

Semantics of swarms

By rule [Local] below, a command's execution updates both local and global logs

$$\frac{\mathbf{S}(i) = \mathbf{M}_{\ell_i} \quad \mathbf{M}_{\ell_i} \xrightarrow{c/l} \mathbf{M}_{\ell'_i} \quad \text{src}(\ell'_i \setminus \ell_i) = \{i\} \quad \ell' \in \ell \bowtie \ell'_i}{(\mathbf{S}, \ell) \xrightarrow{c/l} (\mathbf{S}[i \mapsto \mathbf{M}_{\ell'_i}], \ell')} \text{[Local]}$$

$$\frac{\mathbf{S}(i) = \mathbf{M}_{\ell_i} \quad \ell_i \sqsubseteq \ell' \sqsubseteq \ell \quad \ell_i \subset \ell'}{(\mathbf{S}, \ell) \xrightarrow{\tau} (\mathbf{S}[i \mapsto \mathbf{M}_{\ell'_i}], \ell)} \text{[Prop]}$$

By rule [Prop] above, the propagation of events happens

- ▶ by shipping a **non-deterministically chosen** subset of events in the global log
- ▶ to a **non-deterministically chosen** machine

Semantics at work (I)

If

$$B \boxed{b} \xrightarrow{c/1} B \boxed{b \cdot d \cdot e} \quad \text{with} \quad \vdash d \cdot e : 1$$

Semantics at work (I)

If $B \boxed{b} \xrightarrow{c/1} B \boxed{b \cdot d \cdot e}$ with $\vdash d \cdot e : 1$

then, by [Local] $A \boxed{a} \mid B \boxed{b} \mid C \boxed{c} \mid b \cdot a \cdot c \xrightarrow{c/1} A \boxed{a} \mid B \boxed{b \cdot d \cdot e} \mid C \boxed{c} \mid \ell$

Semantics at work (I)

If $B \boxed{b} \xrightarrow{c/1} B \boxed{b \cdot d \cdot e}$ with $\vdash d \cdot e : 1$

then, by [Local] $A \boxed{a} \mid B \boxed{b} \mid C \boxed{c} \mid b \cdot a \cdot c \xrightarrow{c/1} A \boxed{a} \mid B \boxed{b \cdot d \cdot e} \mid C \boxed{c} \mid \ell$

for all $\ell \in (b \cdot a \cdot c) \bowtie (b \cdot d \cdot e)$

Semantics at work (I)

If $B \boxed{b} \xrightarrow{c/1} B \boxed{b \cdot d \cdot e}$ with $\vdash d \cdot e : 1$

then, by [Local]

$$A \boxed{a} \mid B \boxed{b} \mid C \boxed{c} \mid b \cdot a \cdot c \xrightarrow{c/1} A \boxed{a} \mid B \boxed{b \cdot d \cdot e} \mid C \boxed{c} \mid \ell$$

for all

$$\ell \in (b \cdot a \cdot c) \bowtie (b \cdot d \cdot e)$$

Exercise

Compute $(b \cdot a \cdot c) \bowtie (b \cdot d \cdot e)$

Semantics at work (I)

If $B \boxed{b} \xrightarrow{c/1} B \boxed{b \cdot d \cdot e}$ with $\vdash d \cdot e : 1$

then, by [Local] $A \boxed{a} \mid B \boxed{b} \mid C \boxed{c} \mid b \cdot a \cdot c \xrightarrow{c/1} A \boxed{a} \mid B \boxed{b \cdot d \cdot e} \mid C \boxed{c} \mid \ell$

for all $\ell \in (b \cdot a \cdot c) \bowtie (b \cdot d \cdot e)$

Exercise

Compute $(b \cdot a \cdot c) \bowtie (b \cdot d \cdot e)$

$$\left\{ \begin{array}{l} b \cdot a \cdot c \cdot d \cdot e, \quad b \cdot a \cdot d \cdot c \cdot e, \quad b \cdot a \cdot d \cdot e \cdot c, \\ b \cdot d \cdot a \cdot e \cdot c, \quad b \cdot d \cdot a \cdot c \cdot e, \quad b \cdot d \cdot e \cdot a \cdot c \end{array} \right\}$$

Semantics at work (II)

Exercise

Is $A[a] \mid B[b] \mid C[\] \mid b \xrightarrow{\tau} A[a] \mid B[b] \mid C[\] \mid a \cdot b$ possible? Does it actually make sense?

Semantics at work (II)

Exercise

Is $A[a] | B[b] | C[] | b \xrightarrow{\tau} A[a] | B[b] | C[] | a \cdot b$ possible? Does it actually make sense?

With reference to slide 56, consider

$$A[a] | B[b] | C[c] | b \cdot a \cdot c \xrightarrow{c/1} A[a] | B[b \cdot d \cdot e] | C[c] | \overbrace{b \cdot a \cdot d \cdot e \cdot c}^{=l}$$

Semantics at work (II)

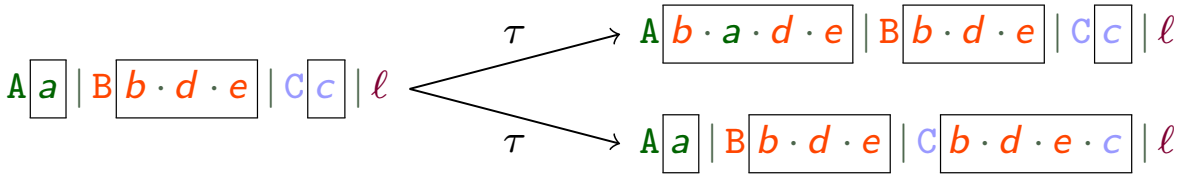
Exercise

Is $A[a] | B[b] | C[] | b \xrightarrow{\tau} A[a] | B[b] | C[] | a \cdot b$ possible? Does it actually make sense?

With reference to slide 56, consider

$$A[a] | B[b] | C[c] | b \cdot a \cdot c \xrightarrow{c/1} A[a] | B[b \cdot d \cdot e] | C[c] | \overbrace{b \cdot a \cdot d \cdot e \cdot c}^{=l}$$

Let's propagate $d \cdot e$: by [Prop] we have to non-deterministically choose a sublog of l



b is shipped too 'cause sublogs must be downward wrt emitting machines

Semantics at work (II)

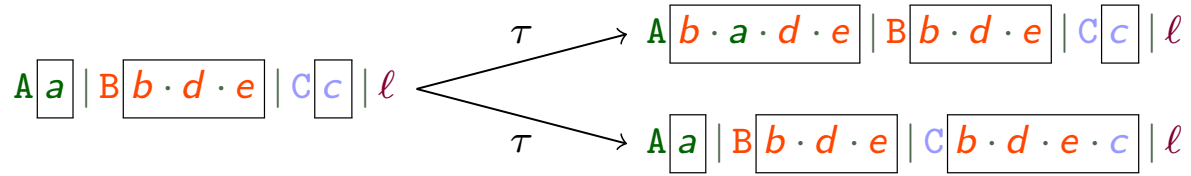
Exercise

Is $A[a] | B[b] | C[] | b \xrightarrow{\tau} A[a] | B[b] | C[] | a \cdot b$ possible? Does it actually make sense?

With reference to slide 56, consider

$$A[a] | B[b] | C[c] | b \cdot a \cdot c \xrightarrow{c/1} A[a] | B[b \cdot d \cdot e] | C[c] | \overbrace{b \cdot a \cdot d \cdot e \cdot c}^{=l}$$

Let's propagate $d \cdot e$: by [Prop] we have to non-deterministically choose a sublog of l



b is shipped too 'cause sublogs must be downward wrt emitting machines

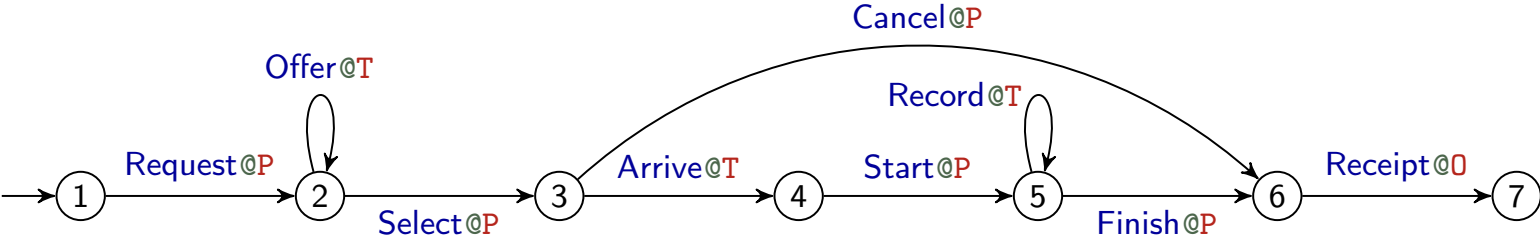
Exercise

Can we propagate just event e ?

Can event a be shipped to C together with events from B ?

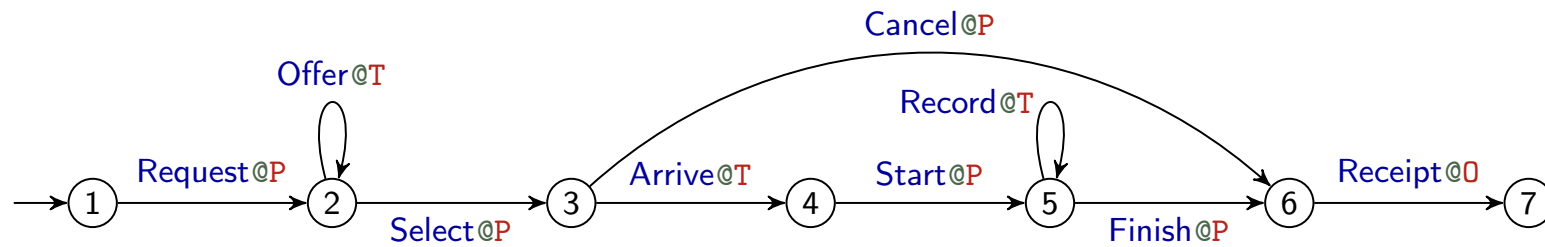
A taxi service

An intuitive auction protocol for a passenger **P** to get a taxi **T**:



A taxi service

An intuitive auction protocol for a passenger P to get a taxi T :

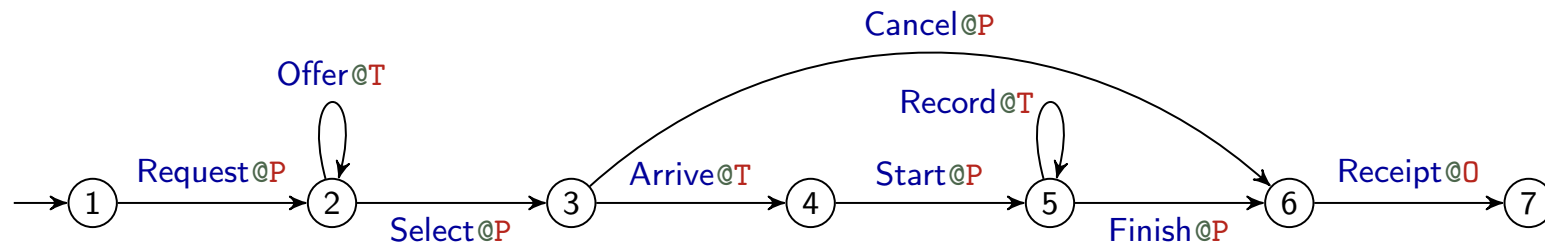


We assume

- ▶ one passenger and one office (for simplicity)

A taxi service

An intuitive auction protocol for a passenger P to get a taxi T :

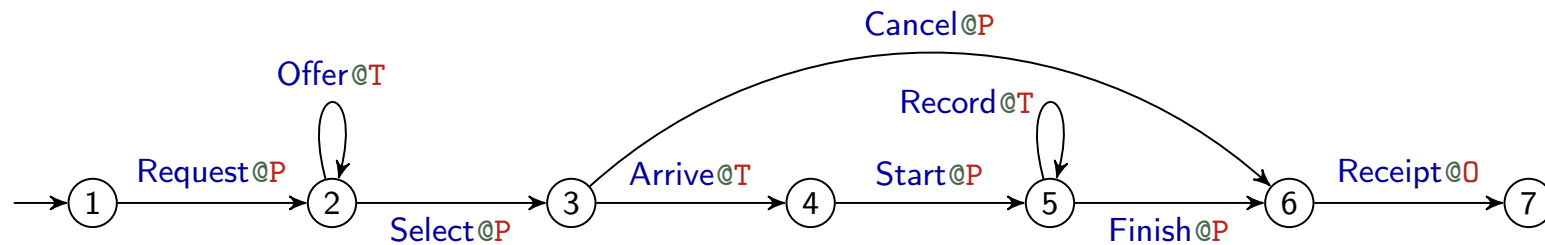


We assume

- ▶ one passenger and one office (for simplicity)
- ▶ but an arbitrary number of taxis

A taxi service

An intuitive auction protocol for a passenger P to get a taxi T :



We assume

- ▶ one passenger and one office (for simplicity)
- ▶ but an arbitrary number of taxis
- ▶ a receipt is issued by the office O at the end of the ride (if any)

Swarm protocols: global type for local-first applications

An **idealised** specification relying on **synchronous communication**

Fix a set of **roles** ranged over by \mathbf{R} (e.g., \mathbf{P} , \mathbf{T} , and \mathbf{O})

The syntax of **swarm protocols** is again given co-inductively:

$$\mathbf{G} ::=^{\text{co}} \sum_{i \in I} c_i @ \mathbf{R}_i \langle \mathbf{1}_i \rangle . \mathbf{G}_i \quad | \quad 0 \quad \text{where } I \text{ is a finite set (of indexes)}$$

An example

A swarm protocol for the taxi scenario:

$$G = \text{Request@P}\langle \text{Requested} \rangle . G_{\text{auction}}$$

$$G_{\text{auction}} = \text{Offer@T}\langle \text{Bid} \cdot \text{BidderID} \rangle . G_{\text{auction}} \\ + \text{Select@P}\langle \text{Selected} \cdot \text{PassengerID} \rangle . G_{\text{choose}}$$

$$G_{\text{choose}} = \text{Arrive@T}\langle \text{Arrived} \rangle . \text{Start@P}\langle \text{Started} \rangle . G_{\text{ride}} \\ + \text{Cancel@P}\langle \text{Cancelled} \rangle . \text{Receipt@O}\langle \text{Receipt} \rangle . 0$$

$$G_{\text{ride}} = \text{Record@T}\langle \text{Path} \rangle . G_{\text{ride}} \\ + \text{Finish@P}\langle \text{Finished} \cdot \text{Rating} \rangle . \text{Receipt@O}\langle \text{Receipt} \rangle . 0$$

An example

A swarm protocol for the taxi scenario:

$$G = \text{Request@P}\langle \text{Requested} \rangle . G_{\text{auction}}$$

$$G_{\text{auction}} = \text{Offer@T}\langle \text{Bid} \cdot \text{BidderID} \rangle . G_{\text{auction}} \\ + \text{Select@P}\langle \text{Selected} \cdot \text{PassengerID} \rangle . G_{\text{choose}}$$

$$G_{\text{choose}} = \text{Arrive@T}\langle \text{Arrived} \rangle . \text{Start@P}\langle \text{Started} \rangle . G_{\text{ride}} \\ + \text{Cancel@P}\langle \text{Cancelled} \rangle . \text{Receipt@O}\langle \text{Receipt} \rangle . 0$$

$$G_{\text{ride}} = \text{Record@T}\langle \text{Path} \rangle . G_{\text{ride}} \\ + \text{Finish@P}\langle \text{Finished} \cdot \text{Rating} \rangle . \text{Receipt@O}\langle \text{Receipt} \rangle . 0$$

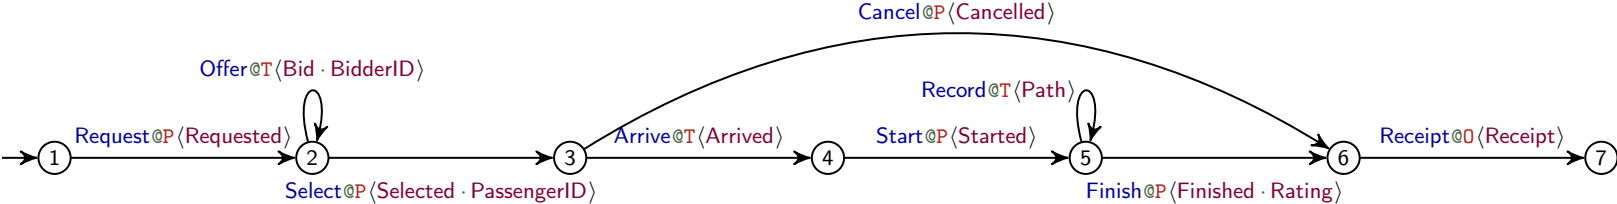
*Notice the log
types in each
prefixes*

Swarm protocols as FSA

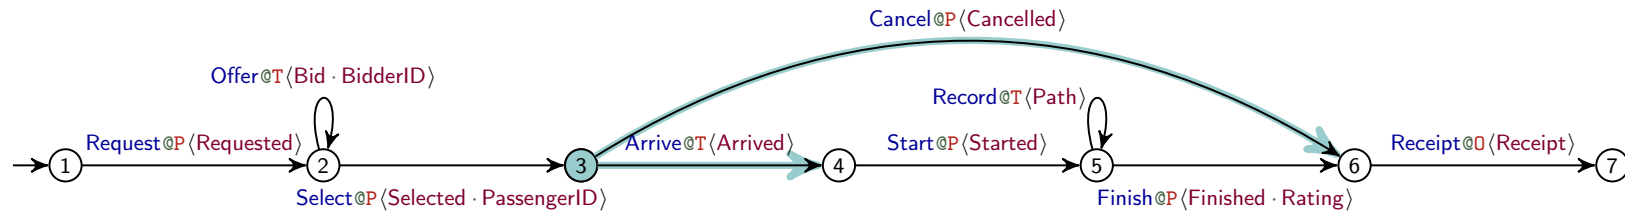
Like for machines, a swarm protocols $G = \sum_{i \in I} c_i @ R_i \langle \mathbf{1}_i \rangle$. G_i has an associated FSA:

- ▶ the set of states consists of G plus the states in G_i for each $i \in I$
- ▶ G is the initial state
- ▶ for each $i \in I$, G has a transition to state G_i labelled with $c_i @ R_i \langle \mathbf{1}_i \rangle$, written $G \xrightarrow{c_i / \mathbf{1}_i} G_i$

An example



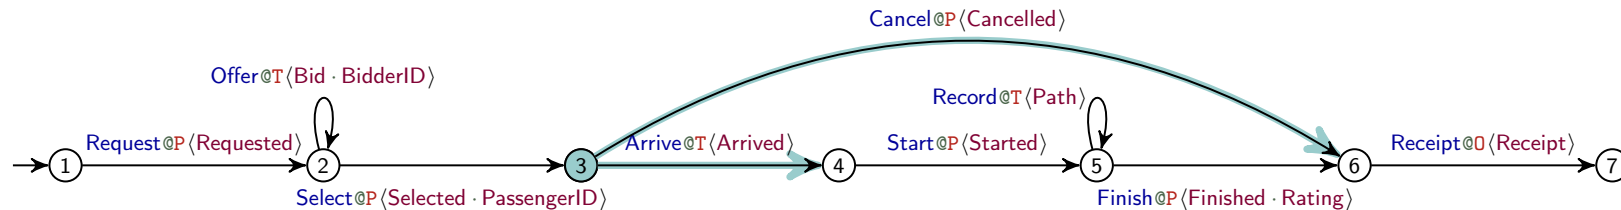
An example



There is a race in state 3!

- ▶ the driver of the selected taxi may invoke **Arrive**
- ▶ **while** **P** loses patience and invokes **Cancel**

An example



There is a race in state 3!

- ▶ the driver of the selected taxi may invoke **Arrive**
- ▶ while **P** loses patience and invokes **Cancel**

This protocol violates well-formedness conditions typically imposed on behavioural types due to the race in state 3 (because it has two active participants, which is also true of states 2 and 5)

Semantics of swarm protocols

One rule only!

$$(G, \ell) \xrightarrow{c/1} (G, \ell) \quad \text{[G-Cmd]}$$

Semantics of swarm protocols

One rule only!

$$\frac{\delta(G, l) \xrightarrow{c/l} G'}{\quad} \text{[G-Cmd]}$$
$$(G, l) \xrightarrow{c/l} (G, l')$$

where

$$\delta(G, l) = \begin{cases} G & \text{if } l = \epsilon \\ \delta(G_1, l_2) & \text{if } G \xrightarrow{c/l} G' \text{ and } l = l_1 \cdot l_2 \text{ with } \vdash l_1 : l \\ \perp & \text{otherwise} \end{cases}$$

Logs have to be consumed "atomically", hence $\delta(G, l)$ may be undefined

Semantics of swarm protocols

One rule only!

$$\frac{\delta(\mathbf{G}, \ell) \xrightarrow{c/1} \mathbf{G}' \quad \vdash \ell' : 1 \quad \ell' \text{ log of fresh events}}{(\mathbf{G}, \ell) \xrightarrow{c/1} (\mathbf{G}, \ell \cdot \ell')} \text{[G-Cmd]}$$

where

$$\delta(\mathbf{G}, \ell) = \begin{cases} \mathbf{G} & \text{if } \ell = \epsilon \\ \delta(\mathbf{G}_1, \ell_2) & \text{if } \mathbf{G} \xrightarrow{c/1} \mathbf{G}' \text{ and } \ell = \ell_1 \cdot \ell_2 \text{ with } \vdash \ell_1 : 1 \\ \perp & \text{otherwise} \end{cases}$$

Logs have to be consumed "atomically", hence $\delta(\mathbf{G}, \ell)$ may be undefined

Semantics of swarm protocols

One rule only!

$$\frac{\delta(\mathbf{G}, \ell) \xrightarrow{c/1} \mathbf{G}' \quad \vdash \ell' : 1 \quad \ell' \text{ log of fresh events}}{(\mathbf{G}, \ell) \xrightarrow{c/1} (\mathbf{G}, \ell \cdot \ell')} \text{[G-Cmd]}$$

where

$$\delta(\mathbf{G}, \ell) = \begin{cases} \mathbf{G} & \text{if } \ell = \epsilon \\ \delta(\mathbf{G}_1, \ell_2) & \text{if } \mathbf{G} \xrightarrow{c/1} \mathbf{G}' \text{ and } \ell = \ell_1 \cdot \ell_2 \text{ with } \vdash \ell_1 : 1 \\ \perp & \text{otherwise} \end{cases}$$

Logs have to be consumed "atomically", hence $\delta(\mathbf{G}, \ell)$ may be undefined

From global to local

Transitions of a swarm protocol G are labelled with a role that may invoke the command

From global to local

Transitions of a swarm protocol G are labelled with a role that may invoke the command

Each machine plays one role

From global to local

Transitions of a swarm protocol G are labelled with a role that may invoke the command

Each machine plays one role



Obtain machines by projecting G on each role

From global to local

Transitions of a swarm protocol G are labelled with a role that may invoke the command

Each machine plays one role



Obtain machines by projecting G on each role

First attempt

$$\left(\sum_{i \in I} c_i @_{R_i} \langle l_i \rangle \cdot G_i \right) \downarrow_{\mathbf{R}} = \kappa \cdot [\&_{i \in I} l_i ? G_i \downarrow_{\mathbf{R}}]$$

where $\kappa = \{(c_i / l_i) \mid R_i = \mathbf{R} \text{ and } i \in I\}$

From global to local

Transitions of a swarm protocol G are labelled with a role that may invoke the command

Each machine plays one role



Obtain machines by projecting G on each role

First attempt

$$\left(\sum_{i \in I} c_i @_{R_i} \langle \mathbf{l}_i \rangle \cdot G_i \right) \downarrow_{\mathbf{R}} = \kappa \cdot [\&_{i \in I} \mathbf{l}_i ? G_i \downarrow_{\mathbf{R}}]$$

where $\kappa = \{(c_i / \mathbf{l}_i) \mid R_i = \mathbf{R} \text{ and } i \in I\}$

simple, but

- ▶ projected machines are large in all but the most trivial cases
- ▶ processing **all** events is undesirable: security and efficiency

Another attempt



Let's subscribe to **subscriptions**: maps from roles to sets of event types

*In pub-sub,
processes
subscribe to
"topics"*

Another attempt



Let's subscribe to **subscriptions**: maps from roles to sets of event types

*In pub-sub,
processes
subscribe to
"topics"*

Given $G = \sum_{i \in I} c_i @_{R_i} \langle 1_i \rangle . G_i$, the **projection of G on a role R with respect to subscription σ** is

$$G \downarrow_R^\sigma = \kappa \cdot [\&_{j \in J} \text{filter}(1_j, \sigma(R)) ? G_j \downarrow_R^\sigma]$$

where

Another attempt



Let's subscribe to **subscriptions**: maps from roles to sets of event types

*In pub-sub,
processes
subscribe to
"topics"*

Given $G = \sum_{i \in I} c_i @_{R_i} \langle \mathbf{l}_i \rangle . G_i$, the **projection of G on a role R with respect to subscription σ** is

$$G \downarrow_{\mathbf{R}}^{\sigma} = \kappa \cdot [\&_{j \in J} \text{filter}(\mathbf{l}_j, \sigma(\mathbf{R}))? G_j \downarrow_{\mathbf{R}}^{\sigma}] \quad \text{where}$$

$$\kappa = \{c_i / \mathbf{l}_i \mid R_i = \mathbf{R} \text{ and } i \in I\}$$

$$J = \{i \in I \mid \text{filter}(\mathbf{l}_i, \sigma(\mathbf{R})) \neq \epsilon\}$$

$$\text{filter}(\mathbf{l}, E) = \begin{cases} \epsilon, & \text{if } \mathbf{l} = \epsilon \\ \mathbf{t} \cdot \text{filter}(\mathbf{l}', E) & \text{if } \mathbf{t} \in E \text{ and } \mathbf{l} = \mathbf{t} \cdot \mathbf{l}' \\ \text{filter}(\mathbf{l}, E) & \text{otherwise} \end{cases}$$

Trading consistency for availability has implications:

Causality

Trading consistency for availability has implications:

Propagation of events is non-atomic (cf. rule [Prop])

⇒ differences in how machines perceive the (state of the) computation

Causality

Fix a subscription σ . For each branch $i \in I$ of $G = \sum_{i \in I} c_i @ R_i \langle \mathbf{1}_i \rangle . G_i$

Explicit re-enabling $\sigma(R_i) \cap \mathbf{1}_i \neq \emptyset$

Command causality if R executes a command in G_i
then $\sigma(R) \cap \mathbf{1}_i \neq \emptyset$ and $\sigma(R) \cap \mathbf{1}_i \supseteq \bigcup_{R' \in \sigma G_i} \sigma(R') \cap \mathbf{1}_i$

Each role explicitly process its emitted events on each of the branches. If R is supposed to enable a command enabled after c_i then $\sigma(R)$ contains some event type emitted by c_i

Causality + Determinacy

Trading consistency for availability has implications:

Propagation of events is non-atomic (cf. rule [Prop])

⇒ different roles may take inconsistent decisions

Causality & Determinacy

Fix a subscription σ . For each branch $i \in I$ of $G = \sum_{i \in I} c_i @ R_i \langle \mathbf{l}_i \rangle . G_i$

Explicit re-enabling $\sigma(R_i) \cap \mathbf{l}_i \neq \emptyset$

Command causality if R executes a command in G_i
then $\sigma(R) \cap \mathbf{l}_i \neq \emptyset$ and $\sigma(R) \cap \mathbf{l}_i \supseteq \bigcup_{R' \in \sigma G_i} \sigma(R') \cap \mathbf{l}_i$

Determinacy $R \in \sigma G_i \implies \mathbf{l}_i[0] \in \sigma(R)$

Causality + Determinacy - Confusion

Trading consistency for availability has implications:

Propagation of events is non-atomic (cf. rule [Prop])

⇒ branches unambiguously identified and events emitted on eventually discharged branches ignored

Causality & Determinacy & Confusion freeness

Fix a subscription σ . For each branch $i \in I$ of $G = \sum_{i \in I} c_i @ R_i \langle \mathbf{l}_i \rangle . G_i$

Explicit re-enabling $\sigma(R_i) \cap \mathbf{l}_i \neq \emptyset$

Command causality if R executes a command in G_i
then $\sigma(R) \cap \mathbf{l}_i \neq \emptyset$ and $\sigma(R) \cap \mathbf{l}_i \supseteq \bigcup_{R' \in \sigma G_i} \sigma(R') \cap \mathbf{l}_i$

Determinacy $R \in \sigma G_i \implies \mathbf{l}_i[0] \in \sigma(R)$

Confusion freeness there is a unique subtree G' of G emitting t
for each t starting a log emitted by a command in G

On correctness



(S, ℓ) faithfully implements G if it produces only logs possibly generated by G

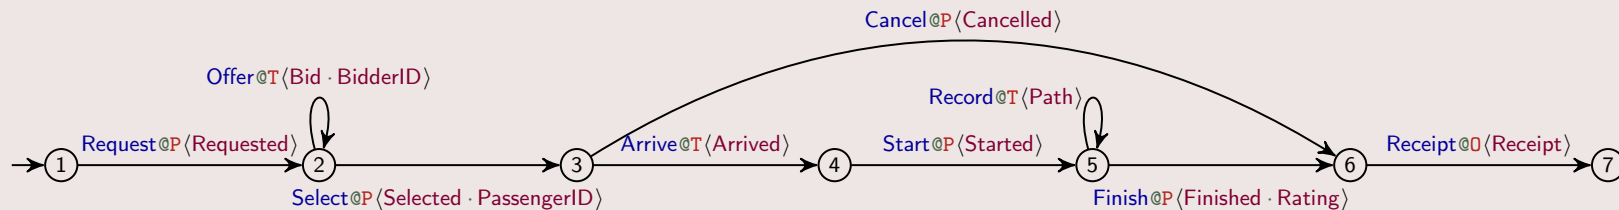
On correctness



(S, ℓ) faithfully implements G if it produces only logs possibly generated by G

Exercise

Take the swarm $S = P \parallel T \parallel O \parallel T$ implementing our taxi protocol



Check that S generates the log

$\ell_{\text{auc}} = \text{requested} \cdot \text{bid} \cdot \text{bidderID} \cdot \text{selected} \cdot \text{bid} \cdot \text{bidderID} \cdot \text{passengerID}$

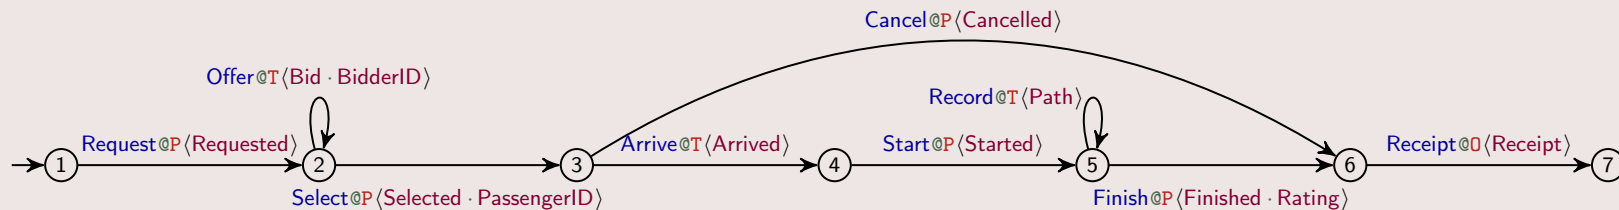
On correctness



(S, ℓ) faithfully implements G if it produces only logs possibly generated by G

Exercise

Take the swarm $S = P \parallel T \parallel O \parallel T$ implementing our taxi protocol



Check that S generates the log

$\ell_{\text{auc}} = \text{requested} \cdot \text{bid} \cdot \text{bidderID} \cdot \text{selected} \cdot \text{bid} \cdot \text{bidderID} \cdot \text{passengerID}$

Too strong a requirement!



Let's consider only "good enough" logs, i.e., those typeable with G 's log types

References I

- [1] Gul Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, MA, USA, 1986.
- [2] Daniel Brand and Pitro Zafiropulo. On Communicating Finite-State Machines. *JACM*, 30(2):323–342, 1983.
- [3] Pierre-Malo Deniélou and Nobuko Yoshida. Multiparty session types meet communicating automata. In *ESOP 2012*, pages 194–213, 2012.
- [4] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, apr 1985.
- [5] Roberto Guanciale and Emilio Tuosto. An abstract semantics of the global view of choreographies. In *Proceedings 9th Interaction and Concurrency Experience, ICE 2016, Heraklion, Greece, 8-9 June 2016.*, pages 67–82, 2016.
- [6] Roberto Guanciale and Emilio Tuosto. Realisability of pomsets. *Journal of Logic and Algebraic Methods in Programming*, 108:69–89, 2019.
- [7] Carl Hewitt, Peter Bohler Bishop, and Richard Steiger. A Universal Modular ACTOR Formalism for Artificial Intelligence. In Nils J. Nilsson, editor, *Proceedings of the 3rd International Joint Conference on Artificial Intelligence. Stanford, CA, USA, August 20-23, 1973*, pages 235–245. William Kaufmann, 1973.
- [8] Nickolas Kavantzas, Davide Burdett, Gregory Ritzinger, Tony Fletcher, and Yves Lafon. Web services choreography description language version 1.0. <http://www.w3.org/TR/2004/WD-ws-cdl-10-20041217>. Working Draft 17 December 2004.

References II

- [9] Roland Kuhn, Hernán C. Melgratti, and Emilio Tuosto. Behavioural types for local-first software. In Karim Ali and Guido Salvaneschi, editors, *37th European Conference on Object-Oriented Programming, ECOOP 2023, July 17-21, 2023, Seattle, Washington, United States*, volume 263 of *LIPICs*, pages 15:1–15:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [10] Julien Lange, Emilio Tuosto, and Nobuko Yoshida. From Communicating Machines to Graphical Choreographies. In *POPL15*, pages 221–232, 2015.
- [11] Emilio Tuosto and Roberto Guanciale. Semantics of global view of choreographies. *Journal of Logic and Algebraic Methods in Programming*, 95:17–40, 2018. Revised and extended version of [5]. available at <http://www.cs.le.ac.uk/people/et52/jlamp-with-proofs.pdf>.