lemma 4.1 (Graph insertion). *If $G \in \text{graph}(V, L)$ is a graph with no undirected cycles and $v_1 \nleftrightarrow_G^* v_2$, then $G \cup \{v_1 \xrightarrow{l}_G v_2\}$ has no undirected cycles.*

Lemma 4.2 (Graph deletion). *If $G \in \text{graph}(V, L)$ is a graph with no undirected cycles and $v_1 \xrightarrow{l}_G v_2$, then $v_1 \nleftrightarrow_H^* v_2$ in the graph $H \triangleq G \setminus \{v_1 \xrightarrow{l}_G v_2\}$.*

We build a library of derived lemmas on top of these two primitive lemmas. A lemma that is crucial for proving the correctness of our separation logic rules in §5 is the exchange lemma, which is used to exchange separation logic resources between vertices of the graph:

Lemma 4.3 (Graph exchange). *Let $G, H \in \text{graph}(V, L)$ be graphs and let $v_1, v_2 \in V$ be vertices. If*

- (1) G has no undirected cycles,
- (2) $v_1 \not\leftrightarrow_G^* v_2$,
- (3) $\text{out}(G, v_1) \uplus \text{out}(G, v_2) = \text{out}(H, v_1) \uplus \text{out}(H, v_2)$, and
- (4) $\text{out}(G, v) = \text{out}(H, v)$ for all $v \in V \setminus \{v_1, v_2\}$.

Then:

- (1) H has no undirected cycles,
- (2) $v_1 \not\leftrightarrow_H^* v_2$, and
- (3) $\text{in}(G, v) = \text{in}(H, v)$ for all $v \in V$.

This lemma is quite a mouthful, so let us go over it step by step. We start with a graph G and we want to exchange outgoing edges between two unconnected vertices v_1 and v_2 to obtain a graph H in which the union of the outgoing edges of v_1 and v_2 stays the same. The lemma tells us that this operation maintains undirected acyclicity and that v_1 and v_2 are unconnected. Furthermore, the labels of incoming edges stay the same for all vertices.

Note that this property only holds because $\text{in}(G, v)$ is a multiset rather than a map that stores the vertices, like we did for $\text{out}(G, v)$. The fact that the local invariants are unaware of the vertices of origin of the incoming edges is what enables local reasoning: exchange of edges only affects the local invariants of v_1 and v_2 . In particular, for a channel it does not matter if its owner changes due to an exchange of resources, because it only matters at which type the channel is owned.

A typical pattern is to compose the lemma for exchange with with lemma for insertion and deletion. For instance, given an edge $v_1 \rightarrow_G^l v_2$, we can first delete the edge using Lemma 4.1 to obtain $v_1 \not\leftrightarrow_H^* v_2$. Then we can apply Lemma 4.3 to exchange some of the outgoing edges of v_1 and v_2 , and then we can re-insert a new edge $v_1 \rightarrow^{l'} v_2$ with a new label l using Lemma 4.1.

The lemmas for insertion and deletion (Lemmas 4.1 and 4.2) can not only be used to prove the acyclicity of modified connectivity graphs, but also to prove structural properties of connectivity graphs. The simplest example is a lemma that connectivity graphs have no self loops, which we give here as an illustration that the lemmas for insertion and deletion suffice.⁶

Lemma 4.4 (No self loops). *A connectivity graph $G \in \text{Cgraph}(V, L)$ has no self loops $v \rightarrow_G^l v$.*

PROOF. Suppose that $v \rightarrow_G^l v$. By Lemma 4.2, $v \not\leftrightarrow_{G'}^* v$ in the connectivity graph $G' \triangleq G \setminus \{v \rightarrow^l v\}$. Since every vertex is connected to itself (by definition), we have a contradiction. $\square$

Another example of a structural property that follows from the lemmas for insertion and deletion is the separation lemma. In §5 this lemma will play an important role in enabling our use of separation logic, where the separating conjunction requires that resources are disjoint.

Lemma 4.5 (Separation). *If $G \in \text{Cgraph}(V, L)$ and $v_1 \not\leftrightarrow_G^* v_2$ or $v_1 \leftrightarrow_G v_2$, then the outgoing edges of v_1 and v_2 are disjoint, i.e., $\text{dom}(\text{out}(G, v_1)) \cap \text{dom}(\text{out}(G, v_2)) = \emptyset$.*

Lastly, we have our generic principle of waiting induction that is key to our progress proof.

Lemma 4.6 (Waiting induction). *Let $G \in \text{Cgraph}(V, L)$ be a connectivity graph, $P \in V \rightarrow \text{Prop}$ a predicate over V , and $R : V \times V \rightarrow \text{Prop}$ a binary relation over V .*

⁶We do actually need this lemma at various points in the Coq proofs of the lemmas in §5.

Then in order to prove $\forall v \in V. P(v)$ it suffices to prove:

$$\begin{aligned} \forall v \in V. (\forall v' \in V. (v \xrightarrow{I}_G v' \wedge R(v, v')) \Rightarrow P(v')) \Rightarrow \\ (\forall v' \in V. (v' \xrightarrow{I}_G v \wedge \neg R(v', v)) \Rightarrow P(v')) \Rightarrow P(v) \end{aligned}$$

In other words, to prove $P(v)$, we can assume that $P(v')$ already holds for outgoing neighbors of v of v that are in relation $R(v, v')$, and we can also assume that $P(v')$ holds for incoming neighbors v' of v that are not in relation $R(v', v)$.

Thus, for neighbors $v \xrightarrow{I} v'$, either the proof of $P(v)$ can assume $P(v')$, or *vice versa*, but not both, and the relation $R(v, v')$ determines which. This induction principle is well founded due to the acyclicity of connectivity graphs. We prove this lemma using a similar lemma for undirected graphs, which we detail in [Jacobs et al. \[2021\]](#).

We call the lemma *waiting induction* because (1) we choose $R \triangleq \text{blocked}_{(\vec{e}, h)}$ from §3.6 and thus R is the waiting relation in our application, and (2) because of the similarity to induction on natural numbers: in order to prove $P(n)$ we can assume that $P(n-1)$ already holds, if $n \neq 0$.

5 LOCAL GRAPH TRANSFORMATION RULES IN SEPARATION LOGIC

We now generalize the well-formedness predicate wf from §3.4 to become parametric in the vertices V and labels L , which involves making it parametric in the local well-formedness predicate to abstract from language-specific aspects. We state separation-logic based rules for the parametric well-formedness predicate so that preservation can be proved using local reasoning. After an initial attempt at a monolithic proof of preservation, we found our approach of separating the graph reasoning from the local language-specific reasoning to be indispensable for mechanization.

Given a local well-formedness predicate $P : V \times \text{Multiset } L \rightarrow \text{Prop}$, we define the generic global well-formedness predicate $\text{wf}(P)$ as follows:

$$\text{wf}(P) \triangleq \exists G : \text{Cgraph}(V, L). \forall v \in V. P(v, \text{in}(G, v))(\text{out}(G, v))$$

We can instantiate the above definition with $P \triangleq \text{wf}_{(\vec{e}, h)}^{\text{local}}$ to obtain the well-formedness predicate from §3.4 that was tied to our concrete language.

Recall from §3.5 that preservation means: if $(\vec{e}, h) \rightsquigarrow_{\text{global}} (\vec{e}', h')$, then $\text{wf}(\text{wf}_{(\vec{e}, h)}^{\text{local}})$ implies $\text{wf}(\text{wf}_{(\vec{e}', h')}^{\text{local}})$. We now present a set of *graph transformation rules* for proving results “ $\text{wf}(P)$ implies $\text{wf}(P')$ ” where P and P' are arbitrary local well-formedness predicates, instead of a concrete local well-formedness predicate. These graph transformation rules perform a transformation of the graph under the hood, but the graphs are encapsulated by the definition of wf , and the rules thus do not mention any graphs. Instead, the premises of these graph transformation rules ask the user of the rule to prove *local* separation logic entailments involving P and P' .

The first of these graph transformation rules allows the user to exchange separation logic resources between two vertices $v_1, v_2 \in V$ in order to prove that $\text{wf}(P)$ implies $\text{wf}(P')$:

Lemma 5.1 (Exchange). *Let $v_1, v_2 \in V$. To prove $\text{wf}(P)$ implies $\text{wf}(P')$, it suffices to prove:*⁷

- (1) $P(v, \Delta) \multimap P'(v, \Delta)$ for all $v \in V \setminus \{v_1, v_2\}$ and $\Delta \in \text{Multiset } L$
- (2) $P(v_1, \Delta_1) \multimap \exists l. \text{own}(v_2 \mapsto l) \multimap \forall \Delta_2 \in \text{Multiset } L. \\ (P(v_2, \{l\} \uplus \Delta_2) \multimap \exists l'. (\text{own}(v_2 \mapsto l') \multimap P'(v_1, \Delta_1)) \multimap P'(v_2, \{l'\} \uplus \Delta_2))$ for $\Delta_1 \in \text{Multiset } L$

This rule generalizes the transformations for send and receive from Figure 8 where resources are exchanged between two vertices. We go over the premises of the rule in detail:

⁷Recall that proving $P \in \text{Prop}$ means proving $P(\emptyset)$ (see §3.3), but in practice (and in Coq) this is done using the proof rules of separation logic.

- (1) The first premise asks the user of the rule to prove the local implication $P(v, \Delta) \multimap P'(v, \Delta)$ for the vertices $v \in V \setminus \{v_1, v_2\}$ that are not involved in the exchange.
- (2) The second premise first gives the user access to the local resources $P(v_1, \Delta_1)$ of vertex v_1 . The rule then asks the user to prove that there exists an edge $v_1 \xrightarrow{I}_G v_2$, by showing that $\exists l. \text{own}(v_2 \mapsto l) * \dots$ follows from the local resources of v_1 . The rule then gives the user access to the local resources $P(v_2, \{l\} \uplus \Delta_2)$ of vertex v_2 , where we have obtained the information that the label $\{l\}$ is part of the incoming edge label multiset of v_2 . The rule then allows the user to pick a new label l' for the edge $v_1 \xrightarrow{l'} v_2$. Subsequently, the user has to restore the local resources of v_1 and v_2 for the new P' . For restoring the local resources $P'(v_1, \Delta_1)$, the user additionally gets the $\text{own}(v_2 \mapsto l')$ of the new edge. For restoring the local resources $P'(v_2, \{l'\} \uplus \Delta_2)$, we get the new label in the incoming edge label multiset.

It may seem like this rule only allows us to change the label on the edge $v_1 \xrightarrow{l}_G v_2$ from l to l' , but the rule in fact allows us to arbitrarily exchange separation logic resources (i.e., outgoing edges) between v_1 and v_2 . The way this works is that the rule gives us access to the old local resources of both v_1 and v_2 , and it asks us to prove the separating conjunction of the new local resources of both v_1 and v_2 . The proof rules of separation logic allow us to use resources stored in the old local resources of v_1 to prove the new local resources of v_2 , and *vice versa*. Thus, the graph transformation that is applied internally in the rule depends on *which* proof of the separation logic entailment is provided by the user of the rule.

A note on the proof of the transformation rule. That the transformation rule is able to offer us access to both local resources simultaneously relies crucially on the acyclicity of the graph. The acyclicity, and the existence of an edge between the two vertices, is what allows us to apply the separation lemma (Lemma 4.5) that allows us to construct the separating conjunction of the two local resources. In the proof of the rule we re-establish the validity of the resources and the acyclicity of the graph using the exchange lemma (Lemma 4.3).

In addition to the preceding transformation rule for changing the label on an edge (and exchanging resources), we have a transformation rule for removing an edge (after exchanging resources). This rule is used in the close case of the preservation proof:

Lemma 5.2 (Deallocation). *Let $v_1, v_2 \in V$. To prove $\text{wf}(P)$ implies $\text{wf}(P')$, it suffices to prove:*

- (1) $P(v, \Delta) \multimap P'(v, \Delta)$ for all $v \in V \setminus \{v_1, v_2\}$ and $\Delta \in \text{Multiset } L$
- (2) $P(v_1, \Delta_1) \multimap \exists l. \text{own}(v_2 \mapsto l) * \forall \Delta_2 \in \text{Multiset } L.$
 $(P(v_2, \{l\} \uplus \Delta_2) \multimap P'(v_1, \Delta_1) * P'(v_2, \Delta_2))$ for $\Delta_1 \in \text{Multiset } L$

We have the following two transformation rules for inserting an outgoing/incoming edge between v_1 and v_2 , respectively. To maintain acyclicity, we have to show that v_2 has no existing outgoing edges. Like the preceding rules, these rules also allow us to transfer resources to v_2 .

Lemma 5.3 (Allocation out). *Let $v_1, v_2 \in V$. To prove $\text{wf}(P)$ implies $\text{wf}(P')$, it suffices to prove:*

- (1) $P(v, \Delta) \multimap P'(v, \Delta)$ for all $v \in V \setminus \{v_1, v_2\}$ and $\Delta \in \text{Multiset } L$
- (2) $P(v_2, \Delta_2) \multimap \top \Delta_2 = \emptyset \top$ for all $\Delta_2 \in \text{Multiset } L$
- (3) $P(v_1, \Delta_1) \multimap \exists l'. (\text{own}(v_2 \mapsto l') \multimap P'(v_1, \Delta_1)) * P'(v_2, \{l'\})$ for all $\Delta_1 \in \text{Multiset } L$

Lemma 5.4 (Allocation in). *Let $v_1, v_2 \in V$. To prove $\text{wf}(P)$ implies $\text{wf}(P')$, it suffices to prove:*

- (1) $P(v, \Delta) \multimap P'(v, \Delta)$ for all $v \in V \setminus \{v_1, v_2\}$ and $\Delta \in \text{Multiset } L$
- (2) $P(v_2, \Delta_2) \multimap \top \Delta_2 = \emptyset \top$ for all $\Delta_2 \in \text{Multiset } L$
- (3) $P(v_1, \Delta_1) \multimap \exists l'. P'(v_1, \Delta_1 \uplus \{l'\}) * (\text{own}(v_1 \mapsto l') \multimap P'(v_2, \emptyset))$ for all $\Delta_1 \in \text{Multiset } L$

Lastly, we have a derived transformation rule that adds two new edges $v_1 \xrightarrow{l'_1} v_2$ and $v_2 \xleftarrow{l'_2} v_3$. We use this rule in the fork case of the preservation proof.

$$\begin{array}{c}
\frac{\Gamma \text{ unrestricted}}{\{x \mapsto \tau\} \uplus \Gamma \vdash x : \tau} \quad \frac{n \in \mathbb{N} \quad \Gamma \text{ unrestricted}}{\Gamma \vdash n : \mathbb{N}} \quad \frac{\Gamma_1 \perp \Gamma_2 \quad \Gamma_1 \vdash e_1 : \tau_1 \quad \Gamma_2 \vdash e_2 : \tau_2}{\Gamma_1 \cup \Gamma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2} \\
\\
\frac{\Gamma \uplus \{x \mapsto \tau_1\} \vdash e : \tau_2 \quad \Gamma \text{ unrestricted}}{\Gamma \vdash \lambda x. e : \tau_1 \rightarrow \tau_2} \quad \frac{\Gamma_1 \perp \Gamma_2 \quad \Gamma_1 \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma_2 \vdash e_2 : \tau_1}{\Gamma_1 \cup \Gamma_2 \vdash e_1 e_2 : \tau_2}
\end{array}$$

Fig. 10. Selected typing rules for unrestricted types.

Lemma 5.5 (Allocation out and in). *Let $v_1, v_2, v_3 \in V$. To prove $\text{wf}(P)$ implies $\text{wf}(P')$, it suffices to prove:*

- (1) $P(v, \Delta) \multimap P'(v, \Delta)$ for all $v \in V \setminus \{v_1, v_2, v_3\}$ and $\Delta \in \text{Multiset } L$
- (2) $P(v, \Delta) \multimap \ulcorner \Delta = \emptyset \urcorner$ for all $v \in \{v_2, v_3\}$ and $\Delta \in \text{Multiset } L$
- (3) $P(v_1, \Delta_1) \multimap \exists l'_1, l'_2. (\text{own}(v_2 \mapsto l'_1) \multimap P'(v_1, \Delta_1 \uplus \{l'\})) \multimap P'(v_2, \{l'_1, l'_2\}) \multimap$
 $\quad (\text{own}(v_2 \mapsto l'_2) \multimap P'(v_3, \emptyset))$
 for all $\Delta_1 \in \text{Multiset } L$

This rule can be proved by applying both allocation out and allocation in. It pays off to prove this in the generic setting, because the intermediate state (in which the channel has been allocated but not yet the thread that will hold the other endpoint) is not well-formed according to our wf^{local} . Instead, we prove $\text{wf}(P) \implies \text{wf}(P')$ by carefully choosing Q and proving $\text{wf}(P) \implies \text{wf}(Q)$ using Lemma 5.3, and $\text{wf}(Q) \implies \text{wf}(P')$ using Lemma 5.4.

6 EXTENSIONS

The programming language for which we have mechanized deadlock and memory leak freedom in Coq [Jacobs et al. 2021] supports more features than described in §2. First, it has more standard features such as sum types, which we do not describe because their rules are standard and the modification of the proof is straightforward. Second, it has unrestricted (non-linear) types, including unrestricted products and sums, and an unrestricted function type (§6.1) and general equi-recursive functional types (which can encode algebraic data types) and equi-recursive recursive session types (which can encode infinite protocols) (§6.2). Furthermore, we prove a deadlock freedom property that is stronger than global progress and also rules out partial deadlock (§6.3).

6.1 Unrestricted Types

We make the types used for conventional functional programming (such as product, sum, and function types) unrestricted (i.e., non-linear) if their components are unrestricted. Instead of introducing separate linear and non-linear products and sums, we introduce the judgment “ τ unrestricted” on types τ , which holds if all the components of τ are unrestricted. Formally, the base types $\mathbf{0}$, $\mathbf{1}$ and $\mathbb{N}$ are unrestricted, and $\tau_1 \times \tau_2$ and $\tau_1 + \tau_2$ are unrestricted if τ_1 and τ_2 are unrestricted. The type $\tau_1 \multimap \tau_2$ is always linear (i.e., not unrestricted), even if τ_1 and τ_2 are unrestricted, because the closure may capture linear variables. We introduce the type $\tau_1 \rightarrow \tau_2$ of unrestricted functions, which is always unrestricted (even if τ_1 and τ_2 are linear), and whose closures are only allowed to capture unrestricted variables. Selected typing rules are shown in Figure 10. The typing rules involve the *disjointness relation* $\Gamma_1 \perp \Gamma_2$, which expresses that Γ_1 and Γ_2 might share unrestricted variables, but otherwise do not overlap. Formally:

$$\Gamma_1 \perp \Gamma_2 \triangleq \forall x \in \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2). \Gamma_1(x) = \Gamma_2(x) \wedge \Gamma_1(x) \text{ unrestricted}$$

Changes to the proof. In order to reason about unrestricted values in the separation logic, we add a standard box modality $\Box P$, defined as $(\Box P)(\Sigma) \triangleq P(\emptyset) \wedge \Sigma = \emptyset$. The box modality asserts that P does not use any linear resources, which allows it to support proof rules for deletion ($\Box P \multimap \text{Emp}$) and duplication ($\Box P \multimap \Box P \ast \Box P$). Lastly, we have the rule $\Box P \multimap P$ to open the box. We use the box modality in the run-time typing rule for the unrestricted function type:

$$\frac{\Box (\Gamma \uplus \{x \mapsto \tau_1\} \vdash e : \tau_2) \quad \ulcorner \Gamma \text{ unrestricted} \urcorner}{\Gamma \vdash \lambda x. e : \tau_1 \rightarrow \tau_2} \ast$$

The box modality makes sure that the closure cannot capture any channels at run-time.

We prove $(\Gamma \vdash e : \tau) \multimap \Box (\Gamma \vdash e : \tau)$ if τ unrestricted. This entailment says that run-time typing judgments for expressions e of unrestricted type τ can be freely deleted and duplicated in the separation logic sense. This is crucial for the main change to our proof—the substitution lemma—in which we now have to consider the case that the type is unrestricted, and that a variable could be substituted in multiple or zero places. We use the preceding entailment and the laws of the box modality to adapt the proof of the substitution lemma.

6.2 Equi-Recursive Types

We extend our type system with equi-recursive functional ($\mu\alpha.\tau$) and session types ($\mu\alpha.s$), in order to be able to encode algebraic data types and infinite protocols, respectively. We extend the type system with the following rule for unfolding recursive types:

$$\frac{\Gamma \vdash e : \tau_1 \quad \tau_1 \equiv \tau_2}{\Gamma \vdash e : \tau_2}$$

The congruence relation ($\equiv$) relates types up to unfolding of $\mu\alpha.\tau \equiv \tau[\mu\alpha.\tau/\alpha]$. Our mechanization (§7) is somewhat more general: we use a coinductive definition of types to allow mutual recursion and recursion through the message type as well as the tail. We also extend unrestricted types to allow recursive types to be unrestricted. We can encode algebraic data types such as lists by using sums and products and recursive types.

Changes to the proof. We do *not* add a rule for unfolding recursive types to the run-time type system. Rather, we define the run-time type system in a syntax directed way so that all constructors respect the congruence relation ($\equiv$), and then *prove* a version of the unfolding rule:

Lemma 6.1. *If $\Gamma_1 \equiv \Gamma_2$ and $\tau_1 \equiv \tau_2$, then $(\Gamma_1 \vdash e : \tau_1) \multimap (\Gamma_2 \vdash e : \tau_2)$.*

Example. The combination of equi-recursive and unrestricted types allows us to type check the call-by-value Y-combinator for constructing recursive functions of type $\tau_1 \rightarrow \tau_2$. Defining recursive functions in terms of a self-referential type is standard [Harper 2016]:

$$\begin{aligned} Y &: ((\tau_1 \rightarrow \tau_2) \rightarrow (\tau_1 \rightarrow \tau_2)) \rightarrow (\tau_1 \rightarrow \tau_2) \\ Y &\triangleq \lambda f. (\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y)) \end{aligned}$$

We use the recursive type $\mu\alpha.(\alpha \rightarrow (\tau_1 \rightarrow \tau_2))$ for x and the type $(\tau_1 \rightarrow \tau_2) \rightarrow (\tau_1 \rightarrow \tau_2)$ for f . Note that while τ_1 and τ_2 can be restricted (linear) types, we *must* use an unrestricted function type for f and x in order to type check the multiple uses of f and x . In fact, a fixed-point combinator for constructing functions $\tau_1 \multimap \tau_2$ with linear function type would violate type safety. Intuitively, a recursive function is allowed to *manipulate* both linear and non-linear resources, but the *definition* of a recursive function is not allowed to *capture* linear resources in its closure because this closure will be invoked multiple times.

6.3 Partial Deadlock and Memory Leak Freedom via Reachability

In the context of session-typed languages with non-termination (e.g., due to recursive types), deadlock freedom is typically stated as *global progress*, which we prove in §3.6. Global progress guarantees that the configuration can either take a step, or is in a final state where all threads have successfully terminated and all channels have been deallocated. Although global progress rules out whole-program deadlocks, as well as memory leaks when all threads have terminated, it admits partial deadlocks as long there is still a thread that can step (e.g., is in an infinite loop). Linear session types actually rule out partial deadlocks and memory leaks even when some threads are still running. Although deadlock freedom and memory leak freedom may seem like separate properties, we state two properties that simultaneously generalize both, namely *partial deadlock/leak freedom* (Definition 6.3) and *full reachability* (Definition 6.5). We prove that these properties are equivalent (Theorem 6.2) and show that full reachability can be proven using the waiting induction principle of our proof method (Theorem 6.3). Finally, we show that they imply global progress.

In order to arrive at a simultaneous generalization of deadlock and memory leak freedom, consider pure memory leaks and pure deadlocks:

- A *pure memory leak* is one in which we have a set S of channels, such that all endpoints of the channels in S are held by the buffers of channels in the same set S .
- A *pure deadlock* is a set S of both threads and channels with empty buffers, such that all threads in S are blocked on one of the channels in the set S , and all of the endpoints of the channels in S are held by threads in the set S .

In general, we can have a mixed partial deadlock/leak situation that is neither a pure memory leak nor a pure deadlock. Intuitively, a partial deadlock and memory leak is a set S of threads and channels such that all threads in S are blocked on one of the channels in S , and all endpoints of channels in S are held by threads and channels in S . To make this formal, we define the set of vertices $\text{refs}_{(\vec{e}, h)}(v) \subseteq V$ that a vertex v references.

Definition 6.1. We let $\text{refs}_{(\vec{e}, h)}(\text{Thread}(i)) \triangleq \{\text{Chan}(a') \mid \text{channel literal } \#(a', t) \text{ occurs in } e_i\}$, and $\text{refs}_{(\vec{e}, h)}(\text{Chan}(a)) \triangleq \{\text{Chan}(a') \mid \text{channel literal } \#(a', t) \text{ occurs in } h(\#(a, 0)) \text{ or } h(\#(a, 1))\}$.

With this function at hand, we can define partial deadlock and memory leak freedom.

Definition 6.2 (Partial deadlock/leak). Given a configuration $(\vec{e}, h)$, a subset $S \subseteq V$ of the threads and channels is in a partial *deadlock/leak* if the following conditions hold:

- (1) We have $\emptyset \subset S \subseteq \text{active}(\vec{e}, h)$ (see Definition 3.1 for the definition of active).
- (2) For all threads $\text{Thread}(i) \in S$, the expression e_i cannot step in the heap h .
- (3) If $\text{Thread}(i) \in S$ and $\text{blocked}_{(\vec{e}, h)}(\text{Thread}(i), \text{Chan}(a))$, then $\text{Chan}(a) \in S$ (see Definition 3.2 for the definition of blocked).
- (4) If $\text{Chan}(a) \in S$ and $\text{Chan}(a) \in \text{refs}_{(\vec{e}, h)}(v)$, then $v \in S$.

Definition 6.3 (Partial deadlock/leak freedom). A configuration $(\vec{e}, h)$ is *deadlock/leak free* if no $S \subseteq V$ is in a partial deadlock/leak in $(\vec{e}, h)$.

In order to prove that well-formed configurations have no partial deadlock/leak, we prove another property that we call *full reachability*, which we show to be equivalent to partial deadlock/leak freedom. Full reachability has the advantage that it can be proved directly using waiting induction. It takes inspiration from the notion of reachability used in garbage collection and memory management, namely that data is said to be reachable if it can be reached by transitively following pointers, starting from any thread's stack frames. Memory leak freedom can then be stated as: all data in the configuration is reachable, i.e., there is never any leaked memory. To incorporate deadlock freedom into this, we strengthen the definition of reachability to only start from stack frames of *threads that*

can step. However, if a thread T1 is blocked on channel C, and the other endpoint of C is held by still running thread T2, then data held by T1 should also be considered transitively reachable: even though this data is held by a thread that (currently) cannot step, further interaction of T2 with the channel C may unblock T1. We formalize this using the following inductive definition:

Definition 6.4 (Reachability). We inductively define the vertices that are *reachable* in $(\vec{e}, h)$:

- (1) Thread(i) is reachable if either
 - the expression e_i can step in the heap h , or
 - there exists an a such that $\text{blocked}_{(\vec{e}, h)}(\text{Thread}(i), \text{Chan}(a))$ and $\text{Chan}(a)$ is reachable.
- (2) $\text{Chan}(a)$ is reachable if there exists a reachable v such that $\text{Chan}(a) \in \text{refs}_{(\vec{e}, h)}(v)$.

It is important that reachability is an inductive definition—a coinductive definition would trivially consider all cycles to be reachable.

Definition 6.5 (Full reachability). A configuration $(\vec{e}, h)$ is *fully reachable* if all $v \in \text{active}(\vec{e}, h)$ are reachable in $(\vec{e}, h)$.

We show equivalence of partial deadlock/leak freedom and full reachability:

Theorem 6.2. A configuration $(\vec{e}, h)$ is deadlock/leak free if and only if it is fully reachable.

For $(\Rightarrow)$, we show that none of the objects in a deadlock/leak are reachable, and for $(\Leftarrow)$, we show that the set of all non-reachable objects is a deadlock/leak.

Theorem 6.3 (Full reachability). If $\text{wf}(\vec{e}, h)$, then $(\vec{e}, h)$ is fully reachable.

This proof goes by waiting induction and closely resembles the global progress proof in §3.6. By using the equivalence between full reachability and partial deadlock/leak freedom, we also obtain that a partial deadlock/leak cannot occur, and can re-prove global progress using reachability.

Corollary 6.3.1 (Partial deadlock/leak freedom). If $\text{wf}(\vec{e}, h)$, then $(\vec{e}, h)$ is deadlock/leak free.

Corollary 6.3.2 (Global progress'). If $\text{wf}(\vec{e}, h)$ and $\text{active}(\vec{e}, h) \neq \emptyset$, then $(\vec{e}, h)$ can step.

The proof of Corollary 6.3.2 uses Theorem 6.3, which gives that active objects are reachable. We then find a thread that can step by straightforward induction on the reachability predicate. Alternatively, we can go via Corollary 6.3.1: if none of the threads can step, then the set of all active threads and channels is a deadlock/leak.

Combined with the proofs that the initial configuration ρ of well-typed program satisfies $\text{wf}(\rho)$, and that $\text{wf}(\rho)$ is preserved by the operational semantics (§3.5), we obtain partial deadlock/leak freedom, full reachability, and global progress for any well-typed program.

7 MECHANIZATION IN COQ

Using the Coq proof assistant [Coq Team 2021] we have mechanized the generic connectivity graph method and its concrete instantiation to our session-typed language. Our mechanization starts with a library for undirected graphs and their acyclicity described in Jacobs et al. [2021]. On top of this, we build a library for connectivity graphs and waiting induction (§4). We combine connectivity graphs with separation logic (§3.3) to define the generic well-formedness predicate and the separation logic local transformation lemmas (§5). We instantiate our library by formalizing the language from §2 with its extensions from §6. This involves defining the syntax, type system, and operational semantics. For the language-specific parts of our deadlock and leak freedom proof, we define the run-time type system (§3.3) and the local well-formedness condition (§3.4). We then prove preservation using our local transformation rules in separation logic (§3.5), and progress using our principle of waiting induction (§3.6). We have also mechanized all the extensions (§6),

including unrestricted types (§6.1), equi-recursive types (§6.2), and the theorems about reachability and partial deadlock/leak freedom (§6.3).

Line counts. The parametric connectivity graph library is 4999 LOC, the language definition is 451 LOC, and the language-specific deadlock and leak freedom proofs are 1688 LOC.

External dependencies and Coq features that we use. We use the `std++` extended standard library for its results on data structures like lists and finite maps [Coq-std++ Team 2021]. We use the Iris Proof Mode for tactics-based separation logic proofs [Krebbers et al. 2017, 2018]. To represent recursive types (§6.2), we use the technique by Gay et al. [2020] based on coinductive types combined with Coq’s generalized rewriting mechanism to reason up to the congruence $\equiv$ [Sozeau 2009].

Archive and GitHub repository. An archive of the Coq mechanization can be found at Jacobs et al. [2021], and the most recent version at <https://github.com/julesjacobs/cgraphs>.

8 RELATED WORK

Session types. The line of works most closely related to ours are derivatives of Wadler [2012]’s GV, a linear functional language with session types inspired by Gay and Vasconcelos [2010]. Whereas Gay and Vasconcelos’s calculus does not enjoy the property of deadlock freedom, Wadler’s GV and its derivatives [Lindley and Morris 2015, 2016b, 2017; Fowler et al. 2019, 2021] do. For Wadler’s GV, deadlock freedom follows from its translation to CP (Classical Processes) Wadler [2012], for which deadlock freedom holds by cut elimination. Lindley and Morris [2015] then concretize the progress statement by introducing the definition of a deadlocked configuration and proving deadlock freedom using a small-step operational semantics. They also give translations between GV and CP and show that both directions of the translation preserve reductions, unlike previous translations from GV to CP. Subsequently, Lindley and Morris [2015]’s GV has been extended to support least and greatest fixed points [Lindley and Morris 2016b], exceptions [Fowler et al. 2019], and polymorphism [Lindley and Morris 2017]. A recent extension of GV [Fowler et al. 2021] moreover simplifies GV’s meta theory by making process equivalence type preserving. The extension adopts the idea of a hypersequent [Avron 1991] from [Montesi and Peressotti 2018; Kokke et al. 2019], yielding Hypersequent GV (HGV).

Like the GV derivatives, our language is a functional language with session-typed channels. Our notion of a connectivity graph moreover bears a resemblance to HGV’s abstract process structure (APS), introduced to reason about the acyclic forest structure of a process configuration. However, whereas abstract process structures are defined over hyperenvironments and channel names, our connectivity graph is parametric in its vertices, labels, and edges. More importantly, our connectivity graph is at the core of a proof method for deadlock freedom, fully mechanized in Coq, that uses separation logic and is parametric in its key results. Besides these conceptual differences, there are various technical differences between our formalization and GV formalizations, and even among the different GV variants (such as a synchronous versus an asynchronous semantics). Our formalization uses a standard operational semantics whereas many GV variants use structural congruences and binders for channel names. In our graph, not only the threads but also channels are vertices, and the edges are directed. Since reasoning about syntax up to equivalence (e.g., structural congruence or α -equivalence) is cumbersome in a proof assistant like Coq, we believe that our operational semantics is better suited for mechanization (and perhaps closer to how these structures are represented on real computers). Orthogonally, we do not tie channel closing to thread termination and allow close everywhere. As a result our language readily accommodates a forest topology without the need for a special connective, such as mix as used by Fowler et al. [2021].

Earlier non-mechanized work has proved deadlock freedom for a π -calculus using a graphical approach [Carbone and Debois 2010]. This is the earliest work that we are aware of that describes an explicit connection between deadlock freedom and acyclicity of a graph. Their graphical representation is, however, an undirected graph between processes, whereas our connectivity graphs (when instantiated for our language) are directed graphs between threads and channels. Furthermore, their graphs are unlabeled, whereas our connectivity graphs are labeled with session types.

More distantly, our language is related to Toninho et al. [2013]; Toninho [2015]’s language SILL, which embeds session-typed processes into a functional core language via a contextual monad. The language is based on the Curry-Howard correspondence established by Caires and Pfenning [2010] between intuitionistic linear logic and session-typed π -calculus. Deadlock freedom of SILL follows thus as a consequence. Due to its modal separation, SILL does not allow mixing of functional and session terms freely, in contrast to GV and our language. The seminal paper by Caires and Pfenning [2010] and Toninho [2015]’s thesis spurred a series of derivatives, similarly to Wadler’s CP and GV, accommodating, for example, polymorphism [Caires et al. 2013; Pérez et al. 2014], work analysis [Das et al. 2018], and information flow control [Derakhshan et al. 2021]. Due to their connection to intuitionistic linear logic, all these works guarantee deadlock freedom. However, unlike ours, none of these deadlock freedom proofs have been mechanized in a proof assistant.

A derivative of SILL, SILL_S [Balzer and Pfenning 2017], introduces a controlled form of aliasing through a stratification of linear and shared session types with adjoint modalities [Pfenning and Griffith 2015; Benton 1994; Reed 2009b] to support multiple-client scenarios. Whereas the resulting language reclaims the expressiveness of the untyped asynchronous π -calculus for session-typed languages [Balzer et al. 2018], it also sacrifices deadlock freedom (which is rectified by its successor SILL_S⁺ [Balzer et al. 2019]). Recent extensions of classical linear logic session types contribute another approach to softening the rigidity of linear session types to support multiple client sessions and nondeterminism [Qian et al. 2021] and memory cells and nondeterministic updates [Rocha and Caires 2021], respectively. Whereas neither of these recent approaches reclaim the full expressiveness of unrestricted sharing, they keep the logical foundation intact and thus uphold deadlock freedom. However, none of these works have been mechanized in a proof assistant.

Prior to the development of logic-based session types [Caires and Pfenning 2010; Wadler 2012], deadlock freedom in session-typed calculi [Vasconcelos 2012] was guaranteed only for processes interacting on a *single* session—interleaving of blocking actions on different sessions could easily result in deadlocks. To address limitations of classical binary session types, Honda et al. [2008] introduced *multiparty* session types, where sessions are described by so-called global types that capture the interactions between an arbitrary number of session participants. Given some well-formedness constraints, global types can ensure that a collection of processes correctly implement the global behavior in a deadlock-free way. However, these global type-based approaches do not ensure deadlock freedom in the presence of higher-order channels, interleaved sessions, dynamic channel creation, or dynamic thread creation. To remedy the deficiency various extensions at increasing degrees of complexity were introduced. For example, Bettini et al. [2008] and Coppo et al. [2016] track usage orders among interleaved multiparty sessions, ruling out cyclic dependencies but also restricting recursion. Our approach instead supports higher-order channels, general recursion, and deadlock freedom solely using a linear type system, by restricting to binary sessions.

Separation logic. Separation logic [O’Hearn et al. 2001] is conventionally used in Hoare-style program logics for proving functional correctness, while we use it to define and reason about (run-time) typing judgments. In conventional separation logic, propositions are predicates over heaps (possibly extended with permissions, ghost state, *etc.*), whereas we consider predicates over the outgoing edges of a connectivity graph (which contain types instead of values). The idea of using

separation logic to define typing judgments for linear languages has been explored by Rouvoet et al. [2020, 2021] in the context of intrinsically-typed programming in Agda. They present separation-logic based programming abstractions to hide types of references in intrinsically-typed interpreters, and to hide types of labels in intrinsically-typed compilers. As a case study, Rouvoet et al. [2020] use their abstractions to define an intrinsically-typed interpreter for a small session-typed language that guarantees type safety by construction (but not deadlock or resource leak freedom).

Separation logic has also been used to define logical relation models of affine type systems. For example, logical relations in the Iris separation logic [Jung et al. 2015, 2018b] have been used for proving memory safety and data race freedom of Rust [Jung et al. 2018a], as well as type safety of session types [Hinrichsen et al. 2021b]. To extend the logical-relations based approach to prove deadlock freedom, a full-fledged separation logic that is capable of proving deadlock freedom is needed. While separation logics and Hoare logics with support for deadlock freedom exist, e.g., [Hamin and Jacobs 2018; Le et al. 2013; Zhang et al. 2016], they use lock-orders, whose logical expressivity is different from session types. Some separation logics have support for pointed-by assertions [Kassios and Kritikos 2013], which can be used to reason about memory leak freedom.

Various extensions of separation logic that incorporate session-type based mechanisms to reason about message-passing programs have been developed, e.g., Francalanza et al. [2011]; Lozes and Villard [2012]; Craciun et al. [2015]; Oortwijn et al. [2016]; Hinrichsen et al. [2020, 2021a]. The goal of these logics is different from ours—they are full-fledged Hoare logics aimed at proving functional correctness instead of deadlock freedom. On the other hand, we use the assertion layer of separation to hide bookkeeping in the definition of run-time typing judgments, and to describe connectivity graph transformations in an abstract and generic way.

Mechanized results of session types. Thiemann [2019] proves type safety of a linear λ -calculus with session types that is inspired by GV. They do not prove deadlock or memory leak freedom. Their mechanization involves an extensive amount of bookkeeping to keep track of resources. Rouvoet et al. [2020] streamlined this approach via separation logic (see discussion above).

Hinrichsen et al. [2021b] prove type safety for a comprehensive session-typed language with locks, subtyping and polymorphism using Iris in Coq. Their type system is affine, which means that deadlocks are considered safe (their receive operation will spin if the buffer is empty). Their proof is based on logical relations instead of progress and preservation (see discussion above).

Tassarotti et al. [2017] prove correctness of a compiler for an affine session-typed language using Iris in Coq. The operational semantics of their source language is similar to ours, while channels are compiled to an implementation involving linked lists in the target. Their compiler is proved to be termination preserving, so a target program deadlocks iff the source deadlocks.

More distantly, there also exist various mechanized results involving π -calculus. Goto et al. [2016] prove type safety for a π -calculus with a polymorphic session type system in Coq. Their type system allows dropping channels, and hence does not enjoy deadlock nor memory leak freedom. Ciccone and Padovani [2020] mechanize dependent binary session session types by embedding them into a π -calculus in Agda. They prove subject reduction (i.e., preservation) and that typing is preserved by structural congruence. Neither deadlock freedom nor leak freedom is proved. Castro-Perez et al. [2020] present a framework for locally-nameless representations of π -calculus in Coq. They use their framework to prove subject reduction (i.e., preservation) of a type system for binary session types. Neither deadlock freedom nor leak freedom is proved. Their framework is used by Castro-Perez et al. [2021] to mechanize a DSL for multiparty communication in Coq based on asynchronous multiparty session types. They prove deadlock freedom w.r.t. a global type, but do not prove deadlock freedom in the presence of higher-order channels, interleaved sessions, dynamic channel creation, or dynamic thread creation.

Gay et al. [2020] study various notions of duality in Agda, and show that distribution laws for duality over the recursion operator are unsound. Unlike the other mechanized results discussed so far, they focus on the static instead of dynamic semantics of session types. We have adapted their approach of using coinductive types for mechanizing general recursive session types (see §7). Keizer et al. [2021] use coalgebras to model session types in a non-mechanized setting.

More distantly related are mechanized versions of cut elimination of linear logic [Reed 2009a; Chaudhuri et al. 2019], which by Curry-Howard relates to deadlock freedom of intuitionistic session types. The authors were incentivized by mistakes in various existing, non-mechanized proofs. However, whereas a cut elimination proof concerns a logical inference system only, our proof of deadlock freedom encompasses a typed programming language with operational semantics, requiring us to reason not only about its statics but also its execution semantics. Moreover, our language includes features such as recursive types (§6.2) that break cut elimination.

Mechanization results, lastly, also exist for choreographic languages [Montesi 2021]. Cruz-Filipe et al. [2021a] mechanize choreography compilation in Coq for the choreographic language Core Choreographies (CC) introduced by Cruz-Filipe et al. [2021b]. CC supports recursion and its semantics has been formalized in Coq by Cruz-Filipe et al. [2021b]. Key results of the formalization include determinism, confluence, and deadlock-freedom by design as well as Turing completeness.

Process calculi. The addition of channel usage information to types in a concurrent, message-passing setting was pioneered by Kobayashi [1997]; Igarashi and Kobayashi [1997]; Kobayashi et al. [1999], who applied the idea to deadlock prevention in the π -calculus and later to more general properties [Igarashi and Kobayashi 2001, 2004], giving rise to a generic system that can be instantiated to produce a variety of concrete typing disciplines for the π -calculus. Typically, types are augmented with the relative ordering of channel actions, with the type system ensuring that the transitive closure of such orderings forms a strict partial order, ensuring deadlock-freedom. Building on this, Kobayashi [2002] proposed type systems that ensure a stronger property, dubbed lock freedom, and variants that are amenable to type inference [Kobayashi et al. 2000]. Kobayashi [2006] extended this to account for recursive processes and type inference. Kobayashi-style systems have also been adopted for session-typed languages [Dardha and Gay 2018; Balzer et al. 2019].

9 FUTURE WORK

We have used our connectivity graph method to give a mechanized proof of deadlock and memory-leak freedom for binary session types. Since connectivity graphs are not restricted to two incoming edges per channel, we would like to explore language designs with a version of multiparty session types that supports dynamic thread and channel creation, and higher order channels, but still enjoys global progress from typing in a manner similar to binary session types (*i.e.*, without additional mechanisms such as orders or priorities). Second, we would like to explore whether other concurrency mechanisms such as locks and barriers could be handled by our method.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful feedback, and especially for the criticism that global progress is a very weak property if one thread is in an infinite loop, which is what led us to develop §6.3. We are grateful to Herman Geuvers, Fabrizio Montesi, Ike Mulder, Arjen Rouvoet, Bernardo Toninho, and Jorge Pérez for discussions about this paper and related work. The second author (Stephanie Balzer) was supported by National Science Foundation Award No. CCF-1718267. The third author (Robbert Krebbers) was supported by the Dutch Research Council (NWO), project 016.Veni.192.259.

REFERENCES

- Arnon Avron. 1991. Hypersequents, Logical Consequence and Intermediate Logics for Concurrency. *Annals of Mathematics and Artificial Intelligence* 4 (1991), 225–248. <https://doi.org/10.1007/BF01531058>
- Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. *PACMPL* 1, ICFP (2017), 37:1–37:29. <https://doi.org/10.1145/3110281>
- Stephanie Balzer, Frank Pfenning, and Bernardo Toninho. 2018. A Universal Session Type for Untyped Asynchronous Communication. In *CONCUR (LIPIcs, Vol. 118)*. 30:1–30:18. <https://doi.org/10.4230/LIPIcs.CONCUR.2018.30>
- Stephanie Balzer, Bernardo Toninho, and Frank Pfenning. 2019. Manifest Deadlock-Freedom for Shared Session Types. In *ESOP (LNCS, Vol. 11423)*. 611–639. https://doi.org/10.1007/978-3-030-17184-1_22
- Nick Benton. 1994. A Mixed Linear and Non-Linear Logic: Proofs, Terms and Models (Extended Abstract). In *CSL (LNCS)*. 121–135. <https://doi.org/10.1007/BFb0022251>
- Lorenzo Bettini, Mario Coppo, Loris D'Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. 2008. Global Progress in Dynamically Interleaved Multiparty Sessions. In *CONCUR (LNCS, Vol. 5201)*. 418–433. https://doi.org/10.1007/978-3-540-85361-9_33
- Luís Caires, Jorge A. Pérez, Frank Pfenning, and Bernardo Toninho. 2013. Behavioral Polymorphism and Parametricity in Session-Based Communication. In *ESOP*. 330–349. https://doi.org/10.1007/978-3-642-37036-6_19
- Luís Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *CONCUR (LNCS, Vol. 6269)*. 222–236. https://doi.org/10.1007/978-3-642-15375-4_16
- Marco Carbone and Søren Debois. 2010. A Graphical Approach to Progress for Structured Communication in Web Services. In *ICE (EPTCS, Vol. 38)*. 13–27. <https://doi.org/10.4204/EPTCS.38.4>
- David Castro-Perez, Francisco Ferreira, Lorenzo Gheri, and Nobuko Yoshida. 2021. Zoooid: A DSL for Certified Multiparty Computation: From Mechanised Metatheory to Certified Multiparty Processes. In *PLDI*. 237–251. <https://doi.org/10.1145/3453483.3454041>
- David Castro-Perez, Francisco Ferreira, and Nobuko Yoshida. 2020. EMTST: Engineering the Meta-theory of Session Types. In *TACAS (2) (LNCS, Vol. 12079)*. 278–285. https://doi.org/10.1007/978-3-030-45237-7_17
- Kaustuv Chaudhuri, Leonardo Lima, and Giselle Reis. 2019. Formalized Meta-Theory of Sequent Calculi for Linear Logics. *Theoretical Computer Science* 781 (2019), 24–38. <https://doi.org/10.1016/j.tcs.2019.02.023>
- Ruofei Chen and Stephanie Balzer. 2020. Ferrite: A Judgmental Embedding of Session Types in Rust. *CoRR abs/2009.13619* (2020). arXiv:2009.13619 <https://arxiv.org/abs/2009.13619>
- Luca Ciccone and Luca Padovani. 2020. A Dependently Typed Linear π -Calculus in Agda. In *PPDP*. 8:1–8:14. <https://doi.org/10.1145/3414080.3414109>
- Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. 2016. Global Progress for Dynamically Interleaved Multiparty Sessions. *MSCS* 26, 2 (2016), 238–302. <https://doi.org/10.1017/S0960129514000188>
- The Coq-std++ Team. 2021. An extended “standard library” for Coq. Available online at <https://gitlab.mpi-sws.org/iris/stdpp>.
- The Coq Team. 2021. *The Coq Proof Assistant*. <https://doi.org/10.5281/zenodo.4501022>
- Florin Craciun, Tibor Kiss, and Andreea Costea. 2015. Towards a Session Logic for Communication Protocols. In *ICECCS*. 140–149. <https://doi.org/10.1109/ICECCS.2015.33>
- Luís Cruz-Filipe, Fabrizio Montesi, and Marco Peressotti. 2021a. Certifying Choreography Compilation. In *ICTAC (LNCS, Vol. 12819)*. 115–133. https://doi.org/10.1007/978-3-030-85315-0_8
- Luís Cruz-Filipe, Fabrizio Montesi, and Marco Peressotti. 2021b. Formalising a Turing-Complete Choreographic Language in Coq. In *ITP (LIPIcs, Vol. 193)*. 15:1–15:18. <https://doi.org/10.4230/LIPIcs.ITP.2021.15>
- Ornela Dardha and Simon J. Gay. 2018. A New Linear Logic for Deadlock-Free Session-Typed Processes. In *FOSSACS (LNCS, Vol. 10803)*. 91–109. https://doi.org/10.1007/978-3-319-89366-2_5
- Ankush Das, Jan Hoffmann, and Frank Pfenning. 2018. Work Analysis with Resource-Aware Session Types. In *LICS*. 305–314. <https://doi.org/10.1145/3209108.3209146>
- Farzaneh Derakhshan, Stephanie Balzer, and Limin Jia. 2021. Session Logical Relations for Noninterference. In *LICS*. 1–14. <https://doi.org/10.1109/LICS52264.2021.9470654>
- Mariangiola Dezani-Ciancaglini, Dimitris Mostrous, Nobuko Yoshida, and Sophia Drossopoulou. 2006. Session Types for Object-Oriented Languages. In *ESOP (LNCS, Vol. 4067)*. 328–352. https://doi.org/10.1007/11785477_20
- Simon Fowler, Wen Kokke, Ornela Dardha, Sam Lindley, and J. Garrett Morris. 2021. Separating Sessions Smoothly. *CoRR abs/2105.08996* (2021). arXiv:2105.08996 <https://arxiv.org/abs/2105.08996>
- Simon Fowler, Sam Lindley, J. Garrett Morris, and Sára Decova. 2019. Exceptional Asynchronous Session Types: Session Types Without Tiers. *PACMPL* 3, POPL (2019), 28:1–28:29. <https://doi.org/10.1145/3290341>
- Adrian Francalanza, Julian Rathke, and Vladimiro Sassone. 2011. Permission-Based Separation Logic for Message-Passing Concurrency. *LMCS* 7, 3 (2011). [https://doi.org/10.2168/LMCS-7\(3:7\)2011](https://doi.org/10.2168/LMCS-7(3:7)2011)
- Simon J. Gay, Peter Thiemann, and Vasco T. Vasconcelos. 2020. Duality of Session Types: The Final Cut. In *PLACES (EPTCS, Vol. 314)*. 23–33. <https://doi.org/10.4204/EPTCS.314.3>

- Simon J. Gay and Vasco Thudichum Vasconcelos. 2010. Linear Type Theory for Asynchronous Session Types. *JFP* 20, 1 (2010), 19–50. <https://doi.org/10.1017/S0956796809990268>
- Matthew A. Goto, Radha Jagadeesan, Alan Jeffrey, Corin Pitcher, and James Riely. 2016. An Extensible Approach to Session Polymorphism. *MSCS* 26, 3 (2016), 465–509. <https://doi.org/10.1017/S0960129514000231>
- Jafar Hamin and Bart Jacobs. 2018. Deadlock-Free Monitors. In *ESOP (LNCS, Vol. 10801)*. 415–441. https://doi.org/10.1007/978-3-319-89884-1_15
- Robert Harper. 2016. *Practical Foundations for Programming Languages* (2nd ed.). Cambridge University Press. <https://doi.org/10.5555/3002812>
- Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2020. Actris: Session-Type Based Reasoning in Separation Logic. *PACMPL* 4, POPL, Article 6 (Dec. 2020), 30 pages. <https://doi.org/10.1145/3371074>
- Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2021a. Actris 2.0: Asynchronous Session-Type Based Reasoning in Separation Logic. (2021). <https://arxiv.org/abs/2010.15030v1> Manuscript.
- Jonas Kastberg Hinrichsen, Daniël Louwink, Robbert Krebbers, and Jesper Bengtson. 2021b. Machine-checked semantic session typing. In *CPP*. 178–198. <https://doi.org/10.1145/3437992.3439914>
- Kohei Honda. 1993. Types for Dyadic Interaction. In *CONCUR (LNCS, Vol. 715)*. 509–523. https://doi.org/10.1007/3-540-57208-2_35
- Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *ESOP (LNCS, Vol. 1381)*. 122–138. <https://doi.org/10.1007/BFb0053567>
- Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty Asynchronous Session Types. In *POPL*. 273–284. <https://doi.org/10.1145/1328438.1328472>
- Atsushi Igarashi and Naoki Kobayashi. 1997. Type-Based Analysis of Communication for Concurrent Programming Languages. In *SAS (LNCS, Vol. 1302)*. 187–201. <https://doi.org/10.1007/BFb0032742>
- Atsushi Igarashi and Naoki Kobayashi. 2001. A Generic Type System for the Pi-calculus. In *POPL*. 128–141. <https://doi.org/10.1145/360204.360215>
- Atsushi Igarashi and Naoki Kobayashi. 2004. A Generic Type System for the Pi-calculus. *Theoretical Computer Science* 311, 1–3 (2004), 121–163. [https://doi.org/10.1016/S0304-3975\(03\)00325-6](https://doi.org/10.1016/S0304-3975(03)00325-6)
- Atsushi Igarashi, Peter Thiemann, Vasco T. Vasconcelos, and Philip Wadler. 2017. Gradual session types. *PACMPL* 1, ICFP (2017), 38:1–38:28. <https://doi.org/10.1145/3110282>
- Keigo Imai, Nobuko Yoshida, and Shoji Yuen. 2019. Session-Ocaml: A Session-Based Library with Polarities and Lenses. *Science of Computer Programming* 172 (2019), 135–159. <https://doi.org/10.1016/j.scico.2018.08.005>
- Keigo Imai, Shoji Yuen, and Kiyoshi Agusa. 2010. Session Type Inference in Haskell. In *PLACES (EPTCS, Vol. 69)*. 74–91. <https://doi.org/10.4204/EPTCS.69.6>
- Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. 2021. Appendix and Coq mechanization of “Connectivity Graphs: A Method for Proving Deadlock Freedom Based on Separation Logic”. The most recent version is at <https://github.com/julesjacobs/cgraphs>.
- Thomas Bracht Laumann Jepsen, Philip Munksgaard, and Ken Friis Larsen. 2015. Session Types for Rust. In *WGP*. 13–22. <https://doi.org/10.1145/2808098.2808100>
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018a. RustBelt: Securing the Foundations of the Rust Programming Language. *PACMPL* 2, POPL (2018), 66:1–66:34. <https://doi.org/10.1145/3158154>
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018b. Iris From the Ground Up: A Modular Foundation for Higher-Order Concurrent Separation Logic. *JFP* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. 637–650. <https://doi.org/10.1145/2676726.2676980>
- Ioannis T. Kassios and Eleftherios Kritikos. 2013. A Discipline for Program Verification Based on Backpointers and Its Use in Observational Disjointness. In *ESOP (LNCS, Vol. 7792)*. 149–168. https://doi.org/10.1007/978-3-642-37036-6_10
- Alex C. Keizer, Henning Basold, and Jorge A. Pérez. 2021. Session Coalgebras: A Coalgebraic View on Session Types and Communication Protocols. In *ESOP (LNCS, Vol. 12648)*. 375–403. https://doi.org/10.1007/978-3-030-72019-3_14
- Naoki Kobayashi. 1997. A Partially Deadlock-Free Typed Process Calculus. In *LICS*. 128–139. <https://doi.org/10.1109/LICS.1997.614941>
- Naoki Kobayashi. 2002. A Type System for Lock-Free Processes. *I&C* 177, 2 (2002), 122–159. <https://doi.org/10.1006/inco.2002.3171>
- Naoki Kobayashi. 2006. A New Type System for Deadlock-Free Processes. In *CONCUR (LNCS, Vol. 4137)*. 233–247. https://doi.org/10.1007/11817949_16
- Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. 1999. Linearity and the pi-calculus. *TOPLAS* 21, 5 (1999), 914–947. <https://doi.org/10.1145/330249.330251>

- Naoki Kobayashi, Shin Saito, and Eijiro Sumii. 2000. An Implicitly-Typed Deadlock-Free Process Calculus. In *CONCUR (LNCS, Vol. 1877)*. 489–503. https://doi.org/10.1007/3-540-44618-4_35
- Wen Kokke. 2019. Rusty Variation: Deadlock-free Sessions with Failure in Rust. In *ICE (EPTCS, Vol. 304)*. 48–60. <https://doi.org/10.4204/EPTCS.304.4>
- Wen Kokke, Fabrizio Montesi, and Marco Peressotti. 2019. Better Late Than Never: a Fully-Abstract Semantics for Classical Processes. *PACMPL* 3, POPL (2019), 24:1–24:29. <https://doi.org/10.1145/3290337>
- Robbert Krebbers, Jacques-Henri Jourdan, Ralf Jung, Joseph Tassarotti, Jan-Oliver Kaiser, Amin Timany, Arthur Charguéraud, and Derek Dreyer. 2018. MoSeL: A General, Extensible Modal Framework for Interactive Proofs in Separation Logic. *PACMPL* 2, ICFP (2018), 77:1–77:30. <https://doi.org/10.1145/3236772>
- Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive Proofs in Higher-Order Concurrent Separation Logic. In *POPL*. 205–217. <https://doi.org/10.1145/3009837.3009855>
- Duy-Khanh Le, Wei-Ngan Chin, and Yong Meng Teo. 2013. An Expressive Framework for Verifying Deadlock Freedom. In *ATVA (LNCS, Vol. 8172)*. 287–302. https://doi.org/10.1007/978-3-319-02444-8_21
- Sam Lindley and J. Garrett Morris. 2015. A Semantics for Propositions as Sessions. In *ESOP (LNCS, Vol. 9032)*. 560–584. https://doi.org/10.1007/978-3-662-46669-8_23
- Sam Lindley and J. Garrett Morris. 2016a. Embedding Session Types in Haskell. In *Haskell Symposium*. 133–145. <https://doi.org/10.1145/2976002.2976018>
- Sam Lindley and J. Garrett Morris. 2016b. Talking Bananas: Structural Recursion For Session Types. In *ICFP*. 434–447. <https://doi.org/10.1145/2951913.2951921>
- Sam Lindley and J. Garrett Morris. 2017. Lightweight Functional Session Types. In *Behavioural Types: from Theory to Tools*. Étienne Lozes and Jules Villard. 2012. Shared Contract-Obedient Endpoints. In *ICE (EPTCS, Vol. 104)*. 17–31. <https://doi.org/10.4204/EPTCS.104.3>
- Fabrizio Montesi. 2021. Introduction to Choreographies. (2021). Accepted for publication by Cambridge University Press.
- Fabrizio Montesi and Marco Peressotti. 2018. Classical Transitions. *CoRR* abs/1803.01049 (2018). [arXiv:1803.01049](https://arxiv.org/abs/1803.01049)
- Peter W. O’Hearn and David J. Pym. 1999. The Logic Of Bunched Implications. *Bulletin of Symbolic Logic* 5, 2 (1999), 215–244. <https://doi.org/10.2307/421090>
- Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. 2001. Local Reasoning about Programs that Alter Data Structures. In *CSL (LNCS, Vol. 2142)*. 1–19. https://doi.org/10.1007/3-540-44802-0_1
- Wytse Oortwijn, Stefan Blom, and Marieke Huisman. 2016. Future-based Static Analysis of Message Passing Programs. In *PLACES (EPTCS, Vol. 211)*. 65–72. <https://doi.org/10.4204/EPTCS.211.7>
- Luca Padovani. 2017. A Simple Library Implementation of Binary Sessions. *JFP* 27 (2017), e4. <https://doi.org/10.1017/S0956796816000289>
- Jorge A. Pérez, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2014. Linear Logical Relations and Observational Equivalences for Session-Based Concurrency. *I&C* 239 (2014), 254–302. <https://doi.org/10.1016/j.ic.2014.08.001>
- Frank Pfenning and Dennis Griffith. 2015. Polarized Substructural Session Types. In *FoSSaCS (LNCS, Vol. 9034)*. 3–22. https://doi.org/10.1007/978-3-662-46678-0_1
- Benjamin C. Pierce. 2002. *Types and Programming Languages* (1st ed.). The MIT Press. <https://doi.org/10.5555/509043>
- Riccardo Pucella and Jesse A. Tov. 2008. Haskell Session Types with (Almost) No Class. In *Haskell Symposium*. 25–36. <https://doi.org/10.1145/1411286.1411290>
- Zesen Qian, G. A. Kavvos, and Lars Birkedal. 2021. Client-Server Sessions in Linear Logic. *PACMPL* 5, ICFP (2021), 1–31. <https://doi.org/10.1145/3473567>
- Jason Reed. 2009a. *A Hybrid Logical Framework*. Ph.D. Dissertation. Carnegie Mellon University.
- Jason Reed. 2009b. A Judgmental Deconstruction of Modal Logic. (2009). <http://www.cs.cmu.edu/~jcreed/papers/jdml.pdf>
- Unpublished manuscript.
- Pedro Rocha and Luís Caires. 2021. *Propositions-as-Types and Shared State*. Technical Report. NOVA LINCS.
- Arjen Rouvoet, Robbert Krebbers, and Eelco Visser. 2021. Intrinsically Typed Compilation With Nameless Labels. *PACMPL* 5, POPL (2021), 1–28. <https://doi.org/10.1145/3434303>
- Arjen Rouvoet, Casper Bach Poulsen, Robbert Krebbers, and Eelco Visser. 2020. Intrinsically-Typed Definitional Interpreters for Linear, Session-Typed Languages. In *CPP*. 284–298. <https://doi.org/10.1145/3372885.3373818>
- Alceste Scalas and Nobuko Yoshida. 2016. Lightweight Session Programming in Scala. In *ECOOP (LIPIcs, 56)*. 21:1–21:28. <https://doi.org/10.4230/LIPIcs.ECOOP.2016.21>
- Mathieu Sozeau. 2009. A New Look at Generalized Rewriting in Type Theory. *JFR* 2, 1 (2009), 41–62. <https://doi.org/10.6092/issn.1972-5787/1574>
- Joseph Tassarotti, Ralf Jung, and Robert Harper. 2017. A Higher-Order Logic for Concurrent Termination-Preserving Refinement. In *ESOP (LNCS, Vol. 10201)*. 909–936. https://doi.org/10.1007/978-3-662-54434-1_34

- Peter Thiemann. 2019. Intrinsically-Typed Mechanized Semantics for Session Types. In *PPDP*. 19:1–19:15. <https://doi.org/10.1145/3354166.3354184>
- Bernardo Toninho. 2015. *A Logical Foundation for Session-Based Concurrent Computation*. Ph.D. Dissertation. Carnegie Mellon University and New University of Lisbon.
- Bernardo Toninho, Luís Caires, and Frank Pfenning. 2013. Higher-Order Processes, Functions, and Sessions: A Monadic Integration. In *ESOP (LNCS, Vol. 7792)*. 350–369. https://doi.org/10.1007/978-3-642-37036-6_20
- Vasco T. Vasconcelos. 2012. Fundamentals of Session Types. *I&C* 217 (2012), 52–70. <https://doi.org/10.1016/j.ic.2012.05.002>
- Philip Wadler. 2012. Propositions as Sessions. In *ICFP*. 273–286. <https://doi.org/10.1145/2364527.2364568>
- Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *I&C* 115, 1 (1994), 38–94. <https://doi.org/10.1006/inco.1994.1093>
- Dan Zhang, Dragan Bosnacki, Mark van den Brand, Cornelis Huizing, Bart Jacobs, Ruurd Kuiper, and Anton Wijs. 2016. Verifying Atomicity Preservation and Deadlock Freedom of a Generic Shared Variable Mechanism Used in Model-To-Code Transformations. In *MODELSWARD (CCIS, Vol. 692)*. 249–273. https://doi.org/10.1007/978-3-319-66302-9_13