


mCRL2 and Multiparty Session Types




mCRL2
analysing system behaviour



Nobuko Yoshida (Oxford)

Multiparty Session Types

Specification-Guided Concurrent and Distributed Programming



1916-1975

- **Christopher Strachey** (sequential computation)
 - ▶ **Types** = abstract and digest computation (data types, polymorphism)
 - ▶ **Structured programming = High-level programming**
- **Session types** (concurrency & communication)
 - ▶ Structured programming = **protocols**

mCRL2 and Multiparty Session Types

- How can we integrate local verifications (by mCRL2) and global protocol conformance?
- What are possible extensions of local verifications of session types?
 1. probabilities
 2. counter-example guided verification
 3. bisimulation checking
 4. asynchrony
 5. mCRL2 graphics tool

Communications are Ubiquitous

- Increasingly, **communications** are the way to organise software and systems.
- Industry trend – programming languages with **explicit message-passing primitives**.



microservices



Problems: Ambiguity

- Protocol descriptions are **ambiguous**
- **SMTP: simple mail transfer protocol**
 - They are written in English, often very long



RFC 821

August 1982
Simple Mail Transfer Protocol

TABLE OF CONTENTS

<u>1.</u>	INTRODUCTION	<u>1</u>
<u>2.</u>	THE SMTP MODEL	<u>2</u>
<u>3.</u>	THE SMTP PROCEDURE	<u>4</u>
<u>3.1.</u>	Mail	<u>4</u>
<u>3.2.</u>	Forwarding	<u>7</u>
<u>3.3.</u>	Verifying and Expanding	<u>8</u>
<u>3.4.</u>	Sending and Mailing	<u>11</u>
<u>3.5.</u>	Opening and Closing	<u>13</u>
<u>3.6.</u>	Relaying	<u>14</u>
<u>3.7.</u>	Domains	<u>17</u>
<u>3.8.</u>	Changing Roles	<u>18</u>
<u>4.</u>	THE SMTP SPECIFICATIONS	<u>19</u>
<u>4.1.</u>	SMTP Commands	<u>19</u>
<u>4.1.1.</u>	Command Semantics	<u>19</u>
<u>4.1.2.</u>	Command Syntax	<u>27</u>
<u>4.2.</u>	SMTP Replies	<u>34</u>
<u>4.2.1.</u>	Reply Codes by Function Group	<u>35</u>
<u>4.2.2.</u>	Reply Codes in Numeric Order	<u>36</u>
<u>4.3.</u>	Sequencing of Commands and Replies	<u>37</u>
<u>4.4.</u>	State Diagrams	<u>39</u>
<u>4.5.</u>	Details	<u>41</u>
<u>4.5.1.</u>	Minimum Implementation	<u>41</u>
<u>4.5.2.</u>	Transparency	<u>41</u>
<u>4.5.3.</u>	Sizes	<u>42</u>

Problems: Ambiguity

- Protocol descriptions are **ambiguous**
- **SMTP: simple mail transfer protocol**
 - They are written in English, often very long



3.1. MAIL

There are three steps to SMTP mail transactions. The transaction is started with a MAIL command which gives the sender identification. A series of one or more RCPT commands follows giving the receiver information. Then a DATA command gives the mail data. And finally, the end of mail data indicator confirms the transaction.

The first step in the procedure is the MAIL command. The <reverse-path> contains the source mailbox.

```
MAIL <SP> FROM:<reverse-path> <CRLF>
```

This command tells the SMTP-receiver that a new mail transaction is starting and to reset all its state tables and buffers, including any recipients or mail data. It gives the reverse-path which can be used to report errors. If accepted, the receiver-SMTP returns a 250 OK reply.

The <reverse-path> can contain more than just a mailbox. The <reverse-path> is a reverse source routing list of hosts and source mailbox. The first host in the <reverse-path> should be the host sending this command.

The second step in the procedure is the RCPT command.

```
RCPT <SP> TO:<forward-path> <CRLF>
```

This command gives a forward-path identifying one recipient. If accepted, the receiver-SMTP returns a 250 OK reply, and stores the forward-path. If the recipient is unknown the receiver-SMTP returns a 550 Failure reply. This second step of the procedure can be repeated any number of times.

Problems: Concurrency Bugs

- Communications increase **concurrency bugs**
 - Survey of 4k users [golang.org]
 - Analysis of 6 large software systems [ASPLOS 19]



GO

Google (2009)



The Go Gopher

CSP_{80'}

*Do not communicate by sharing memory;
share memory by communicating*

– Go Philosophy

Problems: Concurrency Bugs

- Communications increase **concurrency bugs**
 - Survey of 4K users [golang.org]
 - Analysis of 6 large software systems [ASPLOS 19]
Uber code base [PLDI 2022]

More than a half of concurrency bugs in Go are caused by communications.

deadlock

channel errors



The Go Gopher

Problems: Concurrency Bugs

- Communications increase **concurrency bugs**
 - Survey of 4k users [golang.org]
 - Analysis of 6 large software systems [ASPLOS 19]

More than a half of concurrency bugs in Go are caused by communications.

Session Types

- Prevent concurrency bugs.
- Can abstract, implement and manage communications as **Protocols**.
- **Clean, Cheap** and **Retrofittable**.



Why Session Types, Why Now?


Significant academic and industry interests via fundamental breakthroughs

Milner,
Honda, NY



Binary Session Types

ESOP'98

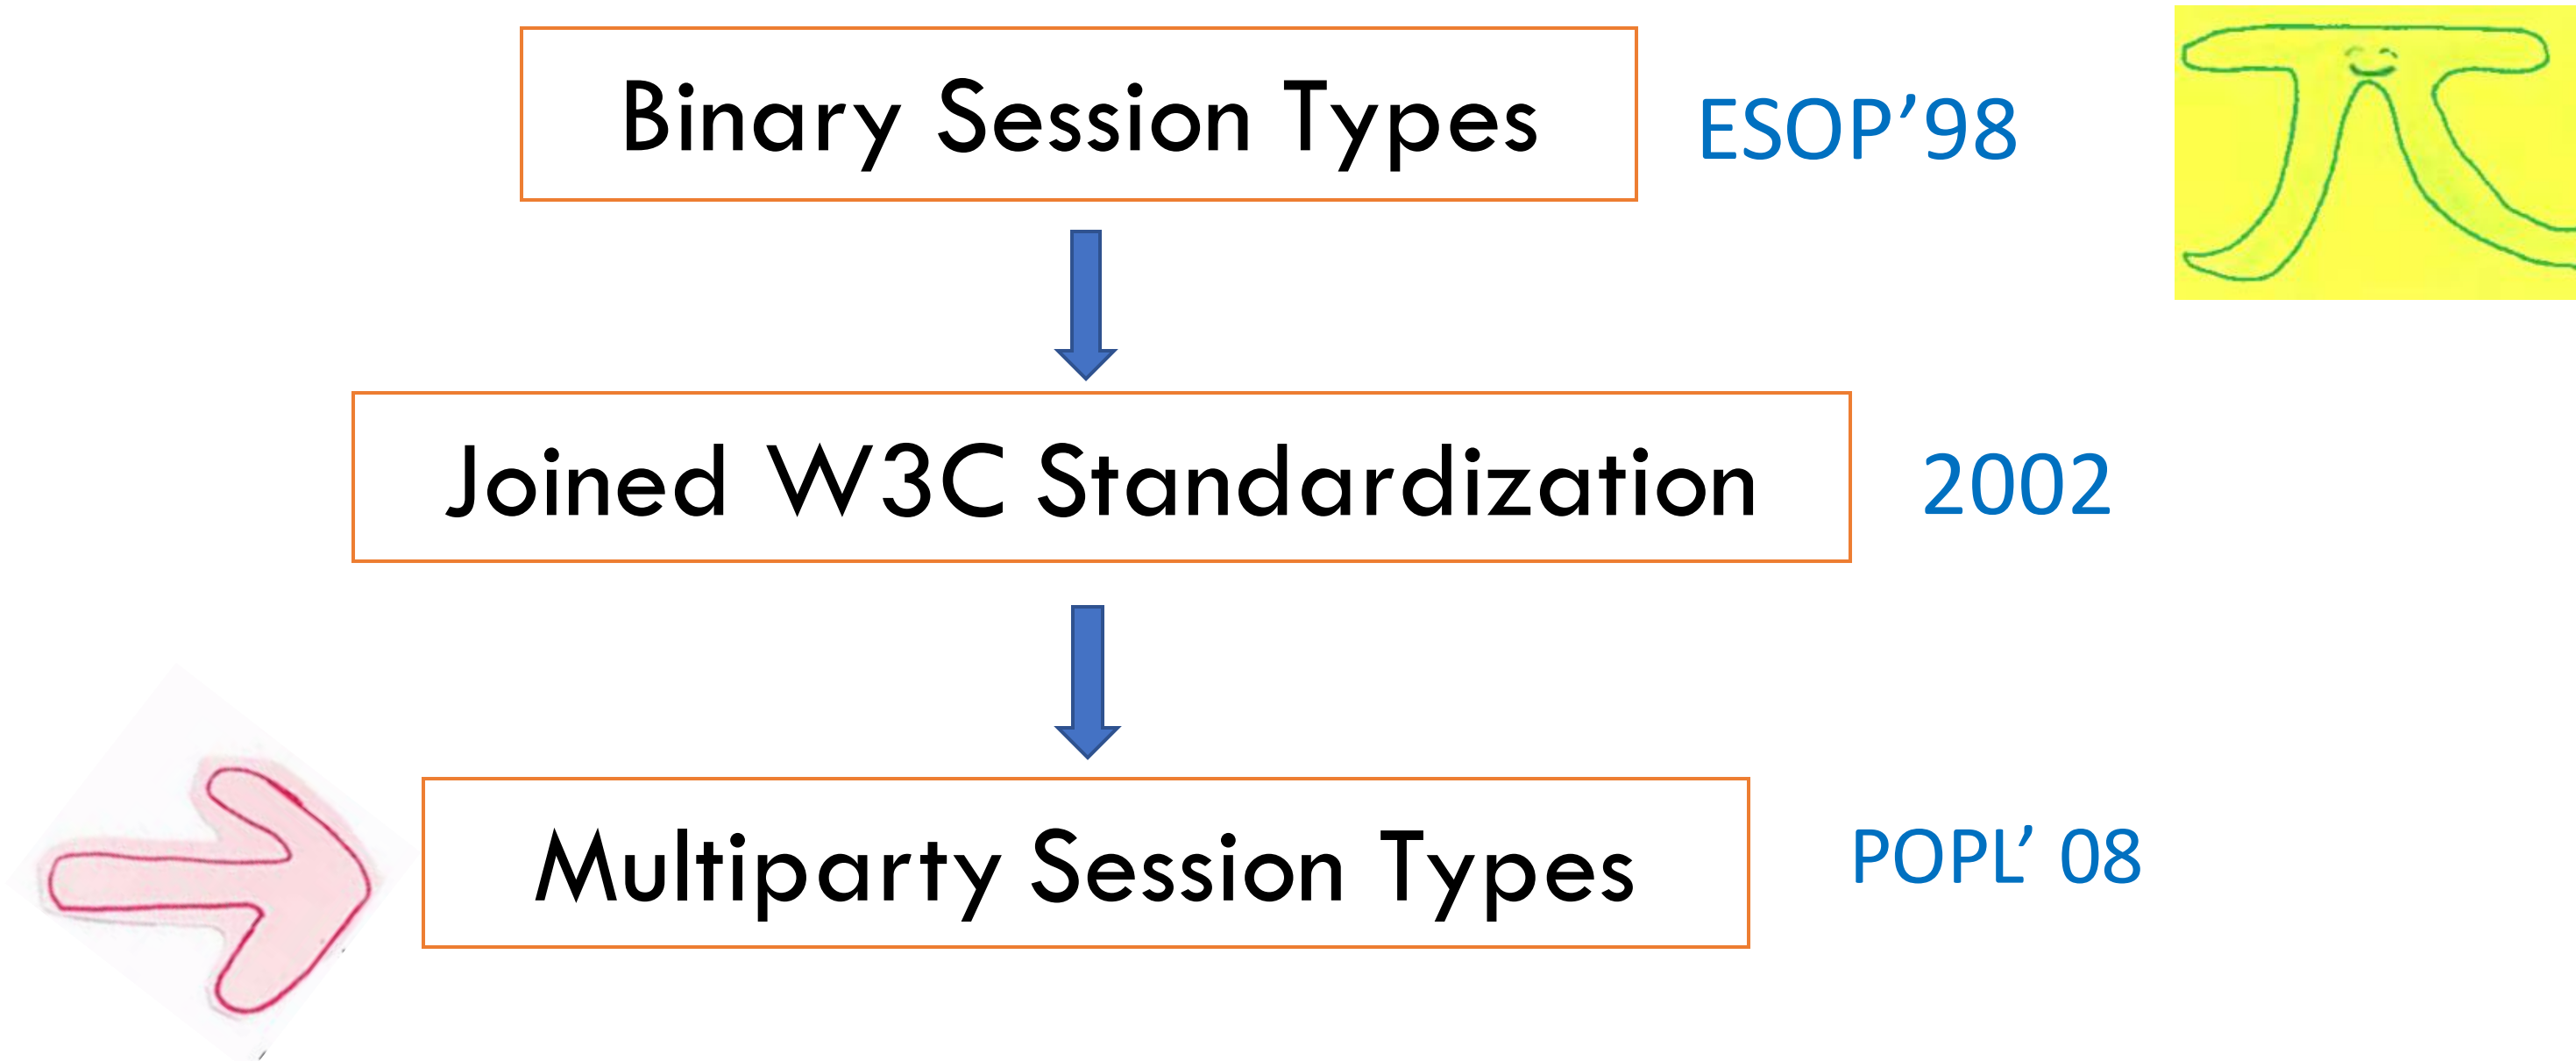


Joined W3C Standardization

2002

Why Session Types, Why Now?

Significant academic and industry interests via fundamental breakthroughs


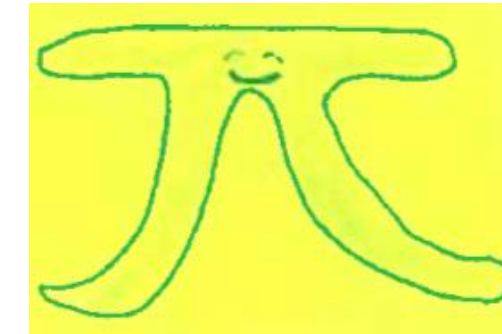


Why Session Types, Why Now?

Significant academic and industry interests via fundamental breakthroughs

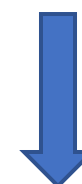
Binary Session Types

ESOP'98



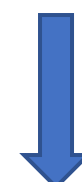
Joined W3C Standardization

2002



Multiparty Session Types

POPL' 08

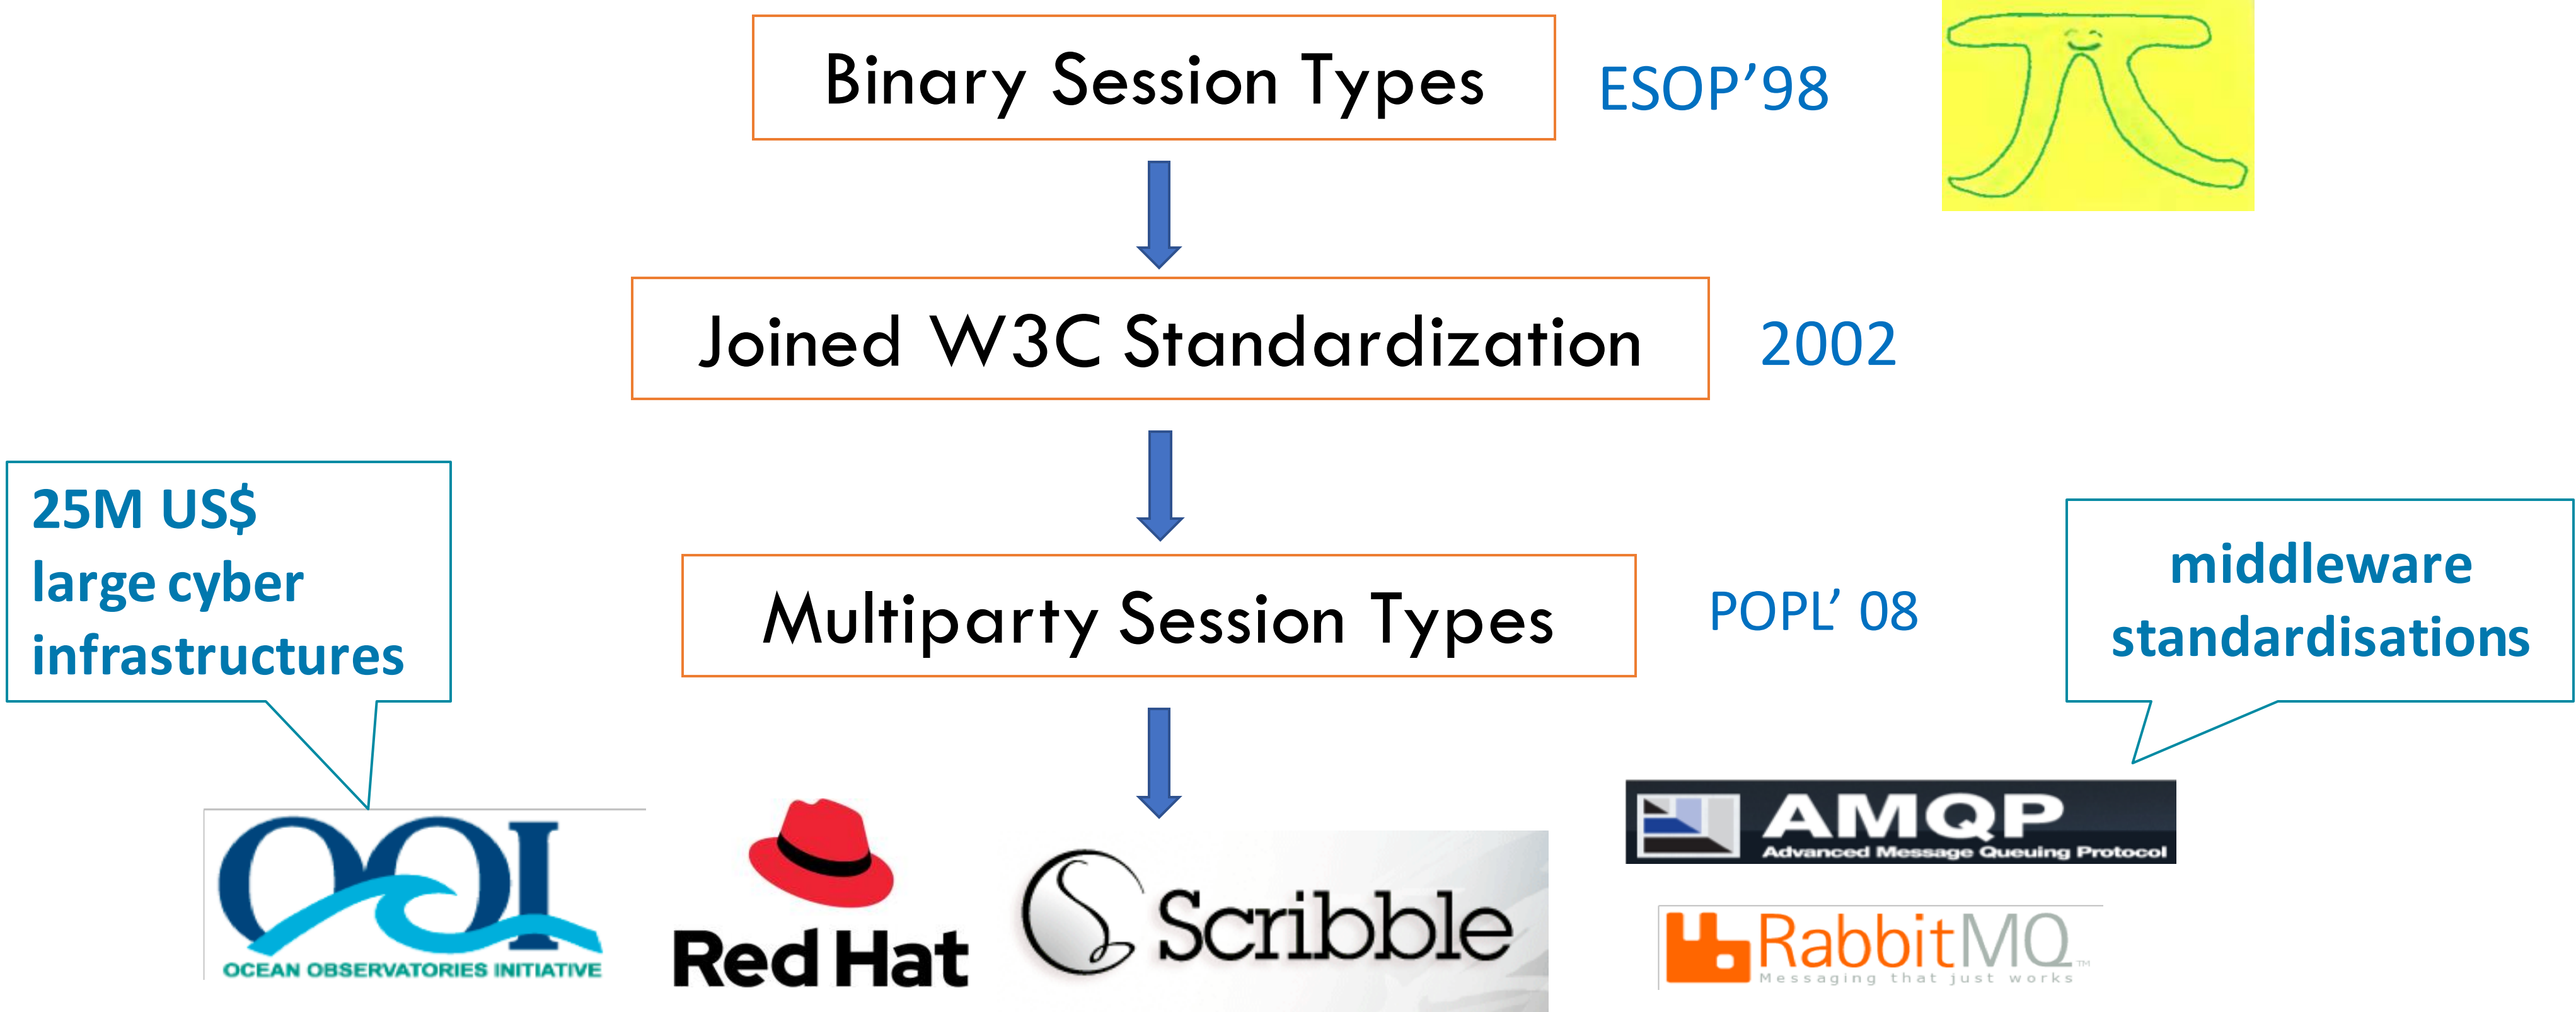


largest open source
company in the world



Why Session Types, Why Now?

Significant academic and industry interests via fundamental breakthroughs




Why Session Types, Why Now?

Significant academic and industry interests via fundamental breakthroughs

Binary Session Types

ESOP'98



Joined W3C Standardization

2002




Multiparty Session Types

POPL'08



TypeScript



Scala

akka



ERLANG

MPI




Why Session Types, Why Now?

Significant academic and industry interests via fundamental breakthroughs

ETAPS Test Time Award 2019

Binary Session Types

ESOP'98



Joined W3C Standardization

2002



Multiparty Session Types

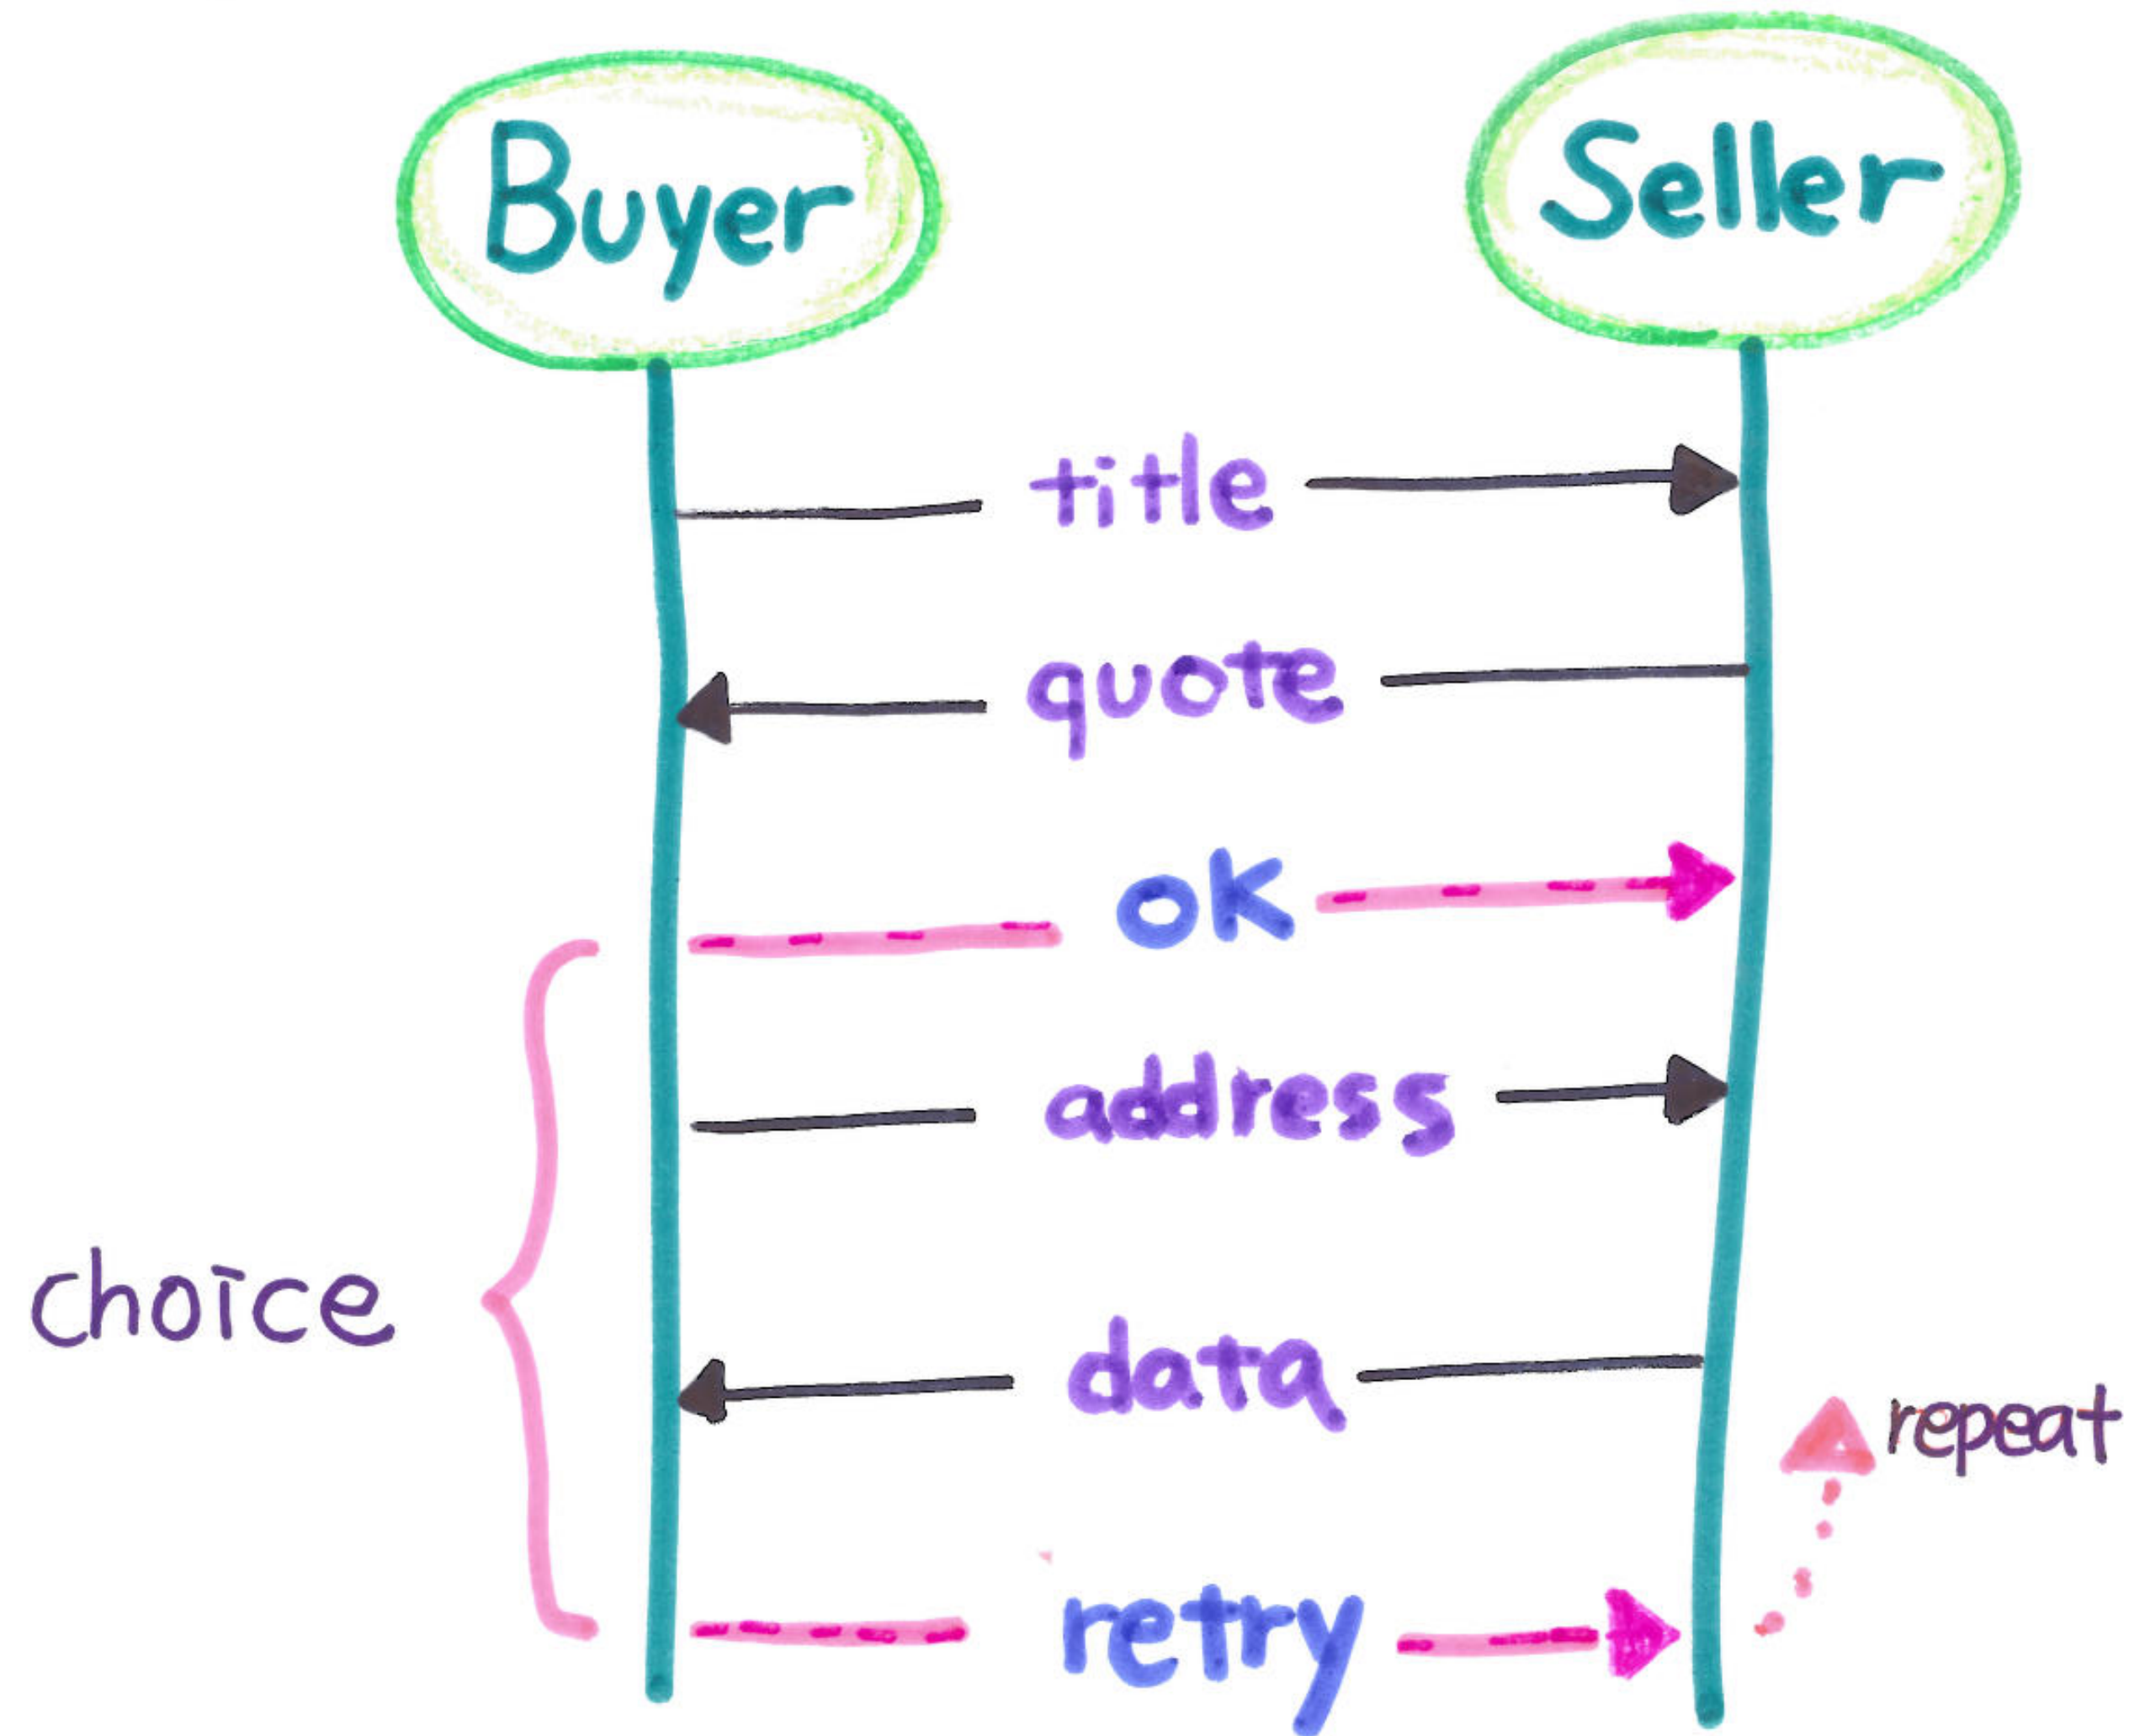
POPL' 08

POPL Influential Paper Award 2018

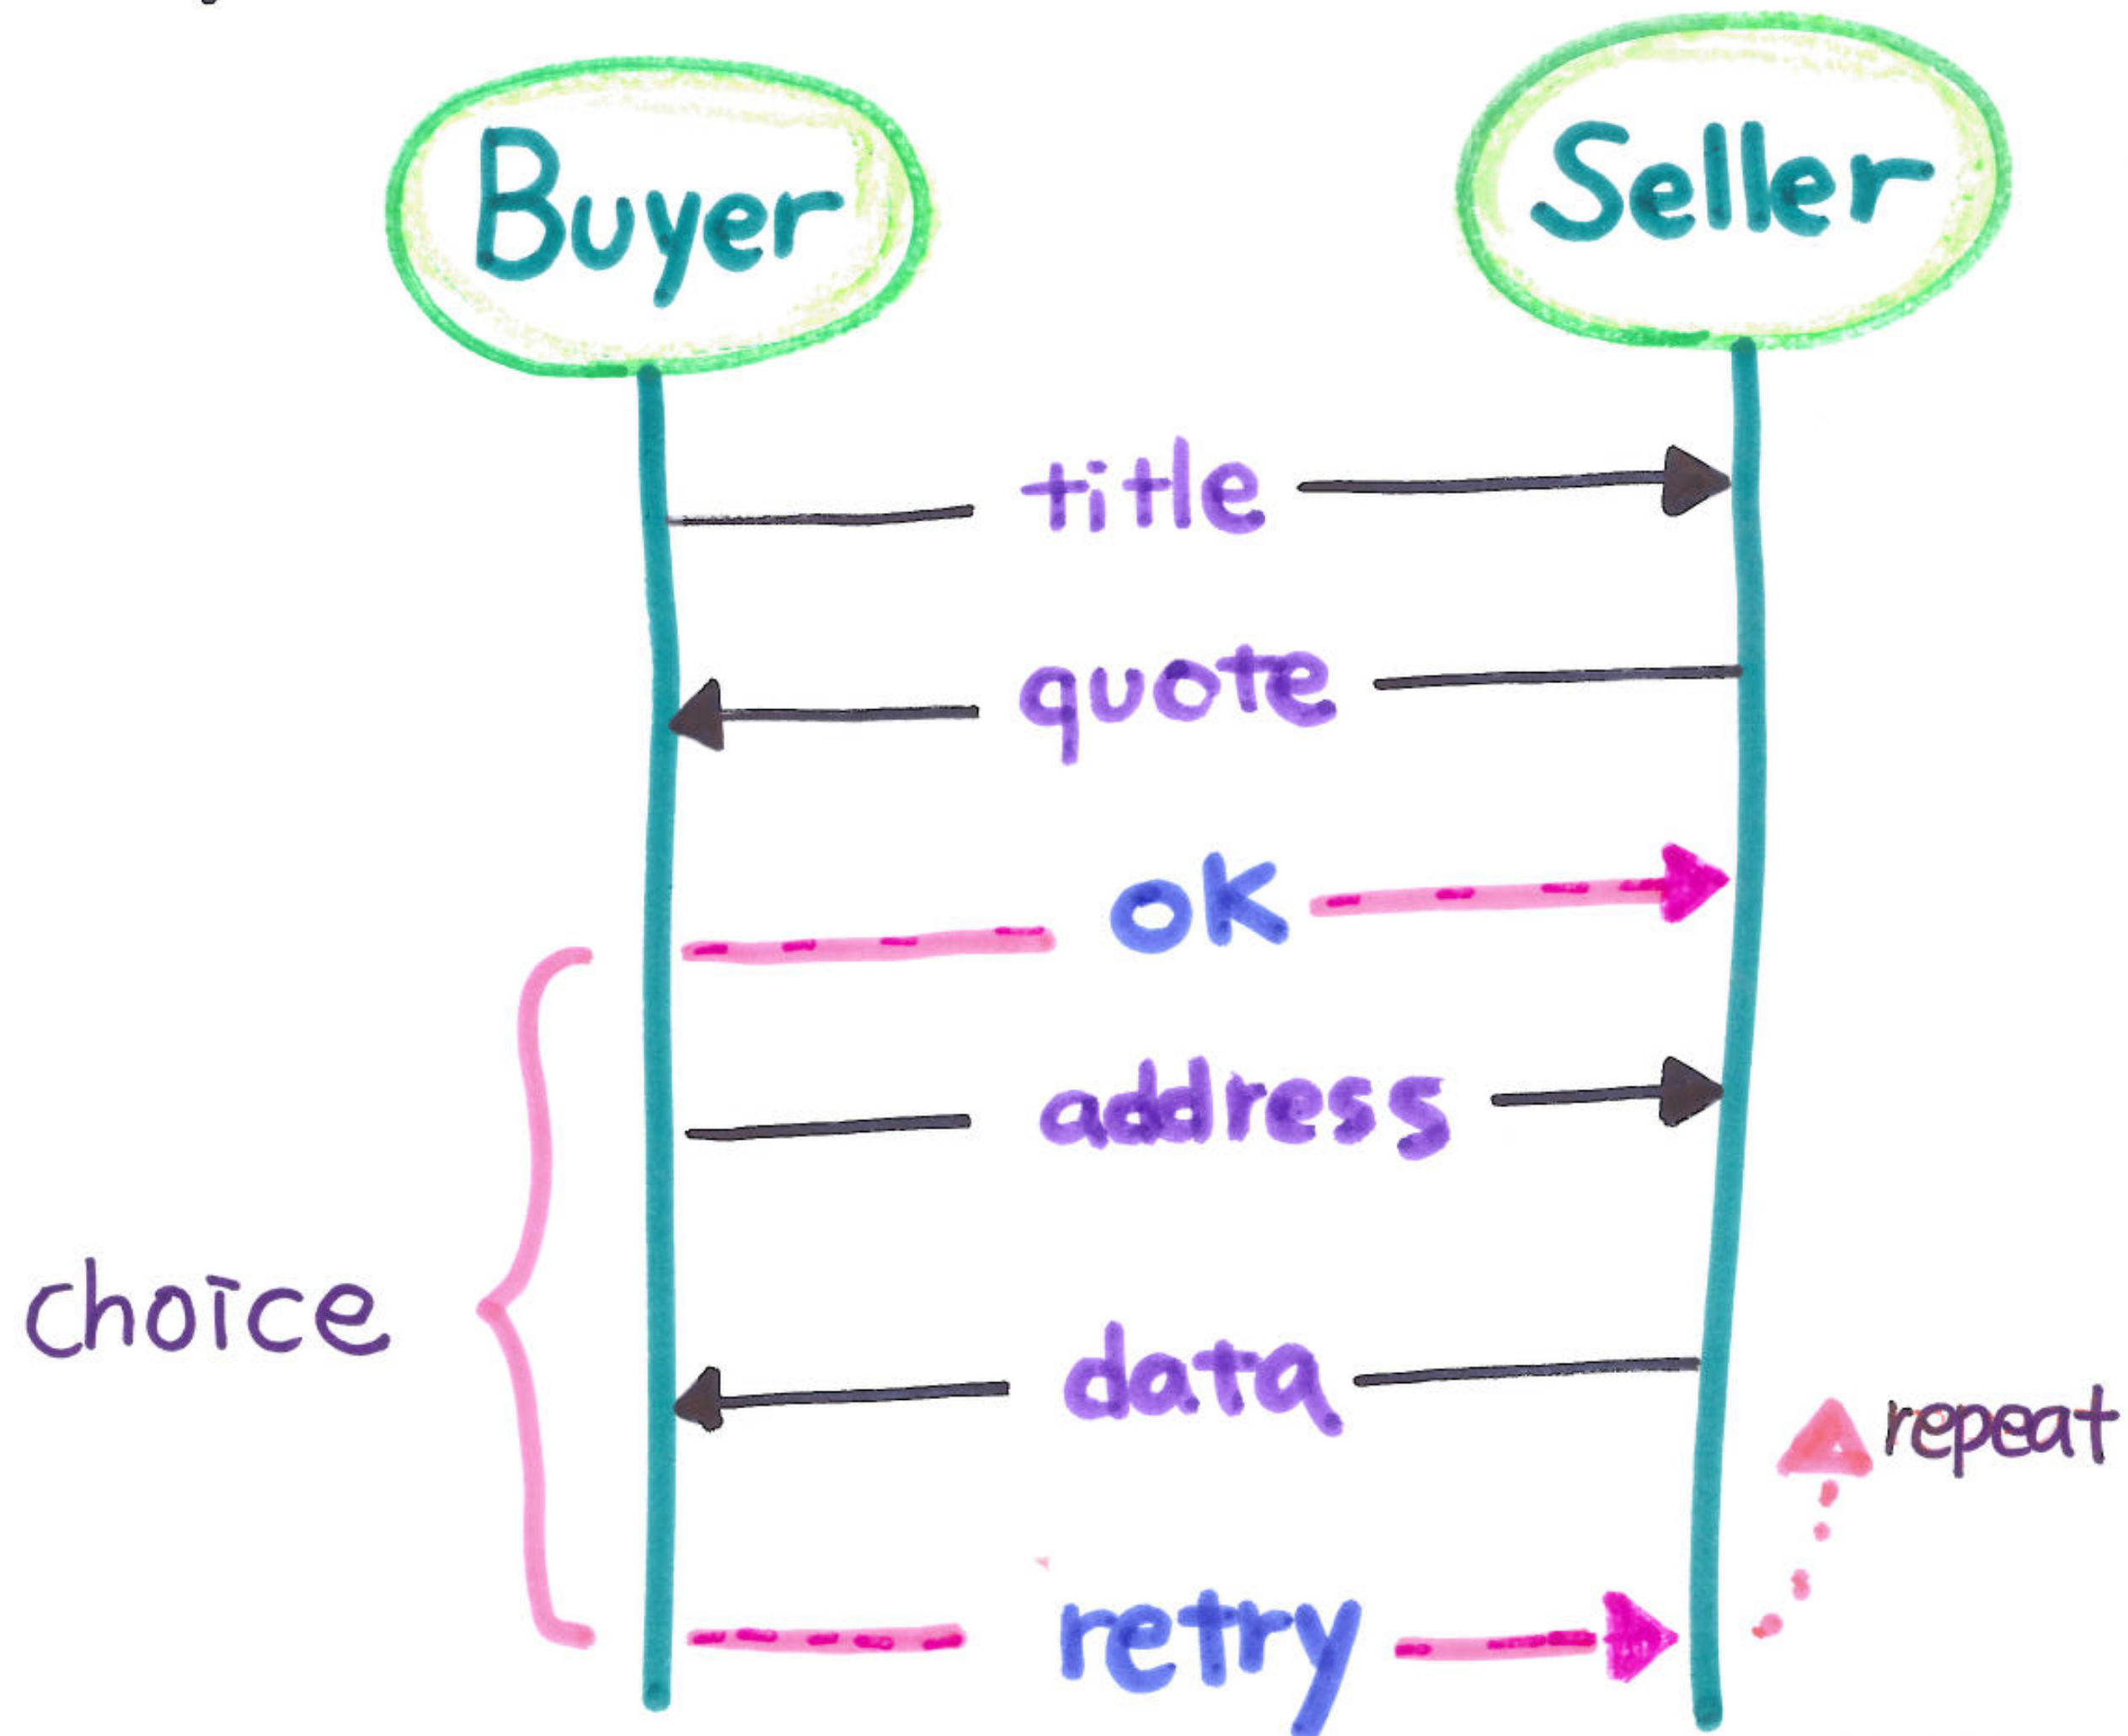


A collection of logos for various programming languages and frameworks, including Java, Go, OOI (Ocean Observatories Initiative), Red Hat, Scribble, AMQP (Advanced Message Queuing Protocol), RabbitMQ, Python, Scala, akka, Erlang, MPI, and OCaml. The logos are arranged in a horizontal line at the bottom of the slide.


Binary Session Types: Buyer - Seller Protocol



Binary Session Types: Buyer - Seller Protocol



nt! Title ; ? Quote ; ! { ok: ! Add ; ? Date, retry: t }

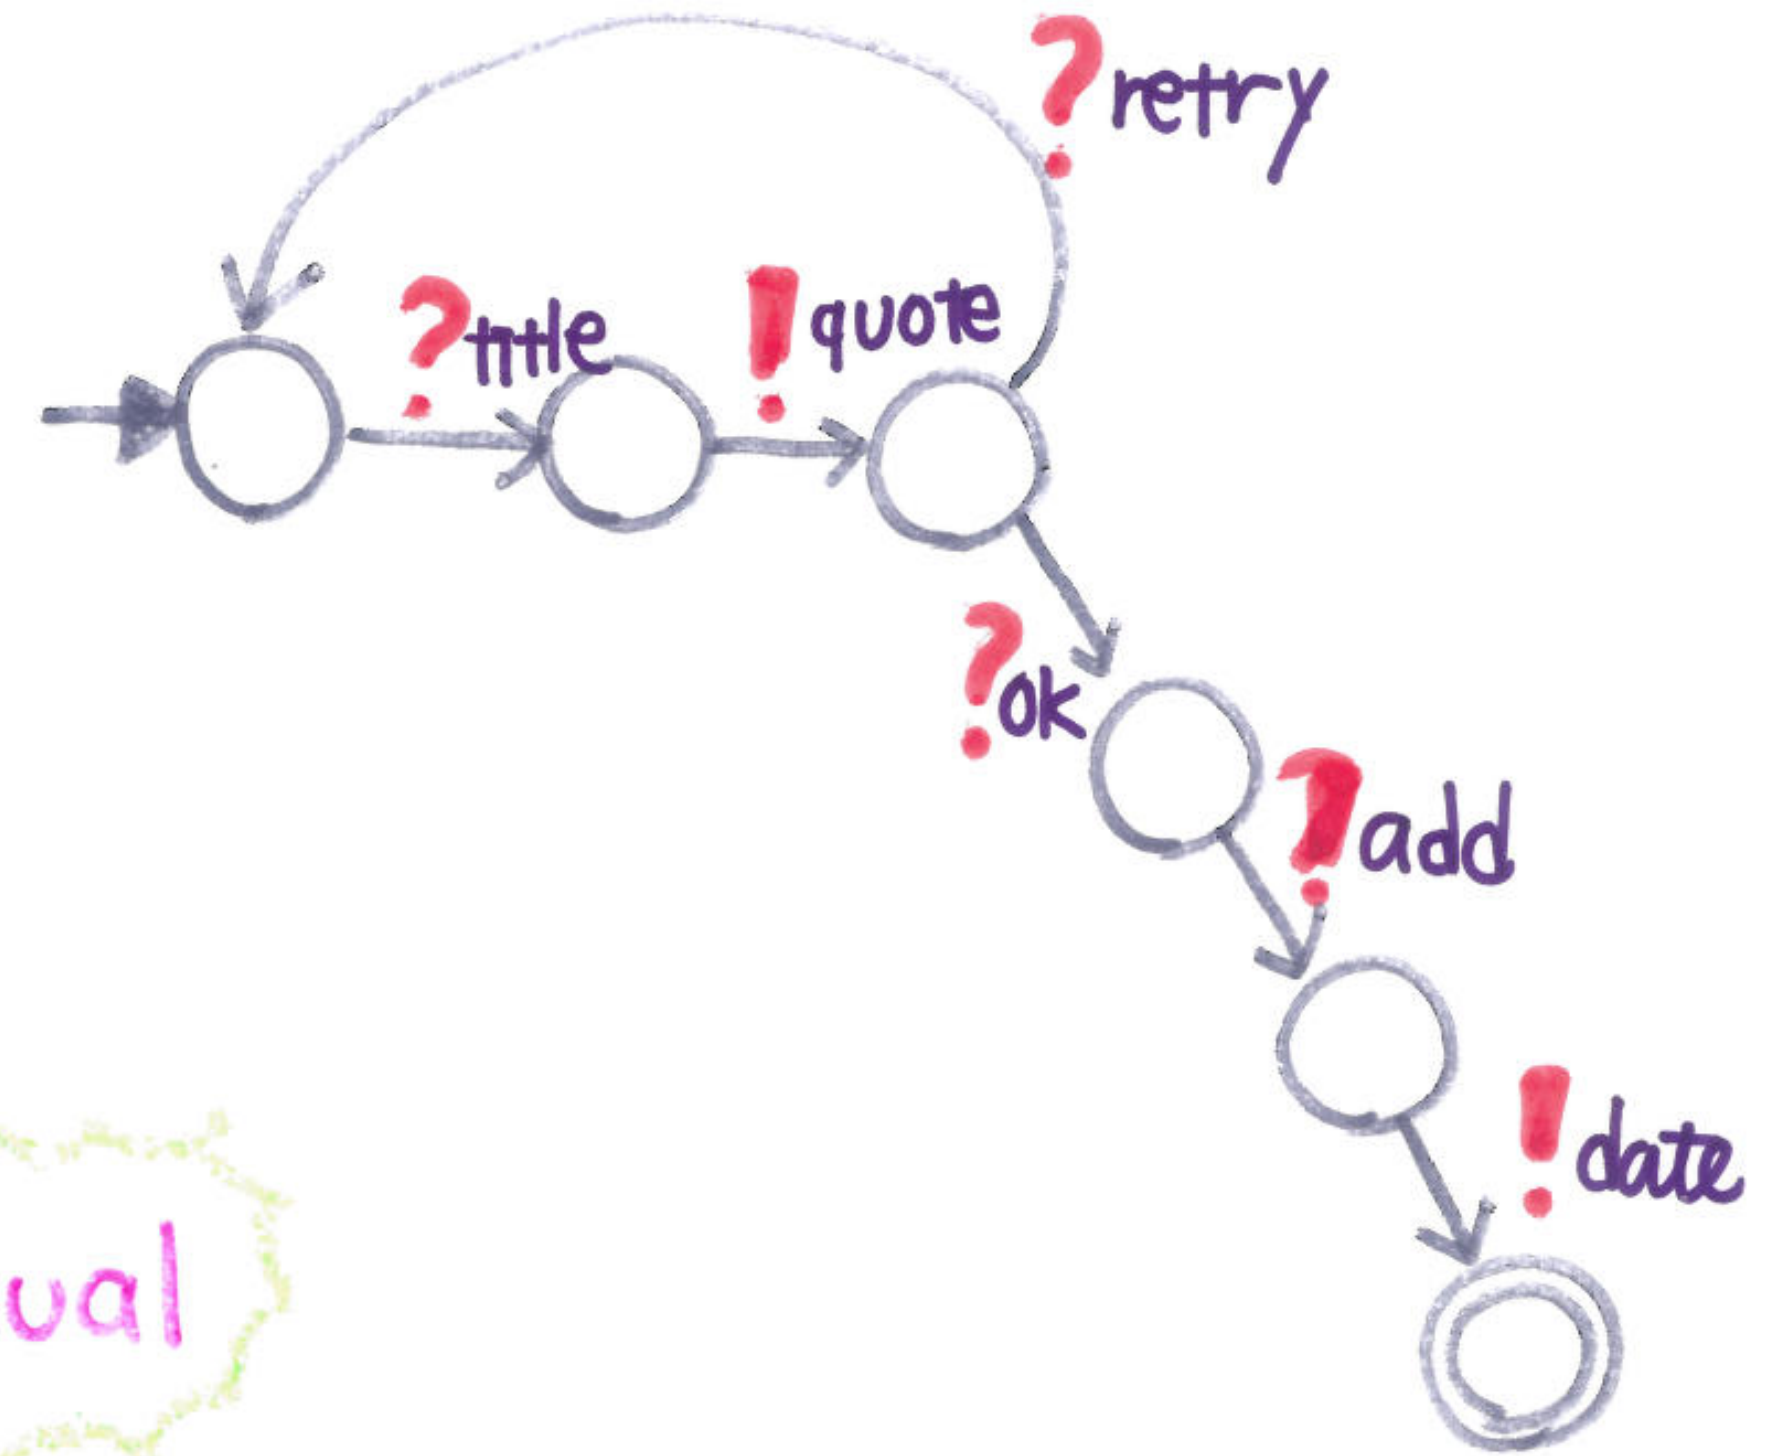



P has T
 Q has \overline{T} *dual*
 P | Q typable

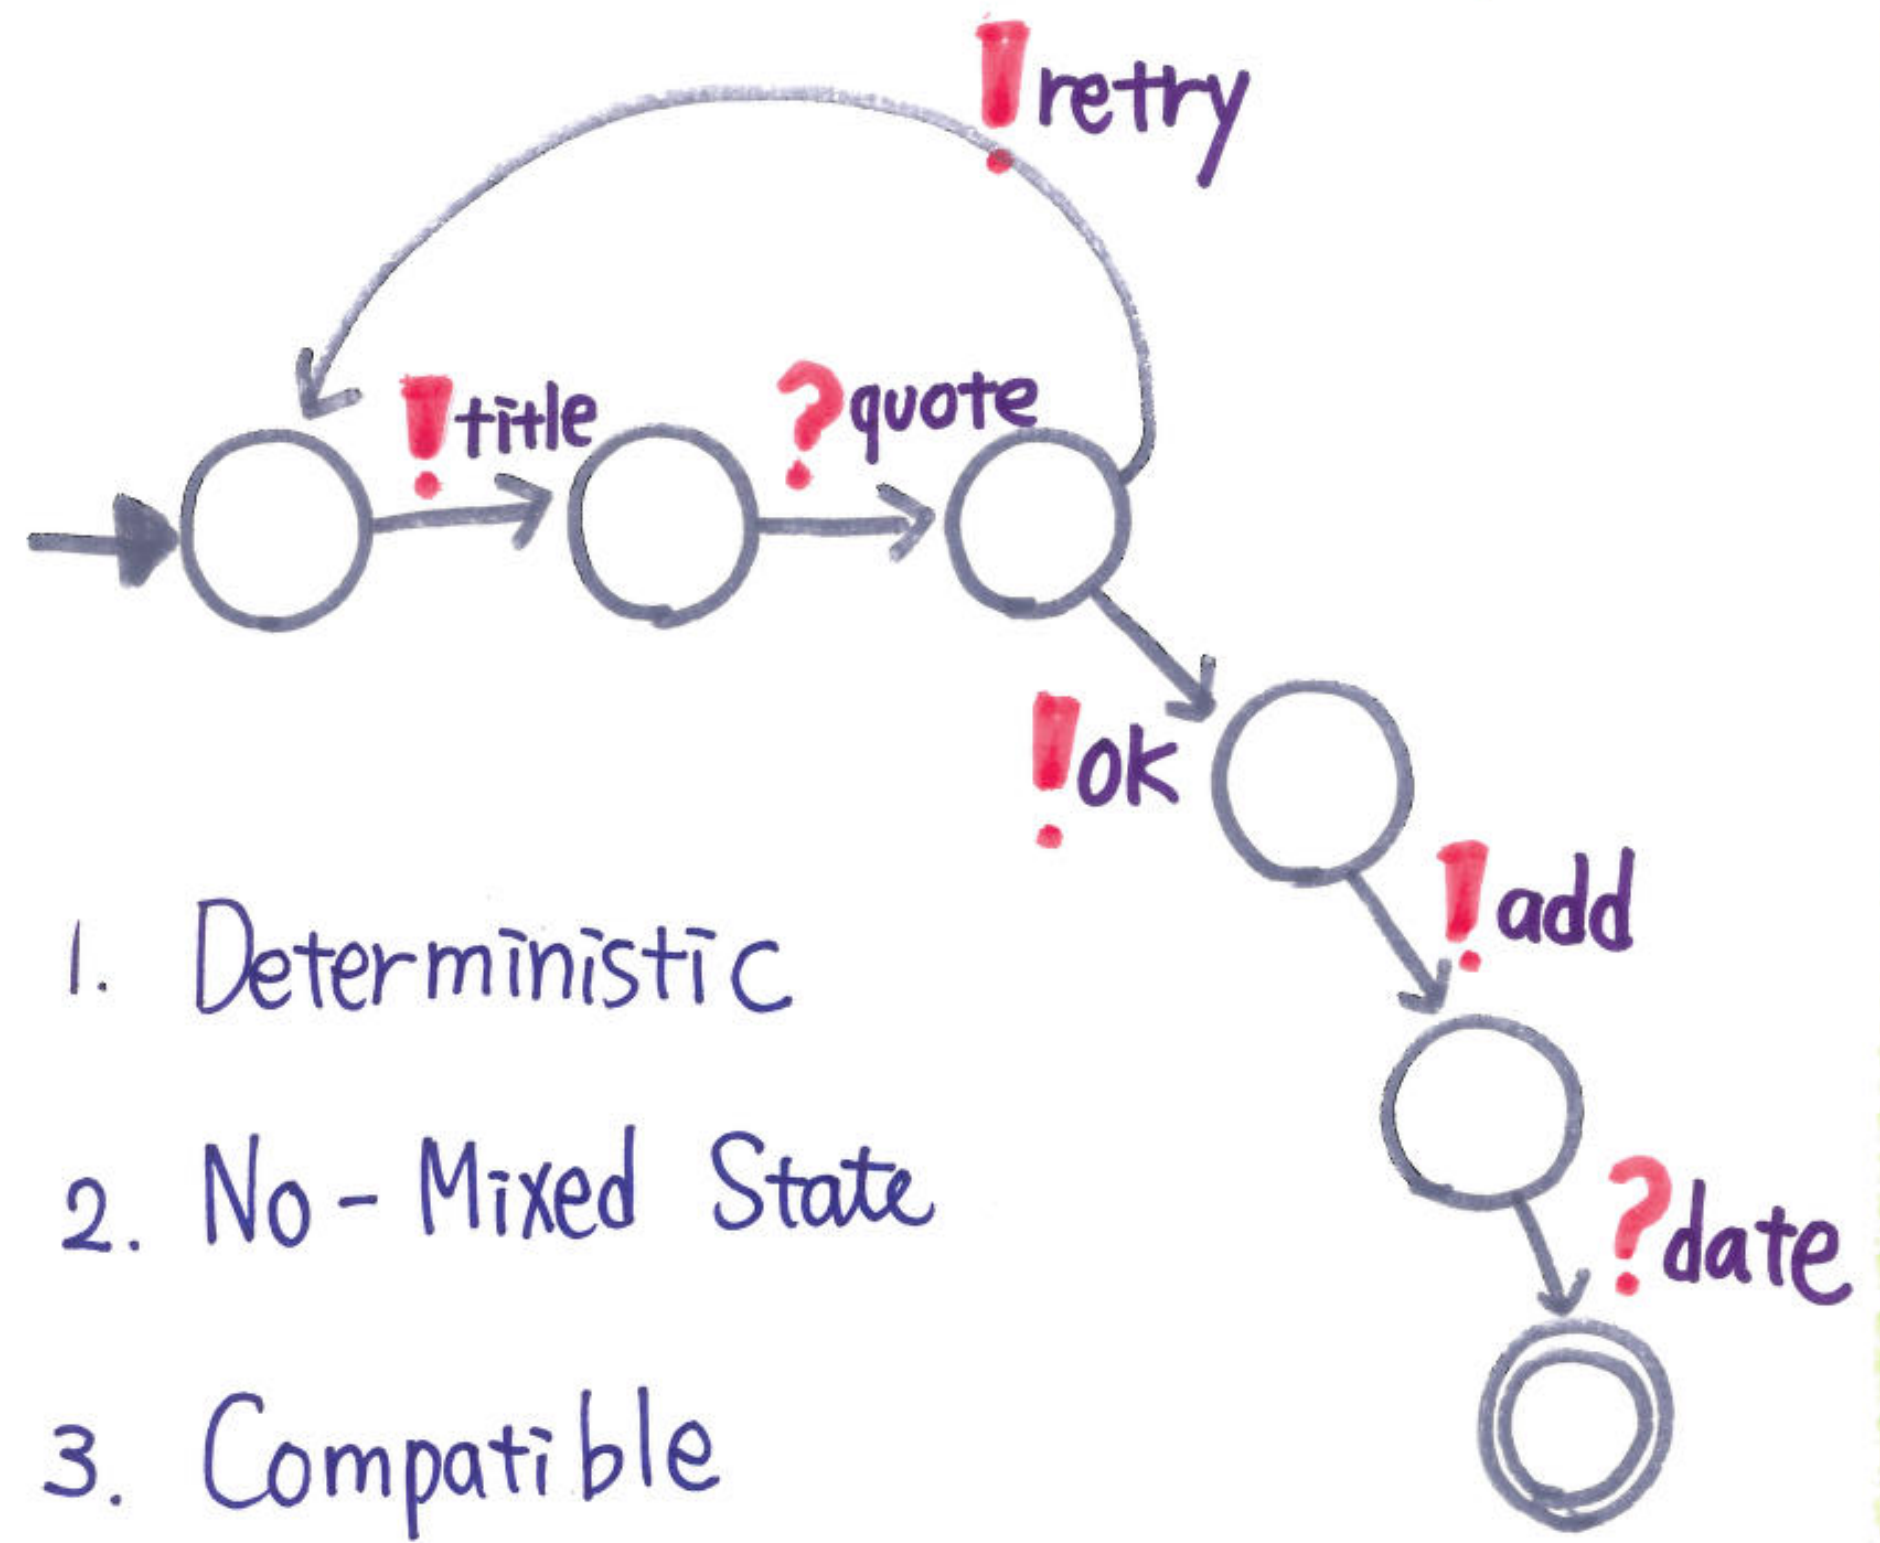
nt! Title ; ? Quote ; ! { ok: ! Add ; ? Date, **retry**: t }

nt? Title ; ! Quote ; ? { ok: ? Add ; ! Date, **retry**: t }


Communicating Automata [1980s]



dual



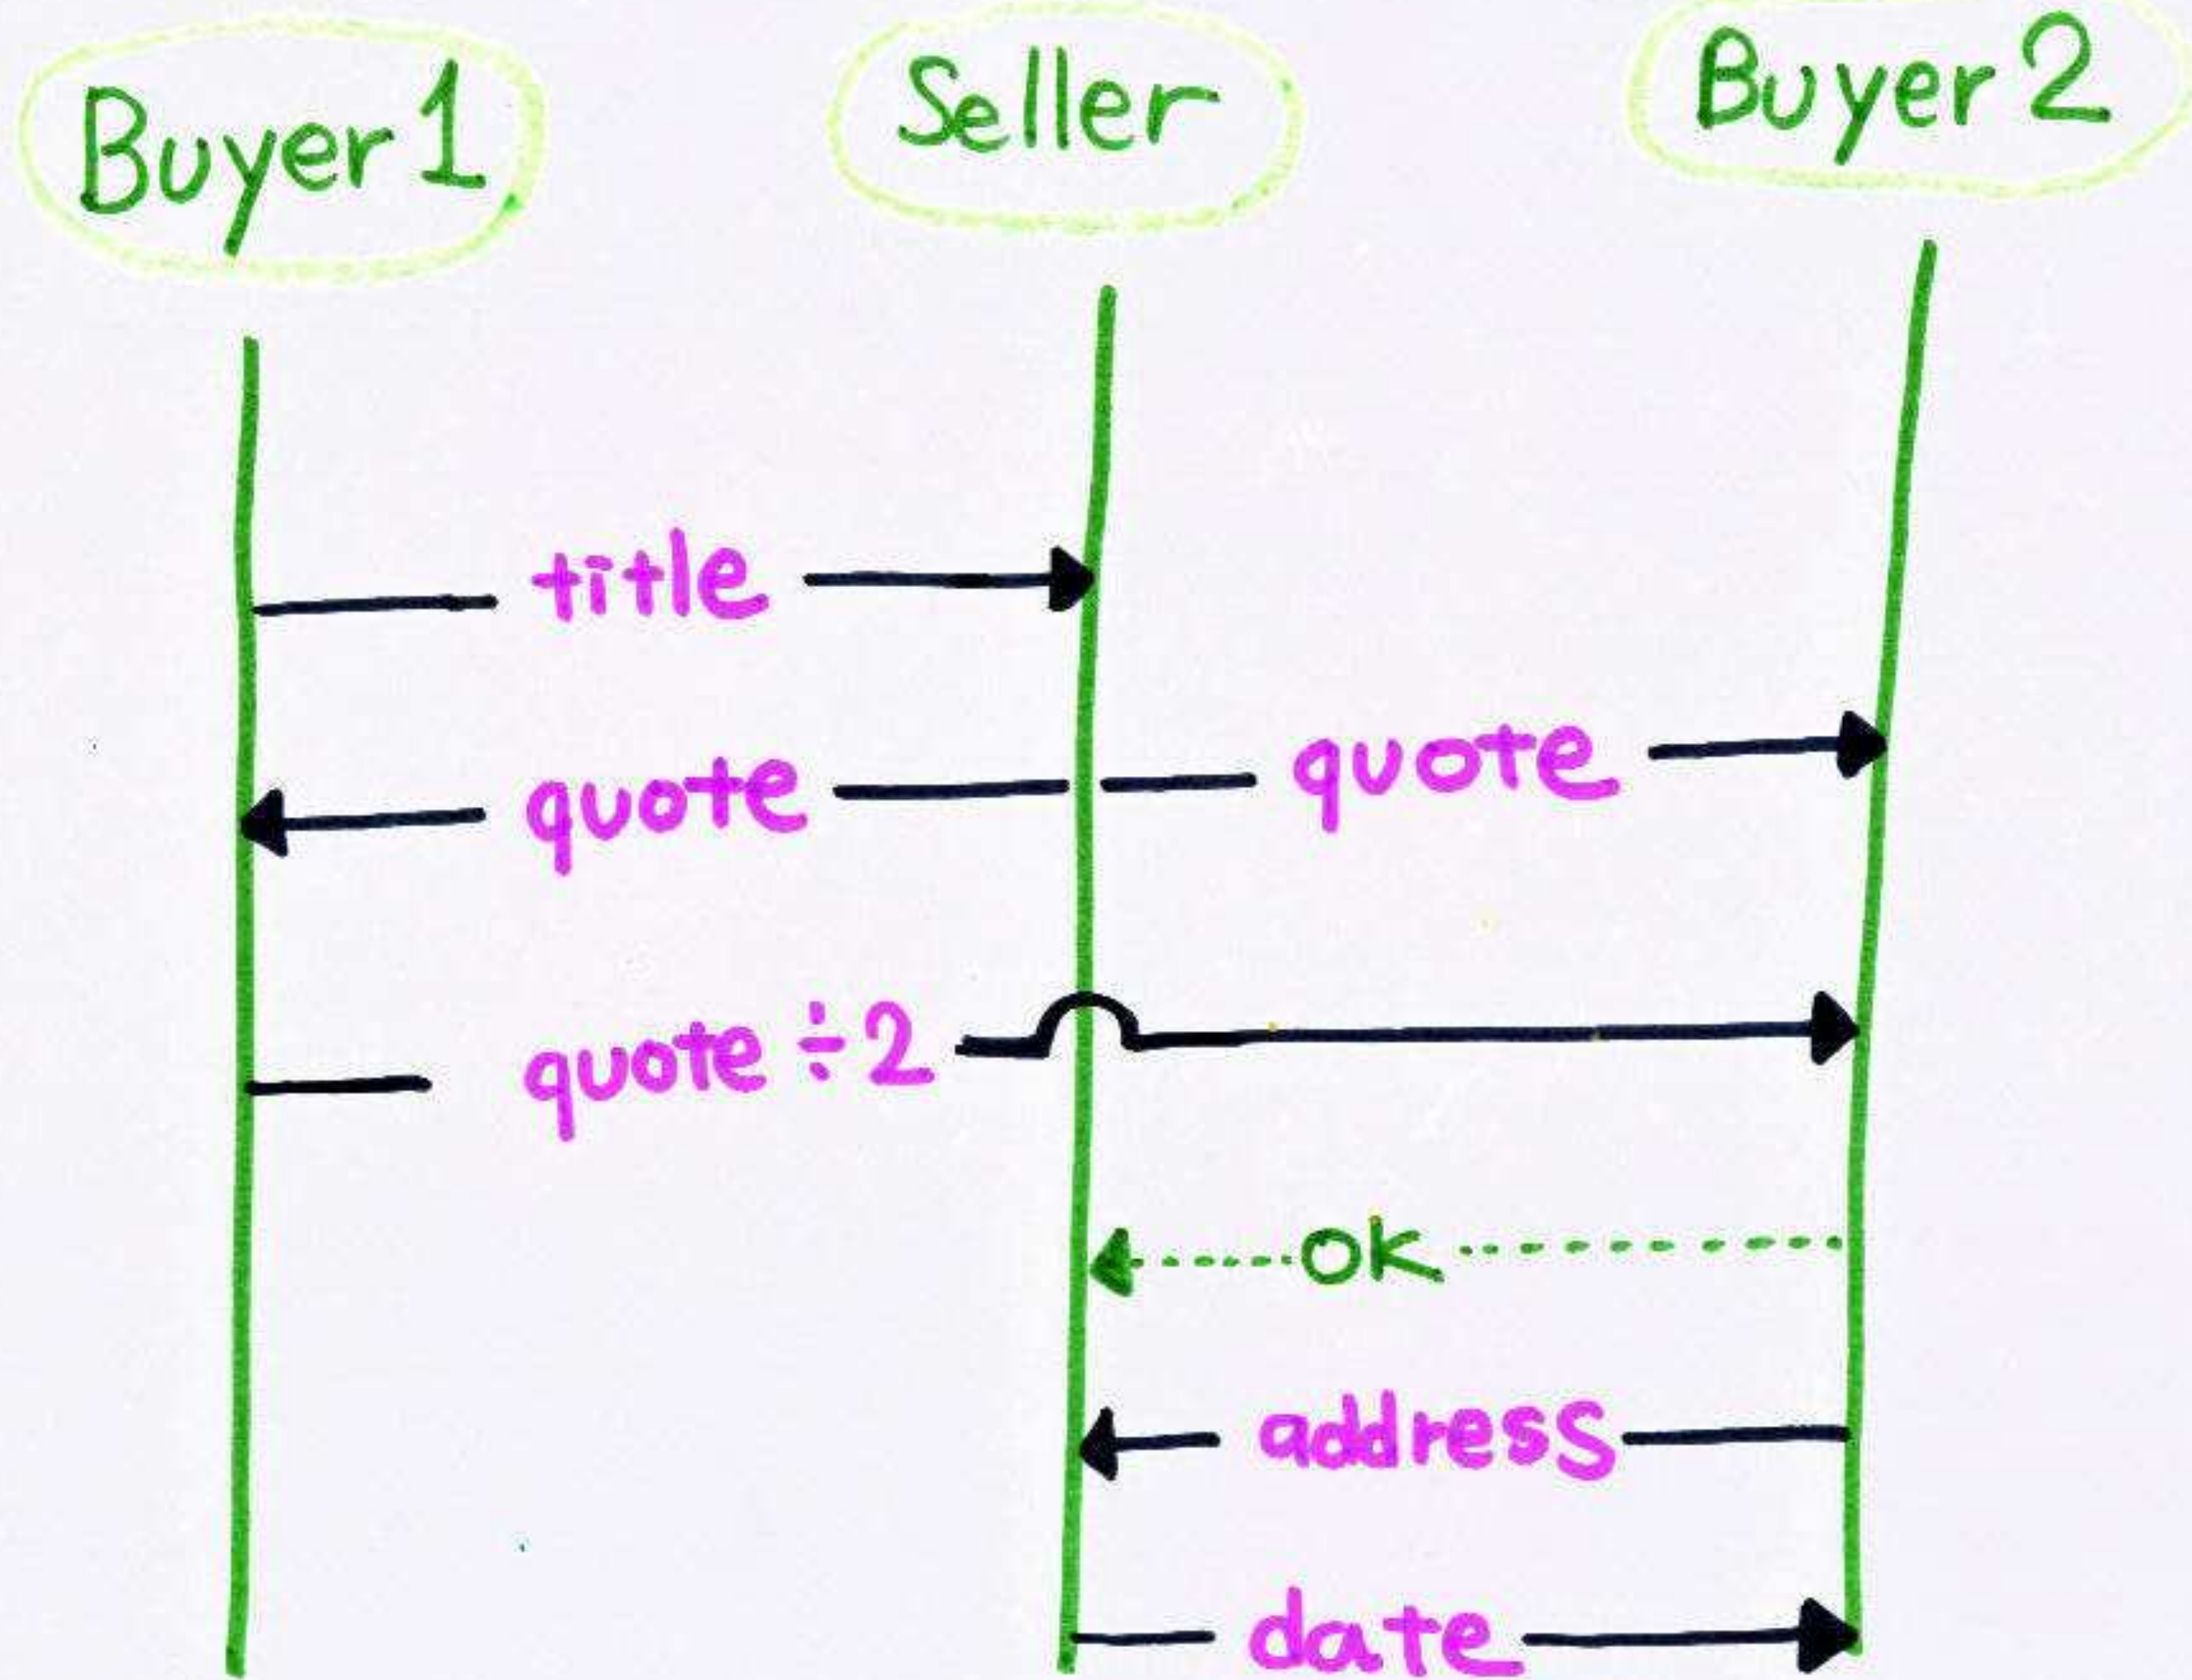
1. Deterministic
2. No - Mixed State
3. Compatible





dual

[Gouda et al 1986] Two compatible machines without mixed states which are deterministic satisfy deadlock-freedom.

Multiparty Session Types








Multiparty Session Types


Buyer1

Seller

Buyer2



Multi party Session Types [Honda, Yoshida, Carbone 2008]



Ⓞ G


BI \rightarrow S Int.

S \rightarrow B2 Char

STEP 1

Write Global Type

Multi party Session Types [Honda, Yoshida, Carbone 2008]



(G) $B_1 \rightarrow S$ Int.
 $S \rightarrow B_2$ Char


STEP 1
Write Global Type

(T) $B_1?Int. B_2!Char$

STEP 2
Project to Local Types

Multi party Session Types

[Honda, Yoshida, Carbone 2008]



(G) $B1 \rightarrow S \text{ Int.}$
 $S \rightarrow B2 \text{ Char}$

(T) $B1? \text{Int. } B2! \text{Char}$

(P) $B1?(x). B2! \langle \text{"apple"} \rangle$

STEP 1

Write Global Type


STEP 2

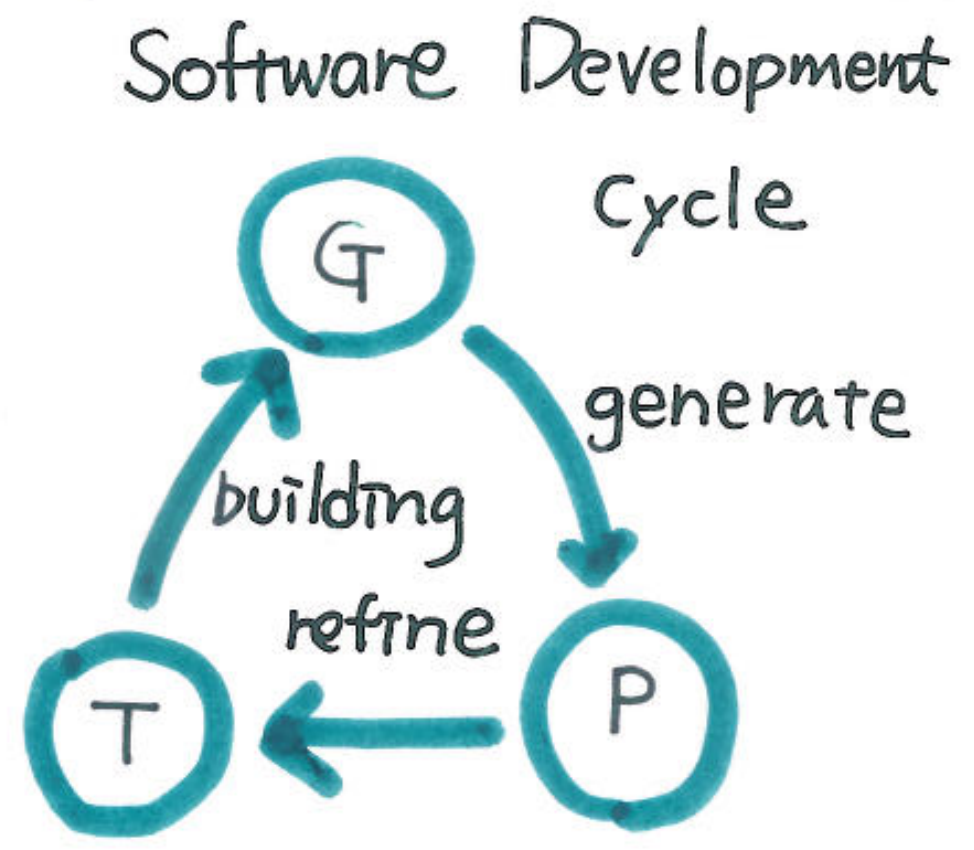
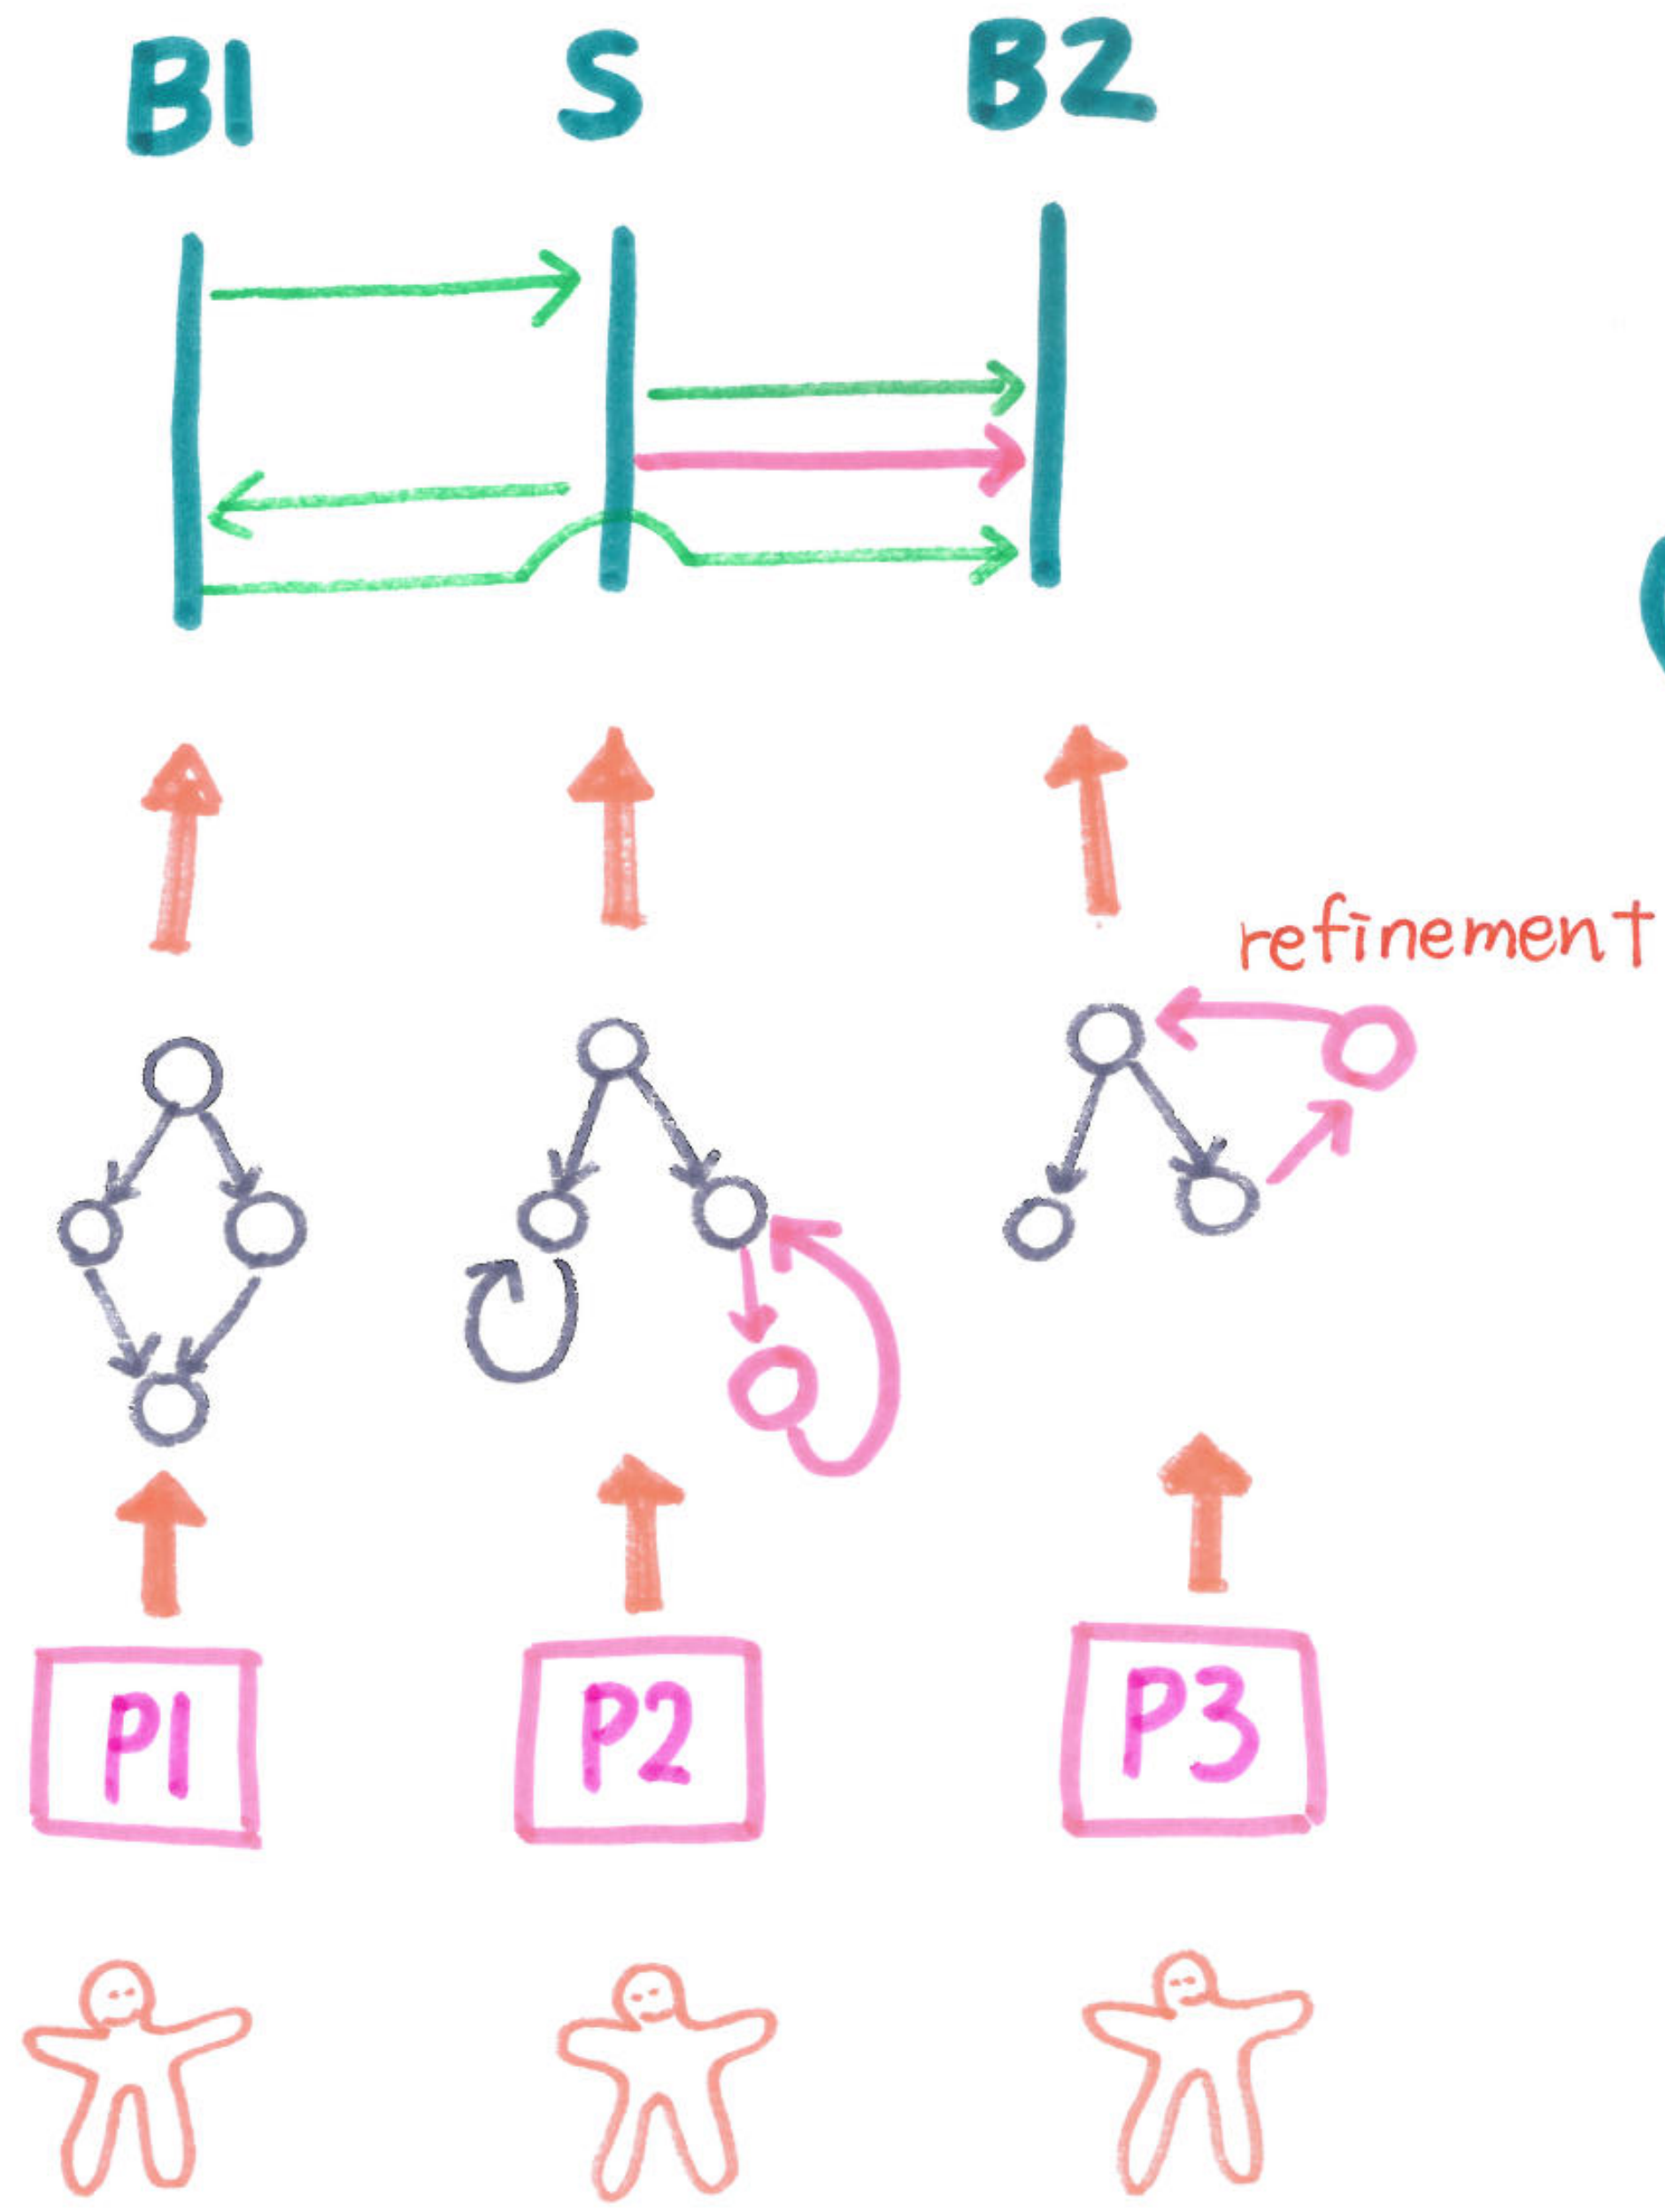
Project to Local Type

STEP 3

- Static Check
- Generate Code
- Run-time check



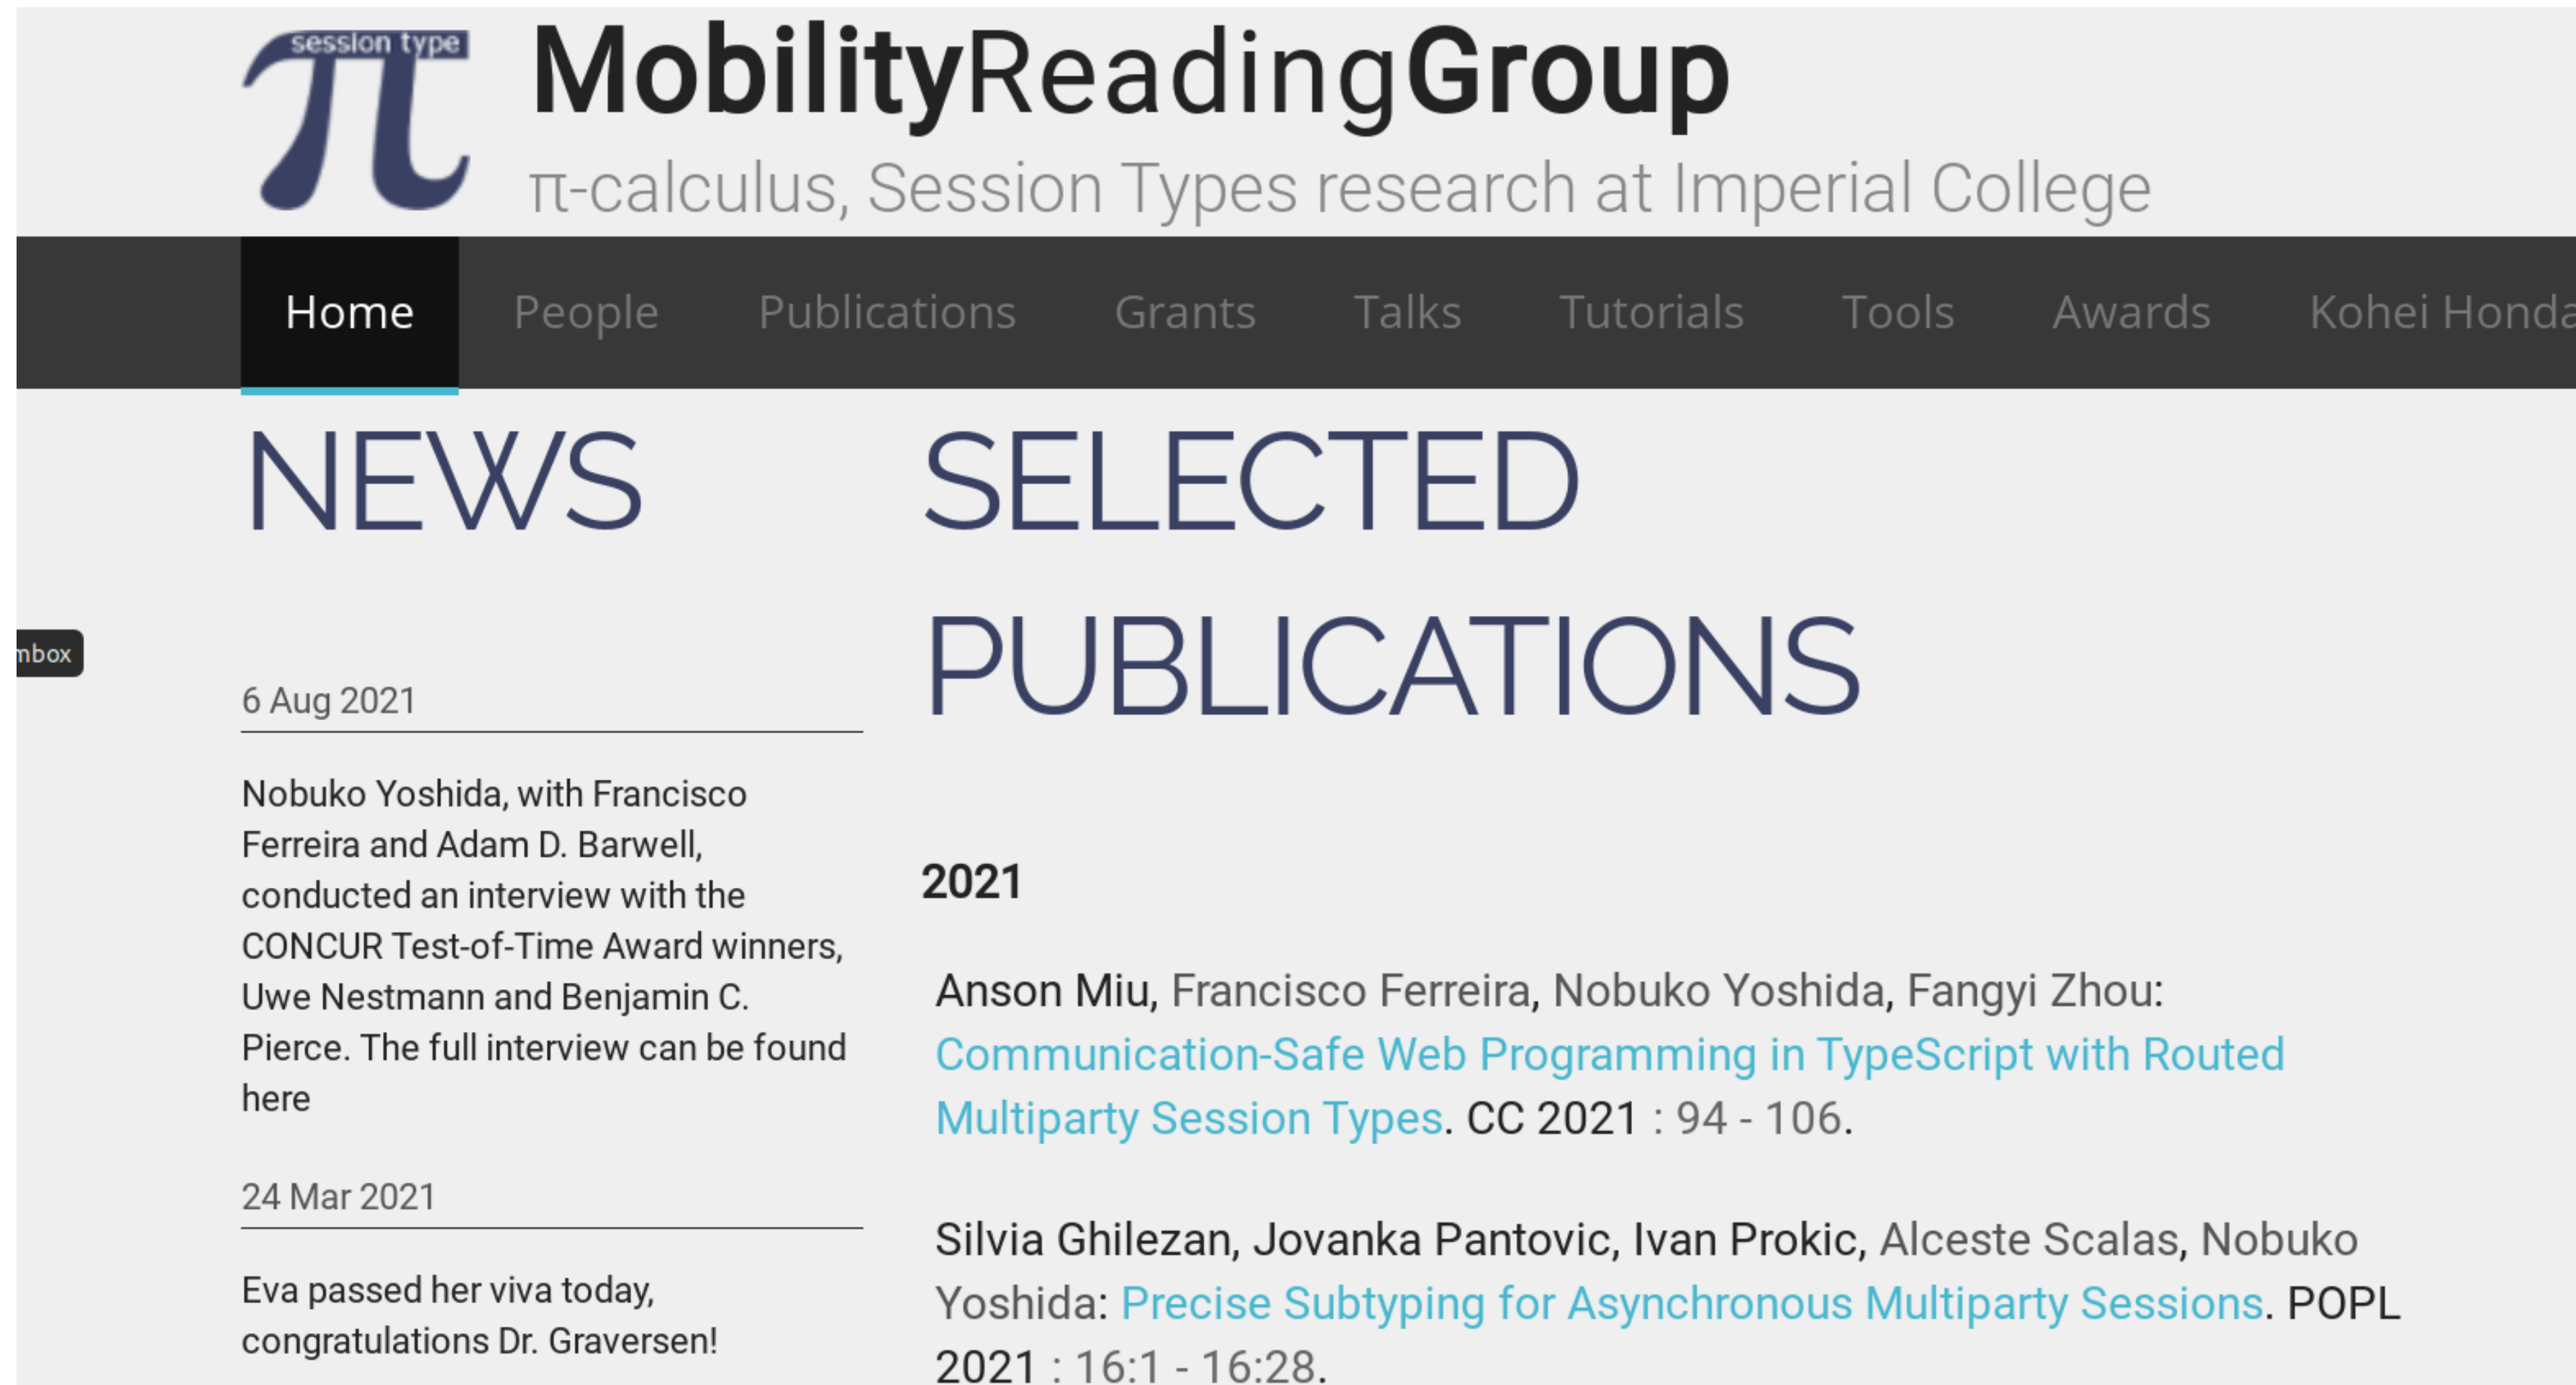




- Optimisation
- refinement
- inference
- Testing

Mobility Reading Group

<http://mrg.doc.ic.ac.uk/>



The screenshot shows the website for the Mobility Reading Group. At the top, there is a logo consisting of a blue pi symbol with the text "session type" above it, followed by the title "MobilityReadingGroup" and the subtitle "π-calculus, Session Types research at Imperial College". Below this is a dark navigation bar with links for "Home", "People", "Publications", "Grants", "Talks", "Tutorials", "Tools", "Awards", and "Kohei Honda". The main content area is split into two columns. The left column is titled "NEWS" and contains two entries: one dated "6 Aug 2021" about an interview with CONCUR award winners, and another dated "24 Mar 2021" congratulating Dr. Graversen. The right column is titled "SELECTED PUBLICATIONS" and lists two papers from 2021: "Communication-Safe Web Programming in TypeScript with Routed Multiparty Session Types" and "Precise Subtyping for Asynchronous Multiparty Sessions".

session type **MobilityReadingGroup**
π-calculus, Session Types research at Imperial College

Home People Publications Grants Talks Tutorials Tools Awards Kohei Honda

NEWS

6 Aug 2021

Nobuko Yoshida, with Francisco Ferreira and Adam D. Barwell, conducted an interview with the CONCUR Test-of-Time Award winners, Uwe Nestmann and Benjamin C. Pierce. The full interview can be found [here](#)

24 Mar 2021

Eva passed her viva today, congratulations Dr. Graversen!

SELECTED PUBLICATIONS

2021

Anson Miu, Francisco Ferreira, Nobuko Yoshida, Fangyi Zhou: [Communication-Safe Web Programming in TypeScript with Routed Multiparty Session Types](#). CC 2021 : 94 - 106.

Silvia Ghilezan, Jovanka Pantovic, Ivan Prokic, Alceste Scalas, Nobuko Yoshida: [Precise Subtyping for Asynchronous Multiparty Sessions](#). POPL 2021 : 16:1 - 16:28.

History of Session Types with mCRL2

- [TACAS'17] Using mCRL2 to benchmark **session subtyping** checking
- [POPL'17] Analysing properties (e.g., deadlock-freedom) of **Go** programs
- [ICSE'18] Analysing properties of **message-passing Go** programs
- [POPL'19] Analysing properties of **synchronous multiparty session pi-calculus**
- [PLDI'19] Analysing properties of **distributed Scala** programs
- [ECOOP'20] Analysing properties of **shared memory Go** programs
- [CONCUR'22] Extension of [POPL'19] with **Stop-Failure**


History of Session Types with mCRL2

- [TACAS'17] Using mCRL2 to benchmark **session subtyping** checking
- [POPL'17] Analysing properties (e.g., deadlock-freedom) of **Go** programs
- [ICSE'18] Analysing properties of **message-passing Go** programs
- [POPL'19] Analysing properties of **synchronous multiparty session pi-calculus**
- [PLDI'19] Analysing properties of **distributed Scala** programs
- [ECOOP'20] Analysing properties of **shared memory Go** programs
- [CONCUR'22] Extension of [POPL'19] with **Stop-Failure**


No gap between type and process properties

Why mCRL2?

- Can analyse properties than the top-down approach
- Can scale




No Failure



with Failure

k-Multiparty Compatibility [CAV'19]



Generalised Multiparty Session Types with Crash-Stop Failures

Adam D. Barwell¹ Alceste Scalas² Nobuko Yoshida¹ Fangyi Zhou¹

¹Imperial College London

²DTU Compute – Technical University of Denmark

33rd International Conference on Concurrency Theory

16 September 2022

Imperial College
London

Technical
University of
Denmark



UK Research
and Innovation



The Case For Session Types

- » Formally describe **communications behaviour** between two or more systems
- » Communications behaviour is **enforced statically**
- » **Guarantee** desirable behavioural properties
 - Deadlock-Freedom
 - Liveness
 - Termination
 - Never-Termination
- » Libraries permit use with **real-world languages**
 - Rust, Haskell, Scala, Erlang, OCaml, Java, Go, Typescript, &c.

The Case **Against** Session Types

- » **Limited practicality** in the real-world
- » Basic assumptions **restrict** what we can express and guarantee
- » One notable example: **fault-tolerant protocols**

Fault-Tolerance or Lack Thereof

- » Most MPST systems assume a **perfect world**
 - No process failures
 - No message loss, duplication, or corruption



Fault-Tolerance or Lack Thereof

- » Most MPST systems assume a **perfect world**
 - No process failures
 - No message loss, duplication, or corruption
- » **Reality** tends to disagree
 - Things break, unfortunately



Fault-Tolerance or Lack Thereof

- » Most MPST systems assume a **perfect world**
 - No process failures
 - No message loss, duplication, or corruption
- » **Reality** tends to disagree
 - Things break, unfortunately

Aim

Represent fault-tolerant protocols.



Fault-Tolerant Session Types

We present a **generalised multiparty session type theory** with **crash-stop failures**.

Fault-Tolerant Session Types

We present a **generalised multiparty session type theory** with **crash-stop failures**.

- » **Crash-stop failures**: our failure model; provides failures to tolerate

Fault-Tolerant Session Types

We present a **generalised multiparty session type theory** with **crash-stop failures**.

- » **Crash-stop failures**: our failure model; provides failures to tolerate
- » **Generalised MPST**: our type system is parametric on a safety invariant
 - Safety can be a behavioural property (e.g. deadlock-freedom)
 - Enables integration with model checking

Fault-Tolerant Session Types

We present a **generalised multiparty session type theory** with **crash-stop failures**.

- » **Crash-stop failures**: our failure model; provides failures to tolerate
- » **Generalised MPST**: our type system is parametric on a **safety invariant**
 - Safety can be a behavioural property (e.g. deadlock-freedom)
 - Enables integration with model checking


Fault-Tolerant Session Types

We present a **generalised multiparty session type theory** with **crash-stop failures**.

1. Session π -calculus with crash-stop failures
2. Multiparty session types with crashes
3. Type-level model checking
4. Optional reliability assumptions


Crash-Stop Failures

- » Processes can **crash arbitrarily**
- » Crashed processes **make no progress**
- » Crashed processes **do not recover**
- » Communications channels deliver messages in order and without losses



Crash-Stop Failures

- » Processes can **crash arbitrarily**
- » Crashed processes **make no progress**
- » Crashed processes **do not recover**
- » Communications channels deliver messages in order and without losses



- » How do we represent crash-stop failures in our π -calculus and types?

A Session π -Calculus with Crash-Stop Failures

'act as participant p in session s '



$c ::= x \mid s[p]$	Variables and session endpoints
$d ::= v \mid c$	Basic value, variable, or session endpoint
$P, Q ::= 0 \mid (\nu s)P \mid P \mid Q$	Inaction, restriction, and parallel composition
$\mid c[q] \oplus_m \langle d \rangle . P$ (where $m \neq \mathbf{crash}$)	Send to participant q as c
$\mid c[q] \& \{m_i(x_i) . P_i\}_{i \in I}$	Receive from participant q as c
$\mid s[p] \downarrow$	Crashed session endpoint

- » $s[p] \downarrow$ represents that participant p has crashed in session s
- » **crash** is a reserved label that cannot be sent


Some standard constructs have been omitted for clarity of presentation.

A Simple DNS Protocol

A client P queries DNS server Q for an IP address.

$$P = s[p][q] \oplus \text{url}.s[p][q] \& \text{res}.0$$

$$Q = s[q][p] \& \text{url}.s[q][p] \oplus \text{res}.0$$



A Simple DNS Protocol

A client P queries DNS server Q for an IP address.

$P = s[p][q] \oplus \text{url}.s[p][q] \&\text{res.0}$

$Q = s[q] \downarrow$

How does P handle Q 's crash?



Handling Crashes

- » Receive operations may contain a **crash-handling branch**

$$s[p][q] \& \{ \dots, \text{crash}.k \}$$

- » Each process has a **failure detector**
 - Strong completeness & accuracy: processes detected to have crashed must have crashed
 - Exact method is abstract; a variety to choose from: timeouts, heartbeat, &c.
 - Relaxing this restriction is future work
- » A crash-handling branch is taken when the sender is detected to have crashed
 - Crashes are handled locally
 - Failures are not implicitly broadcast to peers

Adding Crash-Handling Behaviour to the DNS Protocol


Fail-over server

Activates only when q is detected to have crashed

$$R = s[r][q] \& \mathbf{crash.s[r][p]} \& \mathbf{url.s[r][p]} \oplus \mathbf{res.0}$$

$$P = s[p][q] \oplus \mathbf{url.s[p][q]} \& \left\{ \begin{array}{l} \mathbf{res.0} \\ \mathbf{crash.s[p][r]} \oplus \mathbf{url.s[p][r]} \& \mathbf{res.0} \end{array} \right\}$$

$$Q = s[q] \& \mathbf{\downarrow}$$



Multiparty Session Types with Crash-Stop Failures

T	$::=$	$p \oplus \{m_i(S_i).T_i\}_{i \in I}$	Send to participant p
		$p \& \{m_i(S_i).T_i\}_{i \in I}$	Receive from participant p
		$\mu t.T$ t end	Recursion, type variable, and termination
U	$::=$	T	Pure session type
		stop	Crash type

- » Session types, U , are assigned to session endpoints
- » stop represents a crashed endpoint; a **runtime type**
- » **crash** is a reserved label, denotes crash-handling branches

Basic and payload types have been omitted for clarity of presentation.

Adding Session Types to the DNS Protocol

$$R = s[r][q] \& \mathbf{crash}.s[r][p] \& \mathbf{url}.s[r][p] \oplus \mathbf{res}.0$$

$$P = s[p][q] \oplus \mathbf{url}.s[p][q] \& \left\{ \begin{array}{l} \mathbf{res}.0 \\ \mathbf{crash}.s[p][r] \oplus \mathbf{url}.s[p][r] \& \mathbf{res}.0 \end{array} \right\}$$

$$Q = s[q] \downarrow$$

Adding Session Types to the DNS Protocol

$s[r]$: $q \& \mathbf{crash}.p \& \mathbf{url}.p \oplus \mathbf{res}.end$

R = $s[r][q] \& \mathbf{crash}.s[r][p] \& \mathbf{url}.s[r][p] \oplus \mathbf{res}.0$

P = $s[p][q] \oplus \mathbf{url}.s[p][q] \& \left\{ \begin{array}{l} \mathbf{res}.0 \\ \mathbf{crash}.s[p][r] \oplus \mathbf{url}.s[p][r] \& \mathbf{res}.0 \end{array} \right\}$

Q = $s[q] \downarrow$

Adding Session Types to the DNS Protocol

$s[r]$: $q \& \mathbf{crash}.p \& \mathbf{url}.p \oplus \mathbf{res}.end$

R = $s[r][q] \& \mathbf{crash}.s[r][p] \& \mathbf{url}.s[r][p] \oplus \mathbf{res}.0$

$s[p]$: $q \oplus \mathbf{url}.q \& \left\{ \begin{array}{l} \mathbf{res}.end \\ \mathbf{crash}.r \oplus \mathbf{url}.r \& \mathbf{res}.end \end{array} \right\}$

P = $s[p][q] \oplus \mathbf{url}.s[p][q] \& \left\{ \begin{array}{l} \mathbf{res}.0 \\ \mathbf{crash}.s[p][r] \oplus \mathbf{url}.s[p][r] \& \mathbf{res}.0 \end{array} \right\}$

Q = $s[q] \downarrow$

Adding Session Types to the DNS Protocol

$s[r]$: $q \& \mathbf{crash}.p \& \mathbf{url}.p \oplus \mathbf{res}.end$

$R = s[r][q] \& \mathbf{crash}.s[r][p] \& \mathbf{url}.s[r][p] \oplus \mathbf{res}.0$

$s[p]$: $q \oplus \mathbf{url}.q \& \left\{ \begin{array}{l} \mathbf{res}.end \\ \mathbf{crash}.r \oplus \mathbf{url}.r \& \mathbf{res}.end \end{array} \right\}$

$P = s[p][q] \oplus \mathbf{url}.s[p][q] \& \left\{ \begin{array}{l} \mathbf{res}.0 \\ \mathbf{crash}.s[p][r] \oplus \mathbf{url}.s[p][r] \& \mathbf{res}.0 \end{array} \right\}$

$s[q]$: **stop**

$Q = s[q] \downarrow$

Fault-Tolerant Session Types

We present a **generalised multiparty session type theory** with **crash-stop failures**.

- ✓ Session π -calculus with crash-stop failures

- ✓ Multiparty session types with crashes

⇒ Type-level model checking for behavioural properties

3. Optional reliability assumptions

Desirable Behavioural Properties

Deadlock-Freedom

Lack of progress implies termination.

Liveness

All communications eventually occur.

Terminating

The protocol will eventually cease.

Never-Terminating

The protocol will continue for all eternity.

Desirable Behavioural Properties

Deadlock-Freedom

Lack of progress implies termination.

Liveness

All communications eventually occur.

Terminating

The protocol will eventually cease.

Never-Terminating

The protocol will continue for all eternity.

» All examples of **safety properties**

Desirable Behavioural Properties

Deadlock-Freedom

Lack of progress implies termination.

Liveness

All communications eventually occur.

Terminating

The protocol will eventually cease.

Never-Terminating

The protocol will continue for all eternity.

- » All examples of **safety properties**
- » Abstract properties to one general **safety invariant**, φ

Desirable Behavioural Properties

Deadlock-Freedom

Lack of progress implies termination.

Liveness

All communications eventually occur.

Terminating

The protocol will eventually cease.

Never-Terminating

The protocol will continue for all eternity.

- » All examples of **safety properties**
- » Abstract properties to one general **safety invariant**, φ
- » Typing rules are **parameterised** by φ
 - Enables checking different properties

Desirable Behavioural Properties

Deadlock-Freedom

Lack of progress implies termination.

Liveness

All communications eventually occur.

Terminating

The protocol will eventually cease.

Never-Terminating

The protocol will continue for all eternity.

- » All examples of **safety properties**
- » Abstract properties to one general **safety invariant**, φ
- » Typing rules are **parameterised** by φ
 - Enables checking different properties
- » We check that safety holds inspecting **typing contexts**, Γ

Guaranteeing Safety

Ignore \mathcal{R} for now, we'll come back to it later.

We say that φ is an $(s; \mathcal{R})$ -safety property of typing contexts iff, whenever $\varphi(\Gamma)$, we have:

Communication is possible if p and q try to communicate.

$$[S-\oplus\&] \Gamma \xrightarrow{s[p]:q\oplus m(S)} \text{ and } \Gamma \xrightarrow{s[q]:p\&m'(S')} \text{ implies } \Gamma \xrightarrow{s[p][q]m};$$

Guaranteeing Safety

Ignore \mathcal{R} for now, we'll come back to it later.

We say that φ is an $(s; \mathcal{R})$ -safety property of typing contexts iff, whenever $\varphi(\Gamma)$, we have:

Communication is possible if p and q try to communicate.

$$[S-\oplus\&] \Gamma \xrightarrow{s[p]:q\oplus m(S)} \text{ and } \Gamma \xrightarrow{s[q]:p\& m'(S')} \text{ implies } \Gamma \xrightarrow{s[p][q]m};$$

q has a crash handling branch if it receives from an unreliable p .

$$[S-\not\&] \Gamma \xrightarrow{s[p]\text{stop}} \text{ and } \Gamma \xrightarrow{s[q]:p\& m(S)} \text{ implies } \Gamma \xrightarrow{s[q]\odot p};$$

Guaranteeing Safety

Ignore \mathcal{R} for now, we'll come back to it later.

We say that φ is an $(s; \mathcal{R})$ -safety property of typing contexts iff, whenever $\varphi(\Gamma)$, we have:

Communication is possible if p and q try to communicate.

$$[S-\oplus\&] \Gamma \xrightarrow{s[p]:q\oplus m(S)} \text{ and } \Gamma \xrightarrow{s[q]:p\& m'(S')} \text{ implies } \Gamma \xrightarrow{s[p][q]m};$$

q has a crash handling branch if it receives from an unreliable p .

$$[S-\&\&] \Gamma \xrightarrow{s[p]\text{stop}} \text{ and } \Gamma \xrightarrow{s[q]:p\& m(S)} \text{ implies } \Gamma \xrightarrow{s[q]\odot p};$$

A safe Γ cannot beget an unsafe Γ' .

$$[S-\rightarrow\&] \Gamma \xrightarrow{\& \setminus s; \mathcal{R}} \Gamma' \text{ implies } \varphi(\Gamma').$$

Safety Checking the DNS Protocol

$$\Gamma = \Gamma_p, \Gamma_q, \Gamma_r$$

$$\Gamma_p = s[p]:q \oplus \text{url}.q \& \left\{ \begin{array}{l} \text{res.end} \\ \text{crash}.r \oplus \text{url}.r \& \text{res.end} \end{array} \right\}$$

$$\Gamma_q = s[q]:\text{stop}$$

$$\Gamma_r = s[r]:q \& \text{crash}.p \& \text{url}.p \oplus \text{res.end}$$

By inspection, we know that Γ is safe

Safety Checking the DNS Protocol

$$\begin{aligned}\Gamma &= \Gamma_p, \Gamma_q, \Gamma_r \\ \Gamma_p &= s[p]:q \oplus \text{url}.q \& \left\{ \begin{array}{l} \text{res.end} \\ \text{crash}.r \oplus \text{url}.r \& \text{res.end} \end{array} \right\} \\ \Gamma_q &= s[q]:\text{stop} \\ \Gamma_r &= s[r]:q \& \text{crash}.p \& \text{url}.p \oplus \text{res.end}\end{aligned}$$

By inspection, we know that Γ is safe

$$\begin{array}{lll} [s-\oplus\&] & s[p]:q \oplus \text{url} \dots & s[p]:r \& \text{res.end} \\ & s[q]:p \& \text{url} \dots & s[p]:p \oplus \text{res.end} \end{array}$$

Safety Checking the DNS Protocol

$$\begin{aligned}\Gamma &= \Gamma_p, \Gamma_q, \Gamma_r \\ \Gamma_p &= s[p]:q \oplus \text{url}.q \& \left\{ \begin{array}{l} \text{res.end} \\ \text{crash}.r \oplus \text{url}.r \& \text{res.end} \end{array} \right\} \\ \Gamma_q &= s[q]:\text{stop} \\ \Gamma_r &= s[r]:q \& \text{crash}.p \& \text{url}.p \oplus \text{res.end}\end{aligned}$$

By inspection, we know that Γ is safe

$$\begin{array}{l} [s-\&] \quad s[p]:q \& \left\{ \begin{array}{l} \dots \\ \text{crash}\dots \end{array} \right\} \\ s[q]:\text{stop} \end{array}$$

Safety Checking the DNS Protocol

$$\Gamma = \Gamma_p, \Gamma_q, \Gamma_r$$

$$\Gamma_p = s[p]:q \oplus \text{url}.q \& \left\{ \begin{array}{l} \text{res.end} \\ \text{crash}.r \oplus \text{url}.r \& \text{res.end} \end{array} \right\}$$

$$\Gamma_q = s[q]:\text{stop}$$

$$\Gamma_r = s[r]:q \& \text{crash}.p \& \text{url}.p \oplus \text{res.end}$$

By inspection, we know that Γ is safe

» By further inspection, Γ deadlock-free, live, and terminating.

Safety Checking the DNS Protocol

$$\begin{aligned}\Gamma &= \Gamma_p, \Gamma_q, \Gamma_r \\ \Gamma_p &= s[p]:q \oplus \text{url}.q \& \left\{ \begin{array}{l} \text{res.end} \\ \text{crash}.r \oplus \text{url}.r \& \text{res.end} \end{array} \right\} \\ \Gamma_q &= s[q]:\text{stop} \\ \Gamma_r &= s[r]:q \& \text{crash}.p \& \text{url}.p \oplus \text{res.end}\end{aligned}$$

By inspection, we know that Γ is safe

- » By further inspection, Γ deadlock-free, live, and terminating.
- » How do we determine this **automatically**?

Type-Level Model Checking

Desirable behavioural properties can be checked *via* **model checking**.

- » Off-the-shelf model checker: **mCRL2**
- » Typing contexts are expressed as **models**
- » Behavioural properties are expressed as **modal μ -calculus formulæ**

Type-Level Model Checking

Desirable behavioural properties can be checked *via* **model checking**.

- » Off-the-shelf model checker: **mCRL2**
- » Typing contexts are expressed as **models**
- » Behavioural properties are expressed as **modal μ -calculus formulæ**

$$\begin{aligned} [S-\oplus\&] \quad \Gamma \xrightarrow{s[p]:q\oplus m(S)} \text{ and } \Gamma \xrightarrow{s[q]:p\& m'(S')} \text{ implies } \Gamma \xrightarrow{s[p][q]m}; \\ [S-\&] \quad \Gamma \xrightarrow{s[p]\text{stop}} \text{ and } \Gamma \xrightarrow{s[q]:p\& m(S)} \text{ implies } \Gamma \xrightarrow{s[q]\odot p}; \\ [S-\rightarrow] \quad \Gamma \xrightarrow{\& \setminus s; \mathcal{R}} \Gamma' \text{ implies } \varphi(\Gamma'). \end{aligned}$$

Type-Level Model Checking

Desirable behavioural properties can be checked *via* **model checking**.

- » Off-the-shelf model checker: **mCRL2**
- » Typing contexts are expressed as **models**
- » Behavioural properties are expressed as **modal μ -calculus formulæ**

$$\Gamma \models \nu Z. \left(\begin{array}{l} \forall s, p, q, m, m', S, S'. \left(\langle s[p] \text{stop} \rangle_T \wedge \langle s[q]:p\&m'(S') \rangle_T \Rightarrow \langle s[q] \odot p \rangle_T \right) \\ \wedge \left(\langle s[p]:q \oplus m(S) \rangle_T \wedge \left(\langle s[q]:p\&m'(S') \rangle_T \vee \langle s[q] \text{stop} \rangle_T \right) \Rightarrow \langle s[p][q]m \rangle_T \right) \wedge \phi_{\rightarrow}(Z) \end{array} \right)$$

\downarrow $[S-\&]$

\uparrow $[S-\oplus\&]$

Type-Level Model Checking is Practical

example	states	transitions	safe	df	live	nterm	term
DNS	101	427	12.28 ± 1%	17.14 ± 1%	11.24 ± 1%	15.47 ± 0%	12.33 ± 0%
Adder	37	159	12.43 ± 0%	15.74 ± 0%	12.24 ± 1%	14.46 ± 0%	12.06 ± 1%
TwoBuyers	1409	10248	45.6 ± 0%	88.26 ± 0%	31.33 ± 0%	77.2 ± 0%	45.65 ± 0%
Negotiate	1089	8106	34.61 ± 0%	55.07 ± 0%	25.69 ± 0%	47.46 ± 0%	26.04 ± 0%
Broadcast	161	925	17.99 ± 1%	28.13 ± 0%	14.08 ± 0%	25.72 ± 1%	17.74 ± 0%

- » Property checking takes at < 100ms on our example suite
- » Our prototype tool, MPSTK, is available on GitHub at <https://github.com/alcestes/mpstk-crash-stop>

Type-Level Model Checking is Practical

example	states	transitions	safe	df	live	nterm	term
DNS	101	427	12.28 ± 1%	17.14 ± 1%	11.24 ± 1%	15.47 ± 0%	12.33 ± 0%
Adder	37	159	12.43 ± 0%	15.74 ± 0%	12.24 ± 1%	14.46 ± 0%	12.06 ± 1%
TwoBuyers	1409	10248	45.6 ± 0%	88.26 ± 0%	31.33 ± 0%	77.2 ± 0%	45.65 ± 0%
Negotiate	1089	8106	34.61 ± 0%	55.07 ± 0%	25.69 ± 0%	47.46 ± 0%	26.04 ± 0%
Broadcast	161	925	17.99 ± 1%	28.13 ± 0%	14.08 ± 0%	25.72 ± 1%	17.74 ± 0%

- » Property checking takes at < 100ms on our example suite
- » Our prototype tool, MPSTK, is available on GitHub at <https://github.com/alcestes/mpstk-crash-stop>

Type-Level Model Checking is Practical

example	states	transitions	safe	df	live	nterm	term
DNS	101	427	12.28 ± 1%	17.14 ± 1%	11.24 ± 1%	15.47 ± 0%	12.33 ± 0%
Adder	37	159	12.43 ± 0%	15.74 ± 0%	12.24 ± 1%	14.46 ± 0%	12.06 ± 1%
TwoBuyers	1409	10248	45.6 ± 0%	88.26 ± 0%	31.33 ± 0%	77.2 ± 0%	45.65 ± 0%
Negotiate	1089	8106	34.61 ± 0%	55.07 ± 0%	25.69 ± 0%	47.46 ± 0%	26.04 ± 0%
Broadcast	161	925	17.99 ± 1%	28.13 ± 0%	14.08 ± 0%	25.72 ± 1%	17.74 ± 0%

- » Property checking takes at < 100ms on our example suite
- » Our prototype tool, MPSTK, is available on GitHub at <https://github.com/alcestes/mpstk-crash-stop>

mCRL2 improvement [POPL'19]

Table 2. Average time (in seconds \pm std. dev.) for the verification of the protocols in Fig. 4. Protocols (3) and (4) are instantiated with $n=3$. The outcome of the verification is shown in Table 1. (Benchmarking specs: Intel Core i7-4790 CPU, 3.60GHz, 16 GB RAM, mCRL2 201808.0 invoked 30 times (by mpstk) with: pbes2bool --strategy=2)


	states	safe	deadlock-free	live	live ⁺	live ⁺⁺	never-terminat.	terminat.
(1) OAuth2 fragment	37	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	0.98 \pm 9%
(2) Rec. two-buyers	85	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	0.99 \pm 3%
(3) Rec. map/reduce	2561	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	0.99 \pm 3%
(4) MP workers	442369	1.01 \pm 4%	0.98 \pm 8%	0.98 \pm 9%	1.03 \pm 14%	1.02 \pm 7%	0.99 \pm 6%	1.00 \pm 1%

Fault-Tolerant Session Types

We present a **generalised multiparty session type theory** with **crash-stop failures**.

- ✓ Session π -calculus with crash-stop failures
 - ✓ Multiparty session types with crashes
 - ✓ Type-level model checking for behavioural properties
- ⇒ Optional reliability assumptions

One More Step Towards Realistic Procotols



- » Some participants may represent services that can be **assumed reliable**
 - À la Amazon AWS, Microsoft Azure, Google Cloud Platform, &c.
- » Can we **safely** avoid writing crash handling branches for **reliable** participants?

Optional Reliability Assumptions

- » Specific participants can be marked as **reliable**
- » Reliable participants **will not crash**
- » The set of reliable participants is denoted \mathcal{R}
 - As seen in the safety definition
- » Reduction rules, model checking, &c. conform to these assumptions
 - Behavioural properties can be satisfied for partially reliable protocols

Optionally Assuming Reliability in the DNS Protocol

$$\Gamma = \Gamma_p, \Gamma_q, \Gamma_r$$

$$\Gamma_p = s[p]:q \oplus \text{url}.q \& \left\{ \begin{array}{l} \text{res.end} \\ \text{crash}.r \oplus \text{url}.r \& \text{res.end} \end{array} \right\}$$

$$\Gamma_q = s[q]:\text{stop}$$

$$\Gamma_r = s[r]:q \& \text{crash}.p \& \text{url}.p \oplus \text{res.end}$$

If we assume that **all participants are unreliable**

- » Both p and r can crash arbitrarily
- » Currently missing crash handling branches...
- » The safety property **does not hold**

Optionally Assuming Reliability in the DNS Protocol

$$\Gamma = \Gamma_p, \Gamma_q, \Gamma_r$$

$$\Gamma_p = s[p]:q \oplus \text{url}.q \& \left\{ \begin{array}{l} \text{res.end} \\ \text{crash.r} \oplus \text{url}.r \& \text{res.end} \end{array} \right\}$$

$$\Gamma_q = s[q]:\text{stop}$$

$$\Gamma_r = s[r]:q \& \text{crash.p} \& \text{url}.p \oplus \text{res.end}$$

If we assume that both **p** and **r** are reliable

- » Both **p** and **r** **cannot** crash
- » No further adjustments necessary
- » Safety, deadlock-freedom, liveness, and terminating properties hold

Summary and Future Work

- » Generalised MPST theory with crash-stop failures
- » Model checking support to check for behavioural properties
 - <https://github.com/alcestes/mpstk-crash-stop>
- » Optional Reliability Assumptions
 - Support for fully-reliable to fully-unreliable protocols

In the paper

- » Type system: typing rules and typing context transitions
- » How optional reliability is respected in considering process reductions
- » How properties are formulated as modal μ -calculus formulæ
- » Benchmarks and examples
- » Full version: <https://arxiv.org/abs/2207.02015>

Summary and Future Work

- » Generalised MPST theory with crash-stop failures
- » Model checking support to check for behavioural properties
 - <https://github.com/alcestes/mpstk-crash-stop>
- » Optional Reliability Assumptions
 - Support for fully-reliable to fully-unreliable protocols

Future Work

- » Support more failure models
- » Support a top-down approach to MPST with global types and projection

Summary and Future Work

- » Generalised MPST theory with crash-stop failures
- » Model checking support to check for behavioural properties
 - <https://github.com/alcestes/mpstk-crash-stop>
- » Optional Reliability Assumptions
 - Support for fully-reliable to fully-unreliable protocols

In the paper

- » Type system: typing rules and typing context transitions
- » How optional reliability is respected in considering process reductions
- » How properties are formulated as modal μ -calculus formulæ
- » Benchmarks and examples
- » Full version: <https://arxiv.org/abs/2207.02015>

Aides Memoire: Affine Approaches

- » *Affine Sessions* by Mostrous & Vasconcelos, 2018
 - Binary sessions
- » *Exceptional Asynchronous Session Types* by Fowler et al. 2019
 - λ -calculus
 - Binary sessions
 - Exception handling
- » More exceptions: Cappechi et al. 2016 & Carbone et al. 2008
- » This approach models failures at the application level via throw/catch constructs.
 1. We model arbitrary failures (e.g. hardware failures);
 2. We specify what to do when a failure is detected at the type level;
 3. We support multiparty sessions;
 4. We support a participant's crash whilst handling another's (the do-catch constructs cannot be nested).

Aides Memoire: Coordinator Model Approaches

- » *Session Types for Link Failures* by Adameit et al., 2017
 - Optional blocks and default values
- » *Robust Failure Handling in Distributed Systems* by Chen et al., 2016
 - Synchronisation points to detect and handle failures
- » *Statically Verified Crash Failure Handling* by Viering et al., 2018
 - Try-handle construct, a reliable coordinator, and broadcast failures
- » We do not assume reliable processes, failure broadcasts, or coordination/synchronisation points.

Aides Memoire: Other Approaches


- » *Fault-Tolerant Multiparty Session Types* by Peters et al. 2022
 - Their failure model is different from ours; and they handle failures by continuing the protocol via default branches and values.
 - **crash** is not synonymous to a default branch
- » *Let it Recover* by Neykova & Yoshida, 2017
 - Supervision trees approach to MPST à la Erlang
- » *Fault-Tolerant Event-Driven Distributed Programming* by Viering et al. 2018
 - Processes are monitored and restarted if they fail; needs reliable participants; but they do tolerate false suspicions
- » *Linearity, Control Effects, and Behavioural Types* by Caires et al. 2017
 - Curry-Howard interpretation; binary session types; failures are propagated to relevant peers
- » *Localities* approaches
 - Model distributed systems with failures; no types

mCRL2 and Multiparty Session Types

- How can we integrate local verifications (by mCRL2) and global protocol conformance?
- What are possible extensions of local verifications of session types?
 1. probabilities
 2. counter-example guided verification
 3. bisimulation checking
 4. asynchrony
 5. mCLR2 graphics tool



Thank you! Questions?



CFSMs [1980-] ITU notation SDL · MSCS ...

Def A CFSM $M = (Q, C, q_0, \Sigma, \delta)$

Q a finite set of states

$C = \{ pq \in \text{Participant}^2 \mid p \neq q \}$



q_0 initial state

Σ a finite alphabet of messages



$\delta \subseteq Q \times (C \times \{!, ?\} \times \Sigma) \times Q$ a finite set of transitions

Def CS $S = (M_p)_{p \in \text{Participant}}$

Multiparty Compatibility



Multiparty Compatibility



Def $S = (M_p)_{p \in \text{Participant}}$

$\forall s . s_0 \rightsquigarrow^* s$
1-buffer execution

if M_i does action l

then $(M_{\bar{j}})_{\bar{j} \in P \setminus i}$ do action \bar{l}
 after some \rightsquigarrow^*

Multiparty Compatibility

Definition System $S = (M_p)_{p \in \mathcal{P}}$ is **MC** if for any 1-bound reachable state $s \in RS_1(S)$, and any output action $pq!a$ from s in M_p , there exists an alternation $\varphi.t$ from s in a system where $\text{act}(t) = pq!a$ and $p \notin \text{act}(\varphi)$


(Dual for input)

$S \xrightarrow{t} S'$ configuration $S = (\vec{q}; \vec{w})$
states queues

Send
 $(\dots q_p \dots; \dots W_{pq} \dots) \xrightarrow{pq!l} (\dots q'_p; \dots W_{pq} \cdot l \dots)$
 q_p W_{pq}


Receive
 $(\dots q_q \dots; \dots l \cdot W_{pq} \dots) \xrightarrow{pq?l} (\dots q'_q \dots; \dots W_{pq} \dots)$

Deterministic CFM



Basic CFSMs


A CFSM is **Basic** if **deterministic**
directed, has **no mixed states**



sending




$T = A!\{a, b, c\}$




receiving



$A?\{a, b, c\}$



non
directed



mixed



