



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


Jules Jacobs¹



Robbert Krebbers¹



Stephanie Balzer²



¹Radboud University Nijmegen, The Netherlands

²Carnegie Mellon University, USA

Session types

Message passing concurrency with first-class channels (Honda [1993])

$c : !Nat. ?Bool. !(?String. !Nat. End). End$

Session types

Message passing concurrency with first-class channels (Honda [1993])

$$c : !Nat. ?Bool. !(?String. !Nat. End). End$$
$$\Downarrow \text{dual}$$
$$c' : ?Nat. !Bool. (?String. !Nat. End). End$$

GV: functional programming with session types

(Gay and Vasconcelos [2010], Wadler [2012])

$$\text{fork} : (s \xrightarrow{\text{lin}} 1) \rightarrow \bar{s}$$
$$\text{send} : (!t. s) \times t \xrightarrow{\text{lin}} s$$
$$\text{close} : \text{End} \xrightarrow{\text{lin}} 1$$
$$\text{receive} : ?t. s \xrightarrow{\text{lin}} s \times t$$
$$\text{let } c = \text{fork}(\lambda c'. \dots \text{receive}(c') \dots) \text{ in } \text{send}(c, 23) \dots$$

What makes session types interesting

Linear session types: cannot copy or delete a channel reference before you are done

What makes session types interesting

Linear session types: cannot copy or delete a channel reference before you are done

- ▶ Required for **type safety** (mechanized by Castro-Perez et al. [2020], Ciccone and Padovani [2020], Goto et al. [2016], Hinrichsen et al. [2021], Rouvoet et al. [2020], Thiemann [2019], ...)

What makes session types interesting

Linear session types: cannot copy or delete a channel reference before you are done

- ▶ Required for **type safety** (mechanized by Castro-Perez et al. [2020], Ciccone and Padovani [2020], Goto et al. [2016], Hinrichsen et al. [2021], Rouvoet et al. [2020], Thiemann [2019], ...)
- ▶ But also guarantees **deadlock freedom, global progress** (well-known property, but not yet mechanized for first-class channels, i.e. dynamically allocated and higher order)

Why session types give deadlock freedom

Two owners per channel

- ▶ Duality of channel types \rightarrow no simple deadlocks
- ▶ Linear typing maintains acyclicity of ownership structure \rightarrow no cyclic deadlocks

Why session types give deadlock freedom

Two owners per channel

- ▶ Duality of channel types → no simple deadlocks
- ▶ Linear typing maintains acyclicity of ownership structure → no cyclic deadlocks

Even when channels are asynchronous and first-class values:

- ▶ dynamically created
- ▶ sent as messages over channels
- ▶ stored in data structures
- ▶ captured by closures

Why session types give deadlock freedom

Two owners per channel

- ▶ Duality of channel types → no simple deadlocks
- ▶ Linear typing maintains acyclicity of ownership structure → no cyclic deadlocks

Even when channels are asynchronous and first-class values:

- ▶ dynamically created
- ▶ sent as messages over channels
- ▶ stored in data structures
- ▶ captured by closures

Difficult to reason about typing & graph structure simultaneously

Contribution: connectivity graph proof method

This work: **connectivity graphs**

- ▶ Method for factoring out graph reasoning from reasoning about typing
- ▶ Mechanized in the Coq proof assistant
- ▶ Applied to prove deadlock freedom for feature-rich session-typed language
- ▶ Abstract representation of run-time configuration

Contribution: connectivity graph proof method

This work: **connectivity graphs**

- ▶ Method for factoring out graph reasoning from reasoning about typing
- ▶ Mechanized in the Coq proof assistant
- ▶ Applied to prove deadlock freedom for feature-rich session-typed language
- ▶ Abstract representation of run-time configuration

Run-time configuration ρ

Threads: $\{T_1 \mapsto e_1, \dots, T_6 \mapsto e_6\}$

Channels: $\{C_1 \mapsto \text{buf}_1, \dots, C_5 \mapsto \text{buf}_5\}$

Contribution: connectivity graph proof method

This work: **connectivity graphs**


- ▶ Method for factoring out graph reasoning from reasoning about typing
- ▶ Mechanized in the Coq proof assistant
- ▶ Applied to prove deadlock freedom for feature-rich session-typed language
- ▶ Abstract representation of run-time configuration

Run-time configuration ρ

Threads: $\{T_1 \mapsto e_1, \dots, T_6 \mapsto e_6\}$

Channels: $\{C_1 \mapsto \text{buf}_1, \dots, C_5 \mapsto \text{buf}_5\}$

Connectivity graph G



Contribution: connectivity graph proof method

This work: **connectivity graphs**

- ▶ Method for factoring out graph reasoning from reasoning about typing
- ▶ Mechanized in the Coq proof assistant
- ▶ Applied to prove deadlock freedom for feature-rich session-typed language
- ▶ Abstract representation of run-time configuration


Run-time configuration ρ

Threads: $\{T_1 \mapsto e_1, \dots, T_6 \mapsto e_6\}$

Channels: $\{C_1 \mapsto \text{buf}_1, \dots, C_5 \mapsto \text{buf}_5\}$

$G \vDash \rho$


Connectivity graph G



Connectivity graph proof based on progress and preservation

ρ_1


Connectivity graph proof based on progress and preservation



Connectivity graph proof based on progress and preservation




Connectivity graph proof based on progress and preservation




Connectivity graph proof based on progress and preservation




Connectivity graph proof based on progress and preservation




Connectivity graph proof based on progress and preservation




Connectivity graph proof based on progress and preservation



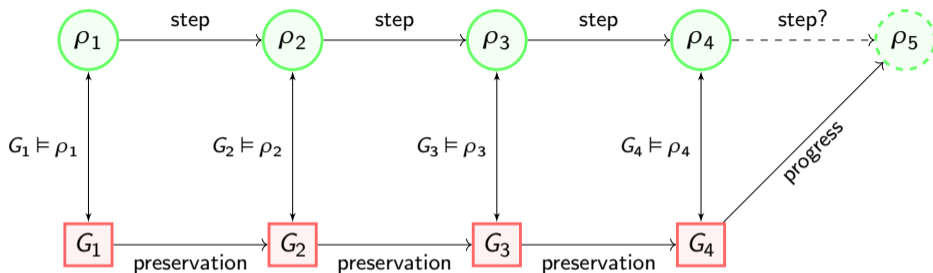
Connectivity graph proof based on progress and preservation



Connectivity graph proof based on progress and preservation



Connectivity graph proof based on progress and preservation



Connectivity graph framework:

- ▶ $Cgraph(V, L)$ data type for acyclic labeled graphs
- ▶ Generic construction for $G \vDash \rho$
 - ▶ Parameterized by local separation logic predicate $P_\rho(v)$ for each vertex $v \in G$
- ▶ Preservation: graph transformations in separation logic
- ▶ Progress: waiting induction principle for $Cgraph(V, L)$

All generic over vertices V and labels L

Linear heap typing in separation logic: (cf. Rouvoet [2020]'s definitional interpreters)

$$\frac{\Sigma_1 \vdash e_1 : \tau_1 \quad \Sigma_2 \vdash e_2 : \tau_2 \quad \Sigma_1 \cap \Sigma_2 = \emptyset}{\Sigma_1 \cup \Sigma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2}$$

Linear heap typing in separation logic: (cf. Rouvoet [2020]'s definitional interpreters)

$$\frac{\Sigma_1 \vdash e_1 : \tau_1 \quad \Sigma_2 \vdash e_2 : \tau_2 \quad \Sigma_1 \cap \Sigma_2 = \emptyset}{\Sigma_1 \cup \Sigma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2} \quad \Rightarrow \quad \frac{e_1 : \tau_1 * e_2 : \tau_2}{(e_1, e_2) : \tau_1 \times \tau_2} *$$

Linear heap typing in separation logic: (cf. Rouvoet [2020]'s definitional interpreters)

$$\frac{\Sigma_1 \vdash e_1 : \tau_1 \quad \Sigma_2 \vdash e_2 : \tau_2 \quad \Sigma_1 \cap \Sigma_2 = \emptyset}{\Sigma_1 \cup \Sigma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2}$$

$$\frac{\Sigma = \{\text{Chan}(a) \mapsto (t, s)\}}{\Sigma \vdash \#a_t : s}$$

$$\Rightarrow \frac{e_1 : \tau_1 * e_2 : \tau_2}{(e_1, e_2) : \tau_1 \times \tau_2} *$$

$$\Rightarrow \frac{\text{own}(\text{Chan}(a) \mapsto (t, s))}{\#a_t : s} *$$

Linear heap typing in separation logic: (cf. Rouvoet [2020]'s definitional interpreters)

$$\frac{\Sigma_1 \vdash e_1 : \tau_1 \quad \Sigma_2 \vdash e_2 : \tau_2 \quad \Sigma_1 \cap \Sigma_2 = \emptyset}{\Sigma_1 \cup \Sigma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2} \quad \Rightarrow \quad \frac{e_1 : \tau_1 * e_2 : \tau_2}{(e_1, e_2) : \tau_1 \times \tau_2} *$$
$$\frac{\Sigma = \{\text{Chan}(a) \mapsto (t, s)\}}{\Sigma \vdash \#a_t : s} \quad \Rightarrow \quad \frac{\text{own}(\text{Chan}(a) \mapsto (t, s))}{\#a_t : s} *$$

For vertex v in the graph, separation logic resource $\Sigma = \text{OutEdges}(v)$

Linear heap typing in separation logic: (cf. Rouvoet [2020]'s definitional interpreters)

$$\begin{array}{c}
 \frac{\Sigma_1 \vdash e_1 : \tau_1 \quad \Sigma_2 \vdash e_2 : \tau_2 \quad \Sigma_1 \cap \Sigma_2 = \emptyset}{\Sigma_1 \cup \Sigma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2} \quad \Rightarrow \quad \frac{e_1 : \tau_1 * e_2 : \tau_2}{(e_1, e_2) : \tau_1 \times \tau_2} * \\
 \\
 \frac{\Sigma = \{\text{Chan}(a) \mapsto (t, s)\}}{\Sigma \vdash \#a_t : s} \quad \Rightarrow \quad \frac{\text{own}(\text{Chan}(a) \mapsto (t, s))}{\#a_t : s} *
 \end{array}$$

For vertex v in the graph, separation logic resource $\Sigma = \text{OutEdges}(v)$

Lemmas in separation logic:

$$(\Sigma \vdash K[e] : B) \iff \exists A, \Sigma_1, \Sigma_2. (\Sigma_1 \cap \Sigma_2 = \emptyset) \wedge (\Sigma = \Sigma_1 \cup \Sigma_2) \wedge (\Sigma_1 \vdash e : A) \wedge \\
 \forall e', \Sigma_3. (\Sigma_2 \cap \Sigma_3 = \emptyset) \wedge (\Sigma_2 \vdash e' : A) \rightarrow (\Sigma_2 \cup \Sigma_3 \vdash K[e'] : B)$$

Linear heap typing in separation logic: (cf. Rouvoet [2020]'s definitional interpreters)

$$\frac{\Sigma_1 \vdash e_1 : \tau_1 \quad \Sigma_2 \vdash e_2 : \tau_2 \quad \Sigma_1 \cap \Sigma_2 = \emptyset}{\Sigma_1 \cup \Sigma_2 \vdash (e_1, e_2) : \tau_1 \times \tau_2} \quad \Rightarrow \quad \frac{e_1 : \tau_1 * e_2 : \tau_2}{(e_1, e_2) : \tau_1 \times \tau_2} *$$

$$\frac{\Sigma = \{\text{Chan}(a) \mapsto (t, s)\}}{\Sigma \vdash \#a_t : s} \quad \Rightarrow \quad \frac{\text{own}(\text{Chan}(a) \mapsto (t, s))}{\#a_t : s} *$$

For vertex v in the graph, separation logic resource $\Sigma = \text{OutEdges}(v)$

Lemmas in separation logic:

$$(\Sigma \vdash K[e] : B) \iff \exists A, \Sigma_1, \Sigma_2. (\Sigma_1 \cap \Sigma_2 = \emptyset) \wedge (\Sigma = \Sigma_1 \cup \Sigma_2) \wedge (\Sigma_1 \vdash e : A) \wedge$$


$$\forall e', \Sigma_3. (\Sigma_2 \cap \Sigma_3 = \emptyset) \wedge (\Sigma_2 \vdash e' : A) \rightarrow (\Sigma_2 \cup \Sigma_3 \vdash K[e'] : B)$$

$$\Rightarrow$$

$$(K[e] : B) \dashv\vdash \exists A. (e : A) * \forall e'. (e' : A) \multimap (K[e'] : B)$$

We use the Iris proof mode to reason in separation logic (Krebbers et al. [2017])


Preservation via local graph transformations



Preserves:

- ▶ Acyclicity
- ▶ Local predicates $P_\rho(v)$ used for $G \models \rho$

Preservation via local graph transformations



Preserves:

- ▶ Acyclicity
- ▶ Local predicates $P_\rho(v)$ used for $G \models \rho$

All in separation logic:

$$P_\rho(v_1) * (\text{own}(v_2 \mapsto \ell) \multimap P_\rho(v_2))$$


$$\vdash$$

$$(\text{own}(v_2 \mapsto \ell') \multimap P_{\rho'}(v_1)) * P_{\rho'}(v_2)$$

Explained in our paper!


Progress via waiting induction

Connectivity graph with *waiting dependencies* (▶)
derived from run-time configuration



Progress via waiting induction

Connectivity graph with *waiting dependencies* (▶)
derived from run-time configuration



Lemma (Waiting induction)

Let $R(v, w)$ be any relation on the vertices. To prove $P(v)$, we may assume $P(w)$ for all w such that $v \rightarrow w$ and $R(v, w)$, or $w \rightarrow v$ and $\neg R(w, v)$

Our language

Functional language + session-typed channels (extension of Wadler [2012]'s GV)

Unrestricted and linear types

- ▶ Unrestricted: numbers, sums, products, unrestricted function type (\rightarrow)
- ▶ Linear: channels, sums, products, linear function type (\multimap)

General recursive types:

- ▶ Recursive session types, including through the message (example: $\mu X. !X.End$)
- ▶ Algebraic data types using recursion + sums + products
- ▶ Recursive types mechanized using coinduction (Gay et al. [2020])

Stronger deadlock and leak freedom result

Global progress is the standard notion that people use

Our POPL reviewers: Can your method prove something stronger?

Stronger deadlock and leak freedom result

Global progress is the standard notion that people use

Our POPL reviewers: Can your method prove something stronger?

Partial deadlock: a set S of threads and channels such that:

1. All threads in S are blocked on a channel in S
2. No references to channels in S from outside S

Stronger deadlock and leak freedom result

Global progress is the standard notion that people use

Our POPL reviewers: Can your method prove something stronger?

Partial deadlock: a set S of threads and channels such that:

1. All threads in S are blocked on a channel in S
2. No references to channels in S from outside S

Strong reachability:

1. A channel is reachable if it is referenced by a reachable channel or thread
2. A thread is reachable if it can step, or is blocked on a reachable channel

Stronger deadlock and leak freedom result

Global progress is the standard notion that people use

Our POPL reviewers: Can your method prove something stronger?

Partial deadlock: a set S of threads and channels such that:

1. All threads in S are blocked on a channel in S
2. No references to channels in S from outside S

Strong reachability:

1. A channel is reachable if it is referenced by a reachable channel or thread
2. A thread is reachable if it can step, or is blocked on a reachable channel

Lemma. All threads and channels are reachable \iff no partial deadlock

Stronger deadlock and leak freedom result

Global progress is the standard notion that people use

Our POPL reviewers: Can your method prove something stronger?

Partial deadlock: a set S of threads and channels such that:

1. All threads in S are blocked on a channel in S
2. No references to channels in S from outside S

Strong reachability:

1. A channel is reachable if it is referenced by a reachable channel or thread
2. A thread is reachable if it can step, or is blocked on a reachable channel

Lemma. All threads and channels are reachable \iff no partial deadlock

Lemma. Any thread or channel is reachable \implies global progress

Stronger deadlock and leak freedom result

Global progress is the standard notion that people use

Our POPL reviewers: Can your method prove something stronger?

Partial deadlock: a set S of threads and channels such that:

1. All threads in S are blocked on a channel in S
2. No references to channels in S from outside S

Strong reachability:

1. A channel is reachable if it is referenced by a reachable channel or thread
2. A thread is reachable if it can step, or is blocked on a reachable channel

Lemma. All threads and channels are reachable \iff no partial deadlock

Lemma. Any thread or channel is reachable \implies global progress

Theorem. For well-typed initial programs, no partial deadlock occurs

Mechanization

Mechanization in Coq:

- ▶ Generic $Cgraph(V, L)$ library: 4999 LOC
- ▶ GV language definition: 451 LOC
- ▶ Language specific deadlock and leak freedom proof: 1688 LOC

<https://github.com/julesjacobs/cgraphs>



Mechanization

Mechanization in Coq:

- ▶ Generic $Cgraph(V, L)$ library: 4999 LOC
- ▶ GV language definition: 451 LOC
- ▶ Language specific deadlock and leak freedom proof: 1688 LOC

<https://github.com/julesjacobs/cgraphs>

Initial direct attempt: **proofs goals got too complex.**

Graph reasoning intertwined with language specifics.

Encapsulating the graph reasoning made it manageable.



Multiparty session types

- ▶ Multiparty session types $>$ binary session types?

Multiparty session types

- ▶ Multiparty session types > binary session types? **No!**

Multiparty session types

- ▶ Multiparty session types $>$ binary session types? **No!**
- ▶ Multiparty literature has focused on single-session deadlock freedom

Multiparty session types

- ▶ Multiparty session types \succ binary session types? **No!**
- ▶ Multiparty literature has focused on single-session deadlock freedom
- ▶ **Exception:** Coppo et al. [2016], order-based deadlock freedom
- ▶ **But we want:** GV with **multiparty** and acyclicity-based deadlock freedom
 - ▶ Deadlock free dynamic session creation & thread spawning
 - ▶ Unrestricted delegation
 - ▶ Free use of multiparty endpoints in data structures & closures
 - ▶ Deadlock freedom by linear typing (no other restrictions)
 - ▶ True global progress
 - ▶ Strictly more general than GV

Multiparty session types

- ▶ Multiparty session types \gt binary session types? **No!**
- ▶ Multiparty literature has focused on single-session deadlock freedom
- ▶ **Exception:** Coppo et al. [2016], order-based deadlock freedom
- ▶ **But we want:** GV with **multiparty** and acyclicity-based deadlock freedom
 - ▶ Deadlock free dynamic session creation & thread spawning
 - ▶ Unrestricted delegation
 - ▶ Free use of multiparty endpoints in data structures & closures
 - ▶ Deadlock freedom by linear typing (no other restrictions)
 - ▶ True global progress
 - ▶ Strictly more general than GV
- ▶ **MPGV:** multiparty session types \gt binary session types!

Multiparty session types

- ▶ Multiparty session types \succ binary session types? **No!**
- ▶ Multiparty literature has focused on single-session deadlock freedom
- ▶ **Exception:** Coppo et al. [2016], order-based deadlock freedom
- ▶ **But we want:** GV with **multiparty** and acyclicity-based deadlock freedom
 - ▶ Deadlock free dynamic session creation & thread spawning
 - ▶ Unrestricted delegation
 - ▶ Free use of multiparty endpoints in data structures & closures
 - ▶ Deadlock freedom by linear typing (no other restrictions)
 - ▶ True global progress
 - ▶ Strictly more general than GV
- ▶ **MPGV:** multiparty session types \succ binary session types!
- ▶ Global progress mechanized in Coq (10400 LOC, 638 lemmas)

Questions?

julesjacobs@gmail.com

- D. Castro-Perez, F. Ferreira, and N. Yoshida. EMTST: engineering the meta-theory of session types. In *TACAS (2)*, volume 12079 of *LNCS*, pages 278–285, 2020. doi: 10.1007/978-3-030-45237-7_17. URL https://doi.org/10.1007/978-3-030-45237-7_17.
- L. Ciccone and L. Padovani. A dependently typed linear π -calculus in agda. In *PPDP*, pages 8:1–8:14, 2020. doi: 10.1145/3414080.3414109. URL <https://doi.org/10.1145/3414080.3414109>.
- M. Coppo, M. Dezani-Ciancaglini, N. Yoshida, and L. Padovani. Global progress for dynamically interleaved multiparty sessions. *MSCS*, 26(2):238–302, 2016. doi: 10.1017/S0960129514000188. URL <https://doi.org/10.1017/S0960129514000188>.
- S. J. Gay and V. T. Vasconcelos. Linear type theory for asynchronous session types. *JFP*, 20(1):19–50, 2010. doi: 10.1017/S0956796809990268. URL <https://doi.org/10.1017/S0956796809990268>.
- S. J. Gay, P. Thiemann, and V. T. Vasconcelos. Duality of session types: The final cut. In *PLACES@ETAPS*, volume 314 of *EPTCS*, pages 23–33, 2020. doi: 10.4204/EPTCS.314.3. URL <https://doi.org/10.4204/EPTCS.314.3>.

- M. A. Goto, R. Jagadeesan, A. Jeffrey, C. Pitcher, and J. Riely. An extensible approach to session polymorphism. *MSCS*, 26(3):465–509, 2016. doi: 10.1017/S0960129514000231. URL <https://doi.org/10.1017/S0960129514000231>.
- J. K. Hinrichsen, D. Louwrik, R. Krebbers, and J. Bengtson. Machine-checked semantic session typing. In *CPP*, pages 178–198, 2021. doi: 10.1145/3437992.3439914. URL <https://doi.org/10.1145/3437992.3439914>.
- K. Honda. Types for dyadic interaction. In *CONCUR*, volume 715 of *LNCS*, pages 509–523, 1993. doi: 10.1007/3-540-57208-2_35. URL https://doi.org/10.1007/3-540-57208-2_35.
- R. Krebbers, A. Timany, and L. Birkedal. Interactive proofs in higher-order concurrent separation logic. In *POPL*, pages 205–217, 2017. doi: 10.1145/3009837.3009855. URL <https://doi.org/10.1145/3009837.3009855>.
- A. Rouvoet, C. Bach Poulsen, R. Krebbers, and E. Visser. Intrinsically-typed definitional interpreters for linear, session-typed languages. In *CPP*, 2020. ISBN 9781450370974. doi: 10.1145/3372885.3373818. URL <https://doi.org/10.1145/3372885.3373818>.

- P. Thiemann. Intrinsically-typed mechanized semantics for session types. In *PPDP*, 2019. ISBN 9781450372497. doi: 10.1145/3354166.3354184. URL <https://doi.org/10.1145/3354166.3354184>.
- P. Wadler. Propositions as sessions. In *ICFP*, pages 273–286, 2012. doi: 10.1145/2364527.2364568. URL <https://doi.org/10.1145/2364527.2364568>.