

Local-First Principles: a Behavioural Types Approach

Emilio Tuosto @ GSSI

Roland Kuhn @ Actyx



joint work with
and

Hernán Melgratti @ UBA



Tutorial at Discotec 2023
Lisbon 23 June, 2023

– Prelude –

Take-away message

To trade consistency for availability in systems of **asymmetric replicated peers**

Take-away message

To trade consistency for availability in systems of **asymmetric replicated peers** you can use **local-first**'s principles to (re-)gain **consistency** ... eventually

Take-away message

To trade consistency for availability in systems of **asymmetric replicated peers**

you can use **local-first**'s principles to (re-)gain **consistency** ... eventually

And get some support by our behavioural typing discipline!

Take-away message

To trade consistency for availability in systems of **asymmetric replicated peers**

you can use **local-first**'s principles to (re-)gain **consistency** ... eventually

And get some support by our behavioural typing discipline!



- **swarm protocols**: systems from a **global** viewpoint
- **machines**: peers
- enforce **good behaviour** via behavioural typing

Take-away message

To trade consistency for availability in systems of **asymmetric replicated peers** you can use **local-first**'s principles to (re-)gain **consistency** ... eventually

And get some support by our behavioural typing discipline!




- **swarm protocols**: systems from a **global** viewpoint
- **machines**: peers
- enforce **good behaviour** via behavioural typing



See our recent ECOOP 2023 paper
(to appear; extended version available at
<https://arxiv.org/abs/2305.04848>)

Distributed coordination



An “old” problem

Distributed agreement


Distributed sharing

Security

Computer-assisted collaborative work

...

Distributed coordination



An “old” problem

Distributed agreement

Distributed sharing

Security

Computer-assisted collaborative work

...

With some “solutions”

Centralisation points

Distributed consensus

Commutative replicated data types

...

Autonomy

Thou shall be autonomous

Thou shall collaborate

Thou shall recognise and embrace conflicts

Thou shall resolve conflicts

Thou shall be consistent

Some implications

- peers are not malicious
- peers can progress at all times...even under partial knowledge
- **purity**: inconsistencies resolved by “replaying” executions (invertible or compensatable actions)
- reliable communications

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Alice's mobile

mascarpone cheese

eggs

Bob's mobile

smoked guanciale

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Alice's mobile

mascarpone cheese

eggs

sugar

Bob's mobile

smoked guanciale

eggs

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Alice's mobile

mascarpone cheese

eggs

sugar

Bob's mobile

smoked guanciale

eggs

pecorino romano cheese

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Alice's mobile

mascarpone cheese

eggs

sugar

Bob's mobile

smoked guanciale

eggs

pecorino romano cheese

spaghetti

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Alice's mobile

mascarpone cheese

eggs

sugar

ground moka coffee

Bob's mobile

smoked guanciale

eggs

pecorino romano cheese

spaghetti

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Alice's mobile

mascarpone cheese

eggs

sugar

ground moka coffee

savoiard biscuits

Bob's mobile

smoked guanciale

eggs

pecorino romano cheese

spaghetti

Local-First at work

Alice and Bob decided to have spaghetti carbonara and tiramisù.
They use a mobile app to agree on a grocery list and decide who buys what.

Alice's mobile

mascarpone cheese

eggs

sugar

ground moka coffee

savoiardis biscuits

Bob's mobile

smoked guanciale

eggs

pecorino romano cheese

spaghetti

Eventually the lists can be merged somehow...But who's going to buy the eggs?

Plan of the talk

A motivating case study

Our formalisation

Our typing discipline


Tool support

Open issues

– Motivations –


A collaborative environment and its execution model

(the pictures are courtesy of Actyx AG)




A collaborative environment and its execution model

(the pictures are courtesy of Actyx AG)



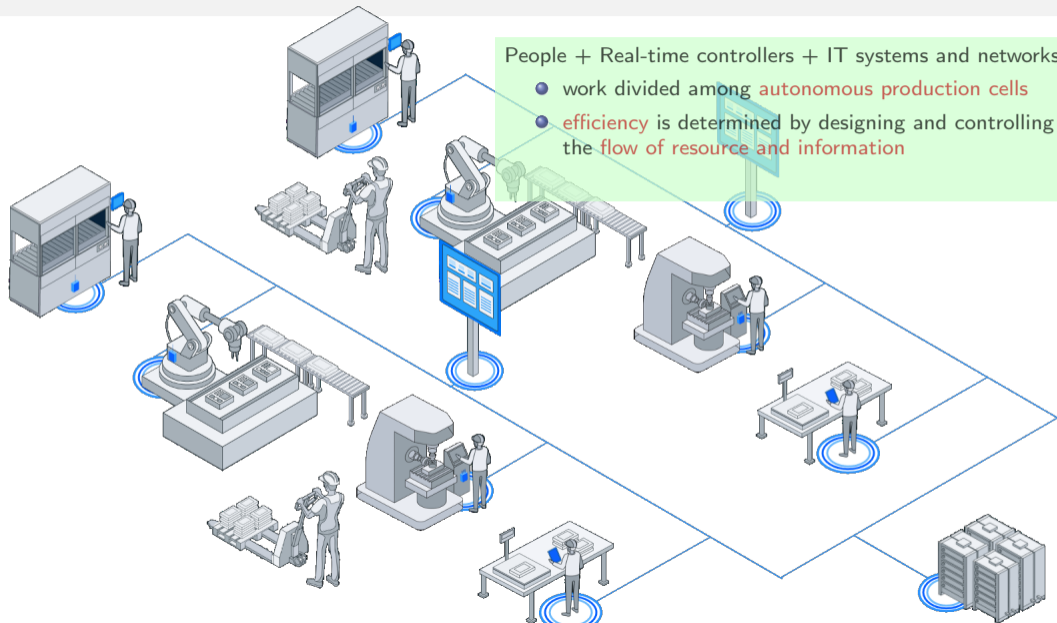
A collaborative environment and its execution model

(the pictures are courtesy of Actyx AG)



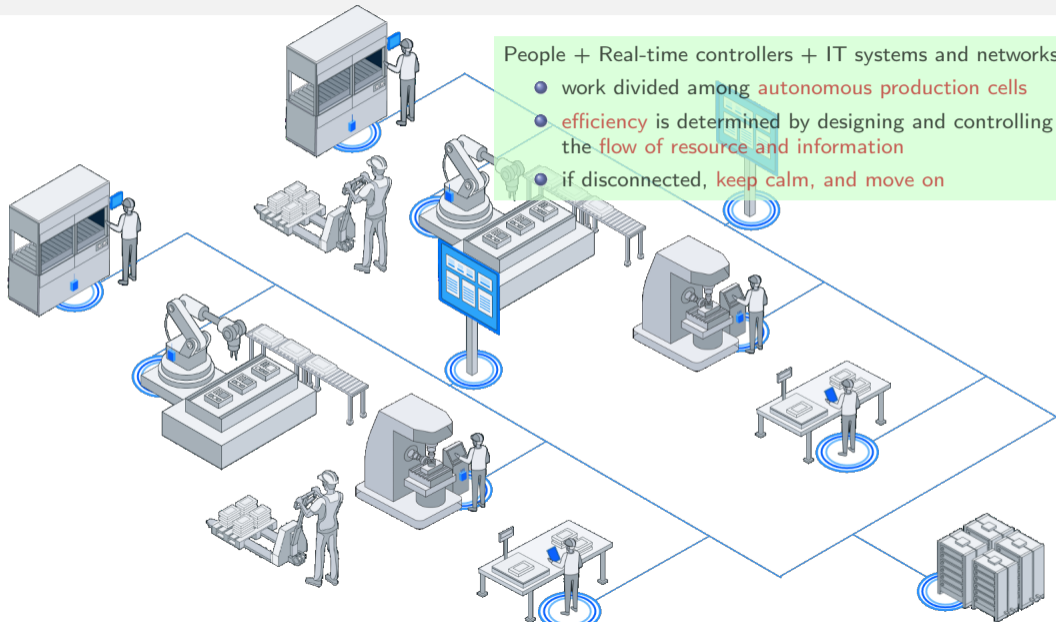
A collaborative environment and its execution model

(the pictures are courtesy of Actyx AG)



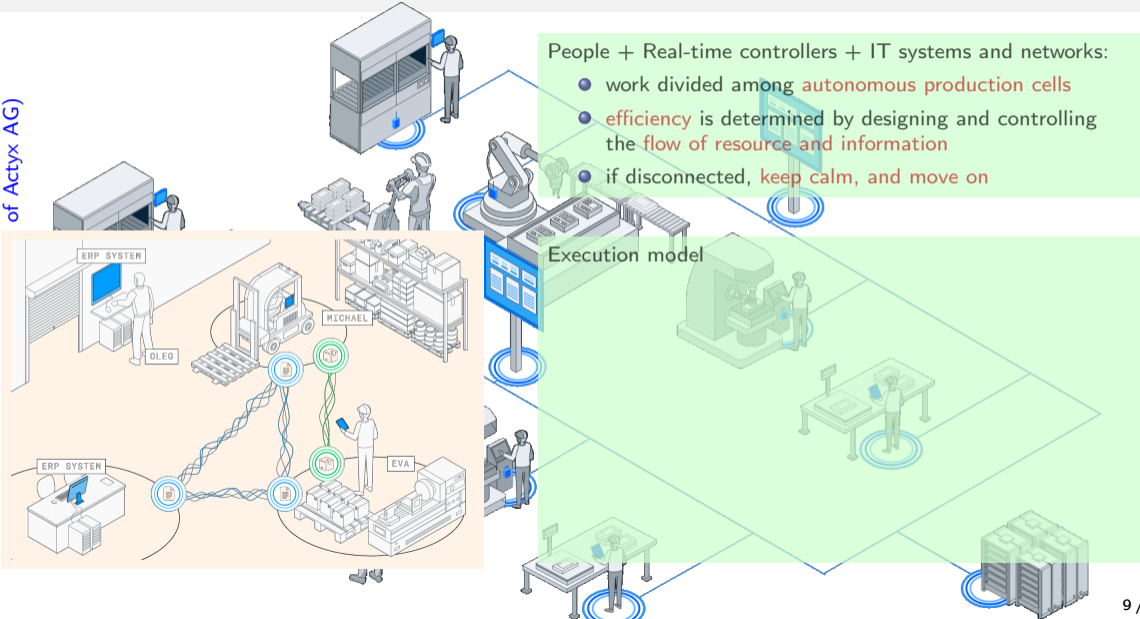
A collaborative environment and its execution model

(the pictures are courtesy of Actyx AG)



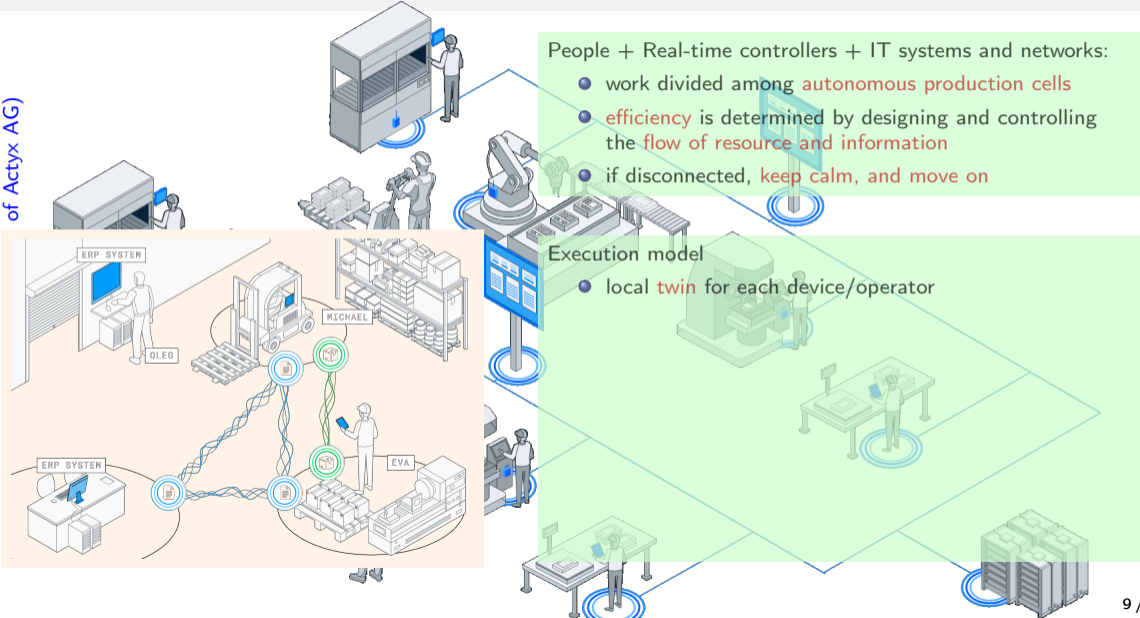
A collaborative environment and its execution model

of Actyx AG)



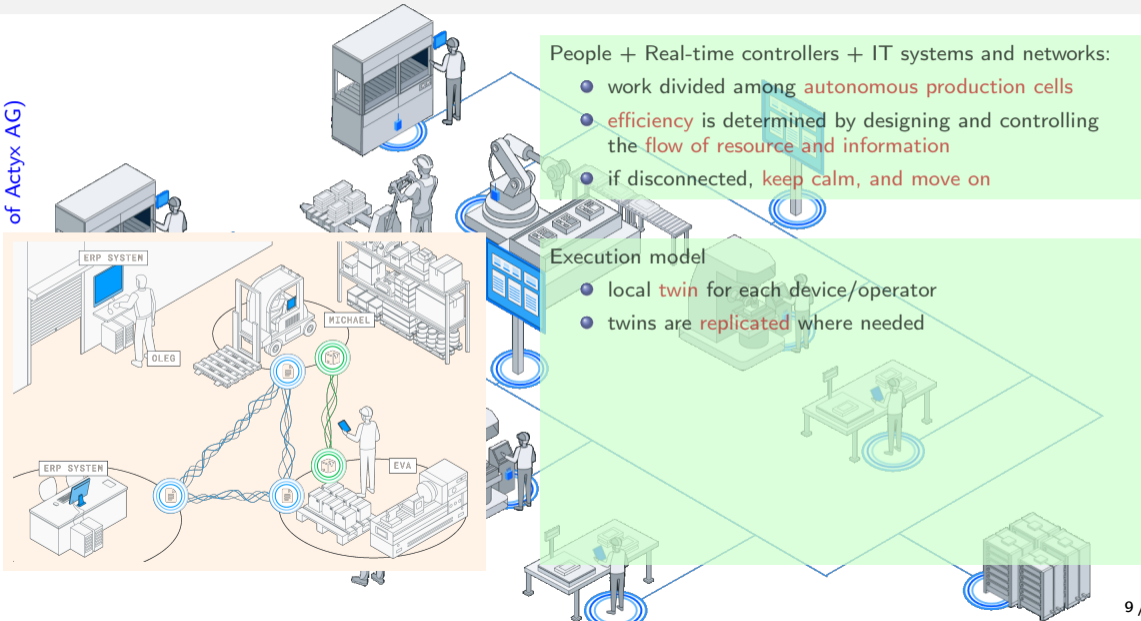
A collaborative environment and its execution model

of Actyx AG)



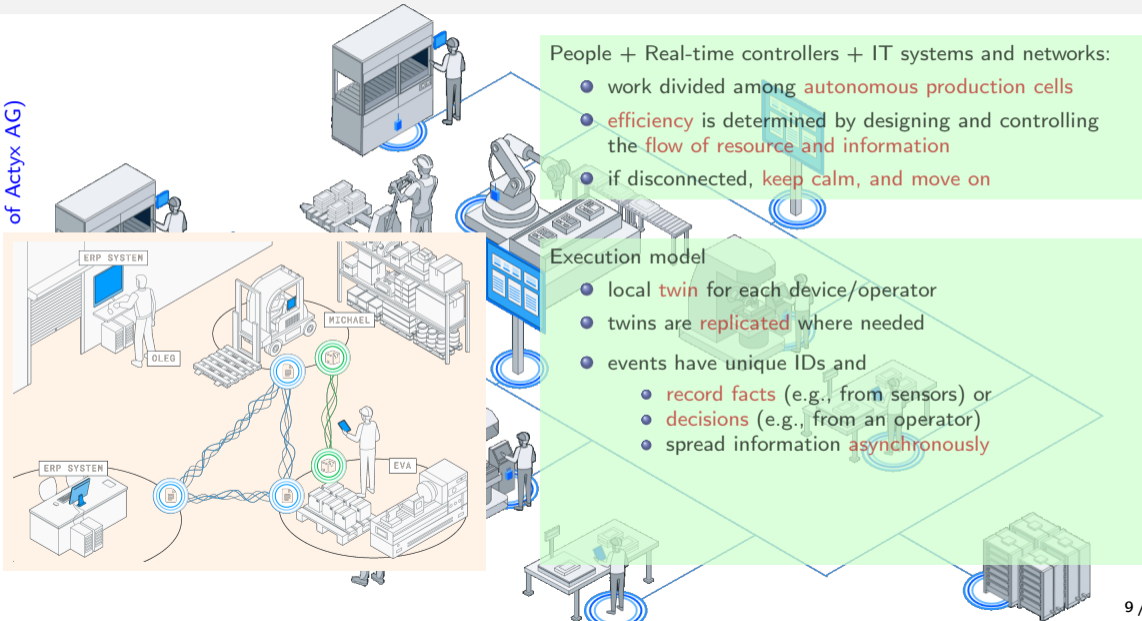
A collaborative environment and its execution model

of Actyx AG)



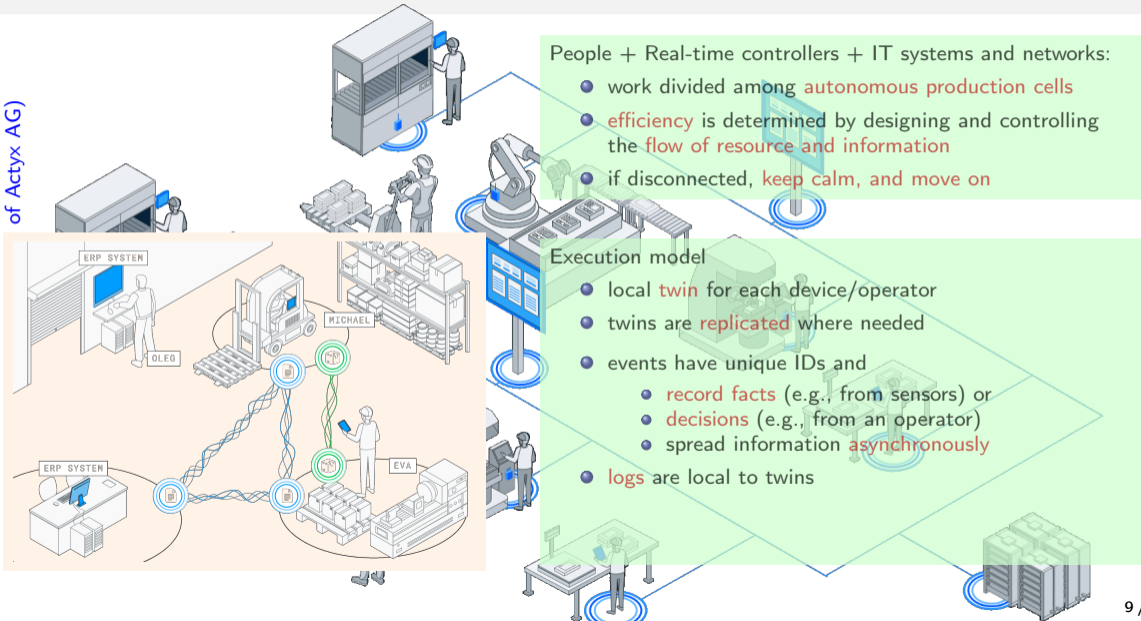
A collaborative environment and its execution model

of Actyx AG)



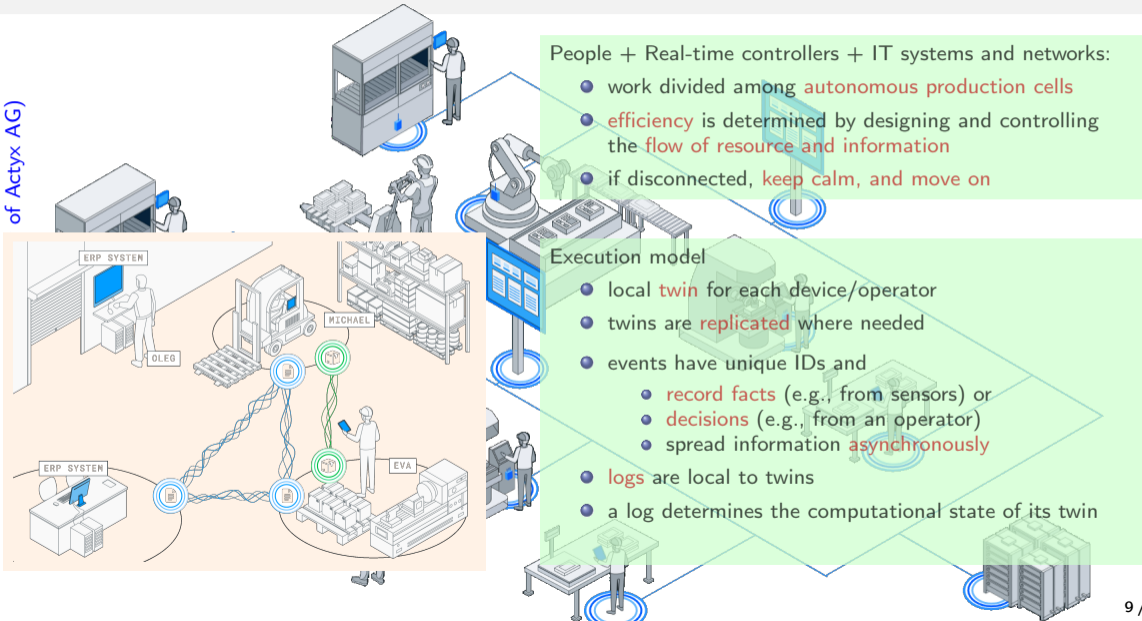
A collaborative environment and its execution model

of Actyx AG)



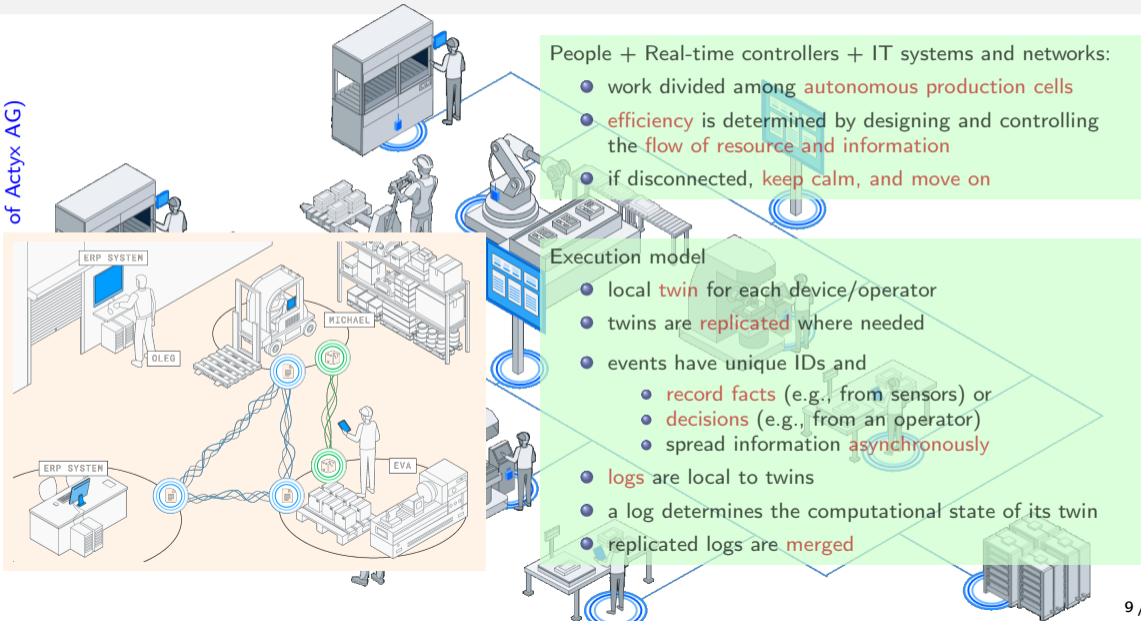
A collaborative environment and its execution model

of Actyx AG)



A collaborative environment and its execution model

of Actyx AG)



execute

+

propagate

+

merge

More applications

Robots (e.g., rescue missions or space applications)

Collaborative applications (<https://automerge.org/>)

Home automation

IoT...really?

Why your fridge and mobile **should go in the cloud** to talk to each other?

Other application domains / motivations

“Anytime, anywhere...” really?

like the AWS's outage on 25/11/2020

or almost all Google services down on 14/12/2020

DSL typical availability of 97% (& some SLA have no **lower bound**)

checkout `https:`

`//www.internetsociety.org/blog/2022/03/what-is-the-digital-divide/`

Other application domains / motivations

Also, taking decisions locally

can reduce downtime

shifts data ownership

gets rid of any centralization point...for real

Challenges

Specify application-level protocols where decisions

- don't require **consensus**

Challenges

Specify application-level protocols where decisions

- don't require **consensus**
- are based on **stale local states**

Challenges

Specify application-level protocols where decisions

- don't require **consensus**
- are based on **stale local states**
- yet, **collaboration** has to be successful

Plan of the talk

A motivating case study

Our formalisation

Our typing discipline

Tool support

Open issues

– A formal model –

Events

e

$src(e)$

Logs

$e_1 \cdot e_2 \cdot \dots$

Events

$\vdash e : t$

$src(e)$

Logs

$\vdash e_1 \cdot e_2 \dots : t_1 \cdot t_2 \dots$

Events

$\vdash e : t$

$src(e)$

Logs

$\vdash e_1 \cdot e_2 \dots : t_1 \cdot t_2 \dots$

order induced by $\ell = e_1 \dots e_n$ $e_i <_{\ell} e_j \iff i < j$

Ingredients (II): log shipping

Machine `Alice` **emits** logs upon **execution** of commands (we'll see how in a moment)


Ingredients (II): log shipping

Machine `Alice` **emits** logs upon **execution** of commands (we'll see how in a moment)
Such events are **appended** to the logs of machines in **two phases**:

Ingredients (II): log shipping

Machine **Alice** emits logs upon **execution** of commands (we'll see how in a moment)
Such events are **appended** to the logs of machines in **two phases**:


Phase I: emitted events are appended to the local log of the emitting machine



Ingredients (II): log shipping

Machine **Alice** emits logs upon **execution** of commands (we'll see how in a moment)
Such events are **appended** to the logs of machines in **two phases**:


Phase I: emitted events are appended to the local log of the emitting machine




Ingredients (II): log shipping

Machine **Alice** emits logs upon **execution** of commands (we'll see how in a moment)
Such events are **appended** to the logs of machines in **two phases**:

Phase I: emitted events are appended to the local log of the emitting machine




Phase II: newly emitted events are shipped to other machines




Ingredients (II): log shipping

Machine **Alice** emits logs upon **execution** of commands (we'll see how in a moment)
Such events are **appended** to the logs of machines in **two phases**:

Phase I: emitted events are appended to the local log of the emitting machine



Phase II: newly emitted events are shipped to other machines



Fix a set of commands ranged over by c

Let κ range over finite maps from commands to non-empty log types

Machines

Fix a set of commands ranged over by c

Let κ range over finite maps from commands to non-empty log types

A machine is a **regular term** of this co-inductive grammar

$$M ::=^{\text{co}} \kappa \cdot [t_1? M_1 \& \dots \& t_n? M_n]$$

for $i \in \{1 \dots, n\}$, the guard of the i -th branch is t_i

An infinite tree is regular when it has finitely-many subtrees. The subtrees of $M = \kappa \cdot [t_1? M_1 \& \dots \& t_n? M_n]$ are M plus the subtrees of each M_i .

An example

Passenger **P** launches an auction for a taxi **T**

InitialP = **Request** \mapsto **Requested** · [**Requested?** **AuctionP**]

AuctionP = **Select** \mapsto **Selected** · **PassengerId** · [
 Bid? **BidderId?** **AuctionP**
 &
 Selected? **PassengerId?** **RideP**
]

RideP = ...

An example

Passenger P launches an auction for a taxi T

$$\text{InitialP} = \text{Request} \mapsto \text{Requested} \cdot [\text{Requested? AuctionP}]$$
$$\begin{aligned} \text{AuctionP} = & \text{Select} \mapsto \text{Selected} \cdot \text{PassengerId} \cdot [\\ & \text{Bid? BidderId? AuctionP} \\ & \& \\ & \text{Selected? PassengerId? RideP} \\ &] \end{aligned}$$
$$\text{RideP} = \dots$$

Notation

- write $t_1?M_1 \& \dots \& t_n?M_n$ when κ is the empty function
- if $n = 0$, $\kappa \cdot 0$ abbreviates $\kappa \cdot [t_1?M_1 \& \dots \& t_n?M_n]$
- write $\&_{1 \leq i \leq n} l_i?M_i$ in place of $t_1?M_1 \& \dots \& t_n?M_n$

Treat κ as its graph and e.g. write $c / l \in \kappa$ for $\kappa(c) = l$ or write κ as $\{c_1 / l_1, \dots, c_h / l_h\}$ when $\kappa : c_i \mapsto l_i$ for $i \in \{1, \dots, h\}$

Machines as automata

A machine $M = \kappa \cdot [t_1?M_1 \& \dots \& t_n?M_n]$ is an FSA where:

- κ yields command-enabling transitions
- a branch $t_i?M_i$ yields a transition $M \xrightarrow{t_i?} M_i$ when an event of type t_i is consumed

Machines as automata

A machine $M = \kappa \cdot [t_1?M_1 \& \dots \& t_n?M_n]$ is an FSA where:

- κ yields command-enabling transitions
- a branch $t_i?M_i$ yields a transition $M \xrightarrow{t_i?} M_i$ when an event of type t_i is consumed

From machines to FSAs

- the states of the automaton are the subtrees of M
- the initial state is M and
 - there is a self-loop transition to M labelled $c / \mathbf{1}$ for each $c / \mathbf{1} \in \kappa$
 - there is a transition labelled $t_i?$ to state M_i for each $i \in \{1 \dots, n\}$
 - and likewise for M_i

Machines as automata

A machine $M = \kappa \cdot [t_1?M_1 \& \dots \& t_n?M_n]$ is an FSA where:

- κ yields command-enabling transitions
- a branch $t_i?M_i$ yields a transition $M \xrightarrow{t_i?} M_i$ when an event of type t_i is consumed

From machines to FSAs

- the states of the automaton are the subtrees of M
- the initial state is M and
 - there is a self-loop transition to M labelled $c / \mathbf{1}$ for each $c / \mathbf{1} \in \kappa$
 - there is a transition labelled $t_i?$ to state M_i for each $i \in \{1 \dots, n\}$
 - and likewise for M_i

This construction yields a finite-state automaton by the regularity of M

An example


Let's build the FSA of the machine `InitialP` on slide 18.



`InitialP` =

An example


Let's build the FSA of the machine `InitialP` on slide 18.



`InitialP` = `Request` \mapsto `Requested`.


An example

Let's build the FSA of the machine `InitialP` on slide 18.


$$\text{InitialP} = \text{Request} \mapsto \text{Requested} \cdot [\text{Requested? } \underline{\text{AuctionP}}]$$

An example

Let's build the FSA of the machine `InitialP` on slide 18.



`InitialP` = `Request` \mapsto `Requested` · [`Requested?` `AuctionP`]

`AuctionP` =

An example

Let's build the FSA of the machine `InitialP` on slide 18.




`InitialP` = `Request` \mapsto `Requested` · [`Requested?` `AuctionP`]

`AuctionP` = `Select` \mapsto `Selected` · `PassengerId` ·

An example

Let's build the FSA of the machine `InitialP` on slide 18.




`InitialP` = `Request` \mapsto `Requested` · [`Requested?` `AuctionP`]

`AuctionP` = `Select` \mapsto `Selected` · `PassengerId` · [
`Bid?` `BidderId?` `AuctionP`]

An example

Let's build the FSA of the machine `InitialP` on slide 18.




`InitialP` = `Request` \mapsto `Requested` · [`Requested?` `AuctionP`]

`AuctionP` = `Select` \mapsto `Selected` · `PassengerId` · [`Bid?` `BidderId?` `AuctionP`]

An example

Let's build the FSA of the machine **InitialP** on slide 18.



InitialP = **Request** \mapsto **Requested** · [**Requested?** **AuctionP**]

AuctionP = **Select** \mapsto **Selected** · **PassengerID** · [
 Bid? **BidderID?** **AuctionP**
 &
 Selected? **PassengerID?** **RideP**
]

RideP = ...

Machines' semantics

So, think of $M = \kappa \cdot [t_1? M_1 \& \dots \& t_n? M_n]$ as an FSA where transitions are

- either self-loops (determined by the κ part)
- or event consumptions (determined by the guards of the branches t_i)

Machines' semantics

So, think of $M = \kappa \cdot [t_1? M_1 \& \dots \& t_n? M_n]$ as an FSA where transitions are

- either self-loops (determined by the κ part)
- or event consumptions (determined by the guards of the branches t_i)

We restrict to **deterministic** machines and treat them as emitters/consumers of events with a semantics given in terms of state transition function :

$$\delta(M, \epsilon) = M$$

$$\delta(M, e \cdot \ell) = \begin{cases} \delta(M', \ell) & \text{if } \vdash e : t, M \xrightarrow{t?} M' \\ \delta(M, \ell) & \text{otherwise} \end{cases}$$

Machines' semantics

So, think of $M = \kappa \cdot [t_1? M_1 \& \dots \& t_n? M_n]$ as an FSA where transitions are

- either self-loops (determined by the κ part)
- or event consumptions (determined by the guards of the branches t_i)

We restrict to **deterministic** machines and treat them as emitters/consumers of events with a semantics given in terms of state transition function :

$$\begin{aligned} \delta(M, \epsilon) &= M \\ \delta(M, e \cdot \ell) &= \begin{cases} \delta(M', \ell) & \text{if } \vdash e : t, M \xrightarrow{t?} M' \\ \delta(M, \ell) & \text{otherwise} \end{cases} \end{aligned}$$

That is

M with local log ℓ is in the implicit state $\delta(M, \ell)$ reached after processing each event in ℓ

Machines' semantics

So, think of $M = \kappa \cdot [t_1? M_1 \& \dots \& t_n? M_n]$ as an FSA where transitions are

- either self-loops (determined by the κ part)
- or event consumptions (determined by the guards of the branches t_i)

We restrict to **deterministic** machines and treat them as emitters/consumers of events with a semantics given in terms of state transition function :

$$\delta(M, \epsilon) = M$$

$$\delta(M, e \cdot \ell) = \begin{cases} \delta(M', \ell) & \text{if } \vdash e : t, M \xrightarrow{t?} M' \\ \delta(M, \ell) & \text{otherwise} \end{cases}$$

$$\frac{\delta(M, \ell) \xrightarrow{c/1} \delta(M, \ell) \quad \ell' \text{ fresh} \quad \vdash \ell' : 1}{(M, \ell) \xrightarrow{c/1} (M, \ell \cdot \ell')}$$

That is

M with local log ℓ is in the implicit state $\delta(M, \ell)$ reached after processing each event in ℓ

Machines' semantics

So, think of $M = \kappa \cdot [t_1? M_1 \& \dots \& t_n? M_n]$ as an FSA where transitions are

- either self-loops (determined by the κ part)
- or event consumptions (determined by the guards of the branches t_i)

We restrict to **deterministic** machines and treat them as emitters/consumers of events with a semantics given in terms of state transition function :

$$\delta(M, \epsilon) = M$$

$$\delta(M, e \cdot \ell) = \begin{cases} \delta(M', \ell) & \text{if } \vdash e : t, M \xrightarrow{t?} M' \\ \delta(M, \ell) & \text{otherwise} \end{cases}$$

$$\frac{\delta(M, \ell) \xrightarrow{c/1} \delta(M, \ell) \quad \ell' \text{ fresh} \quad \vdash \ell' : 1}{(M, \ell) \xrightarrow{c/1} (M, \ell \cdot \ell')}$$

That is


M with local log ℓ is in the implicit state $\delta(M, \ell)$ reached after processing each event in ℓ

That is

after processing the events in ℓ , M reaches a state enabling $c/1$ then the command execution can emit ℓ' of type 1 and append it to the local log of M

An example

Take the machine `InitialP` (slide 20) with a local log $\ell = \text{ignoreMe} \cdot \text{ignoreMeToo}$ where $\nVdash \text{ignoreMe} : \text{Requested}$ and $\nVdash \text{ignoreMeToo} : \text{Requested}$




By definition of δ

- $\delta(\text{InitialP}, \ell) = \text{InitialP}$

An example

Take the machine `InitialP` (slide 20) with a local log $\ell = \text{ignoreMe} \cdot \text{ignoreMeToo}$ where $\not\vdash \text{ignoreMe} : \text{Requested}$ and $\not\vdash \text{ignoreMeToo} : \text{Requested}$




By definition of δ

- $\delta(\text{InitialP}, \ell) = \text{InitialP}$ hence
- $\delta(\text{InitialP}, \ell) \xrightarrow{\text{Request / Requested}} \delta(\text{InitialP}, \ell)$

An example

Take the machine `InitialP` (slide 20) with a local log $\ell = \text{ignoreMe} \cdot \text{ignoreMeToo}$ where $\not\vdash \text{ignoreMe} : \text{Requested}$ and $\not\vdash \text{ignoreMeToo} : \text{Requested}$




By definition of δ

- $\delta(\text{InitialP}, \ell) = \text{InitialP}$ hence
- $\delta(\text{InitialP}, \ell) \xrightarrow{\text{Request} / \text{Requested}} \delta(\text{InitialP}, \ell)$
- $(\text{InitialP}, \ell) \xrightarrow{\text{Request} / \text{Requested}} (\text{InitialP}, \ell \cdot \text{Requested})$ hence with $\vdash \text{Requested} : \text{Request}$ and $\text{src}(\text{Requested}) = \text{P}$ is possible

An example

Take the machine `InitialP` (slide 20) with a local log $\ell = \text{ignoreMe} \cdot \text{ignoreMeToo}$ where $\not\vdash \text{ignoreMe} : \text{Requested}$ and $\not\vdash \text{ignoreMeToo} : \text{Requested}$



By definition of δ

- $\delta(\text{InitialP}, \ell) = \text{InitialP}$ hence
- $\delta(\text{InitialP}, \ell) \xrightarrow{\text{Request} / \text{Requested}} \delta(\text{InitialP}, \ell)$
- $(\text{InitialP}, \ell) \xrightarrow{\text{Request} / \text{Requested}} (\text{InitialP}, \ell \cdot \text{Requested})$ hence with $\vdash \text{Requested} : \text{Request}$ and $\text{src}(\text{Requested}) = \text{P}$ is possible

Exercise

Calculate $\delta(\text{InitialP}, \ell \cdot \text{Requested})$

Some considerations

The commands are enabled only from the state reached **after processing all the events** in the local log of the machine

Some considerations

The commands are enabled only from the state reached **after processing all the events** in the local log of the machine

Deterministic machines may have **non-deterministic** behaviour!
Recall: commands are triggered by the environment

Some considerations

The commands are enabled only from the state reached **after processing all the events** in the local log of the machine

Deterministic machines may have **non-deterministic** behaviour!
Recall: commands are triggered by the environment

We have formalised the emission of events and their consumption
We now focus on the formalisation of **log shipping**

Swarms

A swarm (of size n) is a pair (\mathbf{S}, ℓ) where

- \mathbf{S} maps each index $1 \leq i \leq n$ to a pair (\mathbf{M}_i, ℓ_i)
- ℓ is the (global) log

Swarms

A swarm (of size n) is a pair (\mathbf{S}, ℓ) where

- \mathbf{S} maps each index $1 \leq i \leq n$ to a pair (\mathbf{M}_i, ℓ_i)
- ℓ is the (global) log

Notation

$$\mathbf{M}_1 \boxed{\ell_1} \mid \dots \mid \mathbf{M}_n \boxed{\ell_n} \mid \ell$$

Swarms

A swarm (of size n) is a pair (\mathbf{S}, ℓ) where

- \mathbf{S} maps each index $1 \leq i \leq n$ to a pair (\mathbf{M}_i, ℓ_i)
- ℓ is the (global) log

Notation

$\mathbf{M}_1 \boxed{\ell_1} \mid \dots \mid \mathbf{M}_n \boxed{\ell_n} \mid \ell$

Disclaimer

Seemingly, we've a contradiction: isn't the global log a centralisation point?

Well...no, it isn't: the global log is just a theoretical ploy!

- it abstracts away from low-level technical details for events' dispatching

Log shipping middlewares rely on timestamp mechanisms (Actyx uses Lamport's timestamps) and guarantee that events are in the same order in all the local logs

Swarms

A swarm (of size n) is a pair (\mathbf{S}, ℓ) where

- \mathbf{S} maps each index $1 \leq i \leq n$ to a pair (\mathbf{M}_i, ℓ_i)
- ℓ is the (global) log

Notation

$\mathbf{M}_1 \boxed{\ell_1} \mid \dots \mid \mathbf{M}_n \boxed{\ell_n} \mid \ell$

Disclaimer

Seemingly, we've a contradiction: isn't the global log a centralisation point?

Well...no, it isn't: the global log is just a theoretical ploy!

- it abstracts away from low-level technical details for events' dispatching
- it elegantly (IOHO) models asynchrony

Swarms

A swarm (of size n) is a pair (\mathbf{S}, ℓ) where

- \mathbf{S} maps each index $1 \leq i \leq n$ to a pair (\mathbf{M}_i, ℓ_i)
- ℓ is the (global) log

Notation

$\mathbf{M}_1 \boxed{\ell_1} \mid \dots \mid \mathbf{M}_n \boxed{\ell_n} \mid \ell$

Disclaimer

Seemingly, we've a contradiction: isn't the global log a centralisation point?

Well...no, it isn't: the global log is just a theoretical ploy!

- it abstracts away from low-level technical details for events' dispatching
- it elegantly (IOHO) models asynchrony
- it is not used in our algorithms and tools

Coherence

A swarm $M_1 \ell_1 \mid \dots \mid M_n \ell_n \mid \ell$ is coherent if $\ell = \bigcup_{1 \leq i \leq n} \ell_i$ and $\ell_i \sqsubseteq \ell$ for $1 \leq i \leq n$

Coherence

A swarm $M_1 \boxed{\ell_1} \mid \dots \mid M_n \boxed{\ell_n} \mid \ell$ is coherent if $\ell = \bigcup_{1 \leq i \leq n} \ell_i$ and $\ell_i \sqsubseteq \ell$ for $1 \leq i \leq n$

where $\ell_1 \sqsubseteq \ell_2$ is the sublog relation defined as

- $\ell_1 \subseteq \ell_2$ and $<_{\ell_1} \subseteq <_{\ell_2}$ and

- $e <_{\ell_2} e'$, $src(e) = src(e')$ and $e' \in \ell_1 \implies e \in \ell_1$

That is

all events of ℓ_1 appear in the same order in ℓ_2

That is

the per-source partitions of ℓ_1 are prefixes of the corresponding partitions of ℓ_2

Coherence

A swarm $M_1 \boxed{\ell_1} \mid \dots \mid M_n \boxed{\ell_n} \mid \ell$ is coherent if $\ell = \bigcup_{1 \leq i \leq n} \ell_i$ and $\ell_i \sqsubseteq \ell$ for $1 \leq i \leq n$

where $\ell_1 \sqsubseteq \ell_2$ is the sublog relation defined as

- $\ell_1 \subseteq \ell_2$ and $<_{\ell_1} \subseteq <_{\ell_2}$ and

- $e <_{\ell_2} e'$, $src(e) = src(e')$ and $e' \in \ell_1 \implies e \in \ell_1$

That is

all events of ℓ_1 appear in the same order in ℓ_2

That is

the per-source partitions of ℓ_1 are prefixes of the corresponding partitions of ℓ_2

Hereafter, we assume coherence

Merging logs

Exercise

Recall slide 16 and consider a swarm

$$\dots \mid \text{Alice} \mid \begin{array}{|c|c|c|c|c|c|} \hline e_1 & e_2 & e_3 & a & b & c \\ \hline \end{array} \mid \dots \mid \ell \quad (1)$$

If $\ell = e_1 \cdot e_2 \cdot e_3 \cdot e$, under which condition is (1) coherent?

Merging logs

Exercise

Recall slide 16 and consider a swarm

$$\dots \mid \text{Alice} \mid \begin{array}{|c|c|c|c|c|c|} \hline e_1 & e_2 & e_3 & a & b & c \\ \hline \end{array} \mid \dots \mid \ell \quad (1)$$

If $\ell = e_1 \cdot e_2 \cdot e_3 \cdot e$, under which condition is (1) coherent?

The propagation of newly generated events happens by merging logs:

Log merging: $l_1 \bowtie l_2 = \{l \mid l \subseteq l_1 \cup l_2 \text{ and } l_1 \sqsubseteq l \text{ and } l_2 \sqsubseteq l\}$

Semantics of swarms

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

$$\frac{\mathbf{S}(i) = \mathbf{M}[\ell_i] \quad \mathbf{M}[\ell_i] \xrightarrow{c/1} \mathbf{M}[\ell'_i] \quad \text{src}(\ell'_i \setminus \ell_i) = \{i\} \quad \ell' \in \ell \bowtie \ell'_i}{(\mathbf{S}, \ell) \xrightarrow{c/1} (\mathbf{S}[i \mapsto \mathbf{M}[\ell'_i]], \ell')} \text{[Local]}$$

Semantics of swarms

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

$$\frac{\mathbf{S}(i) = \mathbf{M}_{\boxed{\ell_i}} \quad \mathbf{M}_{\boxed{\ell_i}} \xrightarrow{c/1} \mathbf{M}_{\boxed{\ell'_i}} \quad \text{src}(\ell'_i \setminus \ell_i) = \{i\} \quad \ell' \in \ell \bowtie \ell'_i}{(\mathbf{S}, \ell) \xrightarrow{c/1} (\mathbf{S}[i \mapsto \mathbf{M}_{\boxed{\ell'_i}}], \ell')} \text{[Local]}$$

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

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

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

then, by [Local]

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

Semantics at work (I)

If

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

then, by [Local]

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

for all

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

Semantics at work (I)

If

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

then, by [Local]

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

for all

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

Exercise

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

Semantics at work (II)

Take from slide 28

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

and let's propagate some events

Semantics at work (II)

Take from slide 28

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

and let's propagate some events

Exercise

Can we propagate just event e ?

Semantics at work (II)

Take from slide 28

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

and let's propagate some events

Exercise

Can we propagate just event e ?

By rule [Prop] we can propagate a non-deterministically chosen sublog of $b \cdot d \cdot e$

Semantics at work (II)

Take from slide 28

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

and let's propagate some events

Exercise

Can we propagate just event e ?

By rule [Prop] we can propagate a non-deterministically chosen sublog of $b \cdot d \cdot e$

Let's propagate $d \cdot e$

$$A[a] | B[b \cdot d \cdot e] | C[c] | \ell \begin{array}{l} \xrightarrow{\tau} A[b \cdot a \cdot d \cdot e] | B[b \cdot d \cdot e] | C[c] | \ell \\ \xrightarrow{\tau} A[a] | B[b \cdot d \cdot e] | C[b \cdot d \cdot e \cdot c] | \ell \end{array}$$

Semantics at work (II)

Take from slide 28

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

and let's propagate some events

Exercise

Can we propagate just event e ?

By rule [Prop] we can propagate a non-deterministically chosen sublog of $b \cdot d \cdot e$

Let's propagate $d \cdot e$

$$A[a] | B[b \cdot d \cdot e] | C[c] | \ell \begin{cases} \xrightarrow{\tau} A[b \cdot a \cdot d \cdot e] | B[b \cdot d \cdot e] | C[c] | \ell \\ \xrightarrow{\tau} A[a] | B[b \cdot d \cdot e] | C[b \cdot d \cdot e \cdot c] | \ell \end{cases}$$

Excercise

In both cases b must be shipped too. Why?

And why is event a not shipped to C together with the events from B ?

Plan of the talk

A motivating case study

Our formalisation

Our typing discipline


Tool support

Open issues

– Behavioural types for swarms –


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)

Quoting W3C:

*"[...] a **contract** [...] of the common **ordering conditions and constraints** under which **messages** are exchanged [...] from a **global viewpoint** [...]
Each **party** can then use the global definition to **build and test solutions** [...]
global specification is in turn **realised by combination of** the resulting **local systems**"*

Choreographies

Quoting W3C:

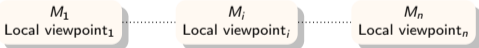
“[...] a *contract* [...] of the common *ordering conditions and constraints* under which *messages* are exchanged [...] from a *global viewpoint* [...]

Each party can then use the global definition to *build and test solutions* [...] global specification is in turn *realised by combination of the resulting local systems*”

Synchrony

Choreography G
global viewpoint

Asynchrony



Choreographies

Quoting W3C:

“[...] a *contract* [...] of the common *ordering conditions and constraints* under which *messages* are exchanged [...] from a *global viewpoint* [...]

Each party can then use the global definition to *build and test solutions* [...]

global specification is in turn *realised by combination of the resulting local systems*”

Synchrony

Choreography G
global viewpoint

Asynchrony

M_1
Local viewpoint₁

M_i
Local viewpoint_i

M_n
Local viewpoint_n

spec, no code

Choreographies

Quoting W3C:

“[...] a *contract* [...] of the common *ordering conditions and constraints* under which *messages* are exchanged [...] from a *global viewpoint* [...]

Each party can then use the global definition to *build and test solutions* [...] global specification is in turn *realised by combination of the resulting local systems*”

Synchrony

Choreography G
global viewpoint

Well-formedness

Asynchrony

M_1
Local viewpoint₁

M_i
Local viewpoint_i

M_n
Local viewpoint_n

spec, no code


Choreographies

Quoting W3C:

“[...] a *contract* [...] of the common *ordering conditions and constraints* under which *messages* are exchanged [...] from a *global viewpoint* [...]

Each party can then use the global definition to *build and test solutions* [...]

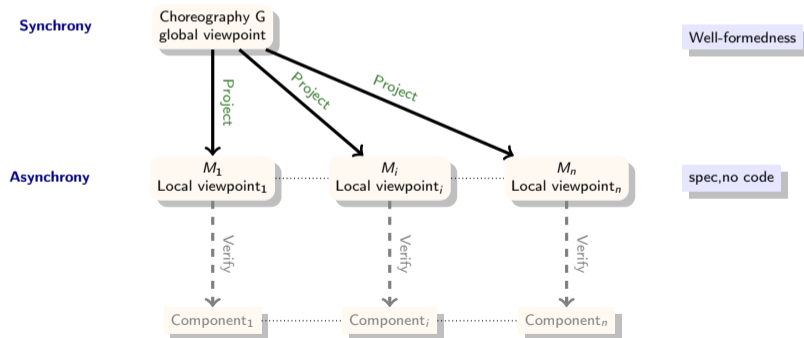
global specification is in turn *realised by combination of the resulting local systems*”



Choreographies

Quoting W3C:

“[...] a *contract* [...] of the common *ordering conditions and constraints* under which *messages* are exchanged [...] from a *global viewpoint* [...] Each *party* can then use the global definition to *build and test solutions* [...] global specification is in turn *realised by combination of the resulting local systems*”



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} on slide 32)

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 on slide 32:

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

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

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

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

An example

A swarm protocol for the taxi scenario on slide 32:

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

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

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

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


*Note 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 $i \in \{1 \dots, n\}$
- 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



Removing log types yields exactly the FSA of the swarm protocol on slide 32


An example



There is a race in state 3!

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

An example



Removing log types yields exactly the FSA of the swarm protocol on slide 32

There is a race in state 3!

- 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 selectors, which is also true of states 2 and 5)*

Semantics of swarm protocols

One rule only!

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

Semantics of swarm protocols

One rule only!

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

where

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

*Logs to be consumed "atomically",
hence $\delta(G, \ell)$ 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

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

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

Semantics of swarm protocols

One rule only!

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

where

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

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

We restrict ourselves to deterministic swarm protocols that is, on different transitions from a same state

- log types start differently
- pairs (command,role) differ

log determinism
command determinism

From swarm protocols to machines

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

From swarm protocols to machines

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

Each machine plays one role

From swarm protocols to machines

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 swarm protocols to machines

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

From swarm protocols to machines

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

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

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

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

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

where

Another attempt



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

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

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

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

where


$$\kappa = \{c_i / \mathbf{1}_i \mid R_i = R \text{ and } i \in I\}$$

$$J = \{i \in I \mid \text{filter}(\mathbf{1}_i, \sigma(R)) \neq \epsilon\}$$

$$\text{filter}(\mathbf{1}, E) = \begin{cases} \epsilon, & \text{if } \mathbf{t} = \epsilon \\ \mathbf{t} \cdot \text{filter}(\mathbf{1}', E) & \text{if } \mathbf{t} \in E \text{ and } \mathbf{1} = \mathbf{t} \cdot \mathbf{1}' \\ \text{filter}(\mathbf{1}, E) & \text{otherwise} \end{cases}$$

An example

A reasonable subscription for **P** is the total one since the passenger should be aware of all events: $\sigma(\mathbf{P})$ contains all event types



An example

Exercise

The taxi driver does not need to bother with the receipt: the subscription for $\sigma(\mathbb{T})$ consists of all messages but **Receipt**; give the projection of the taxi protocol on such subscription for \mathbb{T} .

An example

Exercise

The taxi driver does not need to bother with the receipt: the subscription for $\sigma(\mathbb{T})$ consists of all messages but **Receipt**; give the projection of the taxi protocol on such subscription for \mathbb{T} .

If we want the office to know only the details about the ride we set $\sigma(\mathbb{O}) = \{\mathbf{Started}, \mathbf{Finished}, \mathbf{Receipt}\}$



An example

Exercise

The taxi driver does not need to bother with the receipt: the subscription for $\sigma(\mathbb{T})$ consists of all messages but **Receipt**; give the projection of the taxi protocol on such subscription for \mathbb{T} .

If we want the office to know only the details about the ride we set $\sigma(\mathbb{O}) = \{\mathbf{Started}, \mathbf{Finished}, \mathbf{Receipt}\}$



Exercise (hard)

Is this a good idea?

Well-formedness: sufficient conditions for well-behaviour

Transitory deviations are tolerated provided that consistency is eventually recovered

Well-formedness: sufficient conditions for well-behaviour

Transitory deviations are tolerated provided that consistency is eventually recovered

Example

T may bid after **P** has made their selection if the selection event **T** has not yet been received.

This inconsistency is temporary: when the selection event reaches **T** this inconsistency is recognised and resolved

Well-formedness: sufficient conditions for well-behaviour

Transitory deviations are tolerated provided that consistency is eventually recovered

Example

T may bid after **P** has made their selection if the selection event **T** has not yet been received.

This inconsistency is temporary: when the selection event reaches **T** this inconsistency is recognised and resolved

Convention

Let's write $\mathbf{R} \in_{\sigma} \mathbf{G} = \sum_{i \in I} c_i @ \mathbf{R}_i \langle \mathbf{1}_i \rangle \cdot \mathbf{G}_i$ when there is $i \in I$ such that

$$\mathbf{R} = \mathbf{R}_i \quad \text{or} \quad \sigma(\mathbf{R}) \cap \mathbf{1}_i \neq \emptyset \quad \text{or} \quad \mathbf{R} \in_{\sigma} \mathbf{G}_i$$

and set $\text{roles}(\mathbf{G}, \sigma) = \{\mathbf{R} \mid \mathbf{R} \in_{\sigma} \mathbf{G}\}$ and

Well-formedness

Trading consistency for availability has implications:

Well-formedness = 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 $\sigma(R_i) \cap \mathbf{1}_i \neq \emptyset$

If R should have c enabled after c' then $\sigma(R)$ contains some event type emitted by c'

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

Well-formedness = Causality + Determinacy

Trading consistency for availability has implications:

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

\implies different roles may take inconsistent decisions

Causality & Determinacy

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

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 $\sigma(R) \cap \mathbf{1}_i \supseteq \bigcup_{R' \in \sigma G_i} \sigma(R') \cap \mathbf{1}_i$

Determinacy $R \in \sigma G_i \implies \mathbf{1}_i[0] \in \sigma(R)$

Well-formedness = 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 $G = \sum_{i \in I} c_i @ R_i \langle \mathbf{1}_i \rangle . G_i$

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 $\sigma(R) \cap \mathbf{1}_i \supseteq \bigcup_{R' \in \sigma G_i} \sigma(R') \cap \mathbf{1}_i$

Determinacy $R \in \sigma G_i \implies \mathbf{1}_i[0] \in \sigma(R)$

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

Some considerations

Further consequences:

- **Unspecified receptions** are just ignored according to the δ transition function of machines
- It is fine to violate **session fidelity**, provided that consistency is eventually attained

Some considerations

Further consequences:

- **Unspecified receptions** are just ignored according to the δ transition function of machines
- It is fine to violate **session fidelity**, provided that consistency is eventually attained

Care is therefore necessary

- for the definition of **correctness**
- and for the **correct realisation** of swarm protocols

Of course we appeal to projections

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 $S = \text{P} \parallel \text{T} \parallel \text{O} \parallel \text{T}$ implementing



(i.e., the swarm protocol G on slide 37). Check that S generates the log

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


On correctness



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

Exercise

Take the swarm $S = \text{P} \parallel \text{T} \parallel \text{O} \parallel \text{T}$ implementing



(i.e., the swarm protocol G on slide 37). 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

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle 1_i \rangle . G_i) = \bigcup_{i \in I} \{R_i\}$

l has effective type 1 wrt G and σ if $G, \epsilon \vdash_{\sigma} l \triangleright 1$ is provable; where

$$G, \epsilon \vdash_{\sigma} e \cdot l \triangleright t \cdot 1$$

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle 1_i \rangle . G_i) = \bigcup_{i \in I} \{R_i\}$

l has effective type 1 wrt G and σ if $G, \epsilon \vdash_{\sigma} l \triangleright 1$ is provable; where

$\vdash e : t$

$G, \epsilon \vdash_{\sigma} e \cdot l \triangleright t \cdot 1$

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle l_i \rangle . G_i) = \bigcup_{i \in I} \{R_i\}$

l has effective type 1 wrt G and σ if $G, \epsilon \vdash_{\sigma} l \triangleright 1$ is provable; where

$$\frac{\vdash e : t \in \sigma(\text{roles}(G, \sigma)) \quad G \xrightarrow{c/t \cdot 1'} G'}{G, \epsilon \vdash_{\sigma} e \cdot l \triangleright t \cdot 1}$$

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle \mathbf{1}_i \rangle . G_i) = \bigcup_{i \in I} \{R_i\}$

ℓ has effective type $\mathbf{1}$ wrt G and σ if $G, \epsilon \vdash_{\sigma} \ell \triangleright \mathbf{1}$ is provable; where

$$\frac{\vdash e : t \in \sigma(\text{roles}(G, \sigma)) \quad G \xrightarrow{c/t \cdot \mathbf{1}'} G' \quad G', \text{filter}(\mathbf{1}', \sigma(\text{active}(G'))) \vdash_{\sigma} \ell \triangleright \mathbf{1}}{G, \epsilon \vdash_{\sigma} e \cdot \ell \triangleright t \cdot \mathbf{1}}$$

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle l_i \rangle \cdot G_i) = \bigcup_{i \in I} \{R_i\}$

l has effective type 1 wrt G and σ if $G, \epsilon \vdash_{\sigma} l \triangleright 1$ is provable; where

$$\frac{\vdash e : t \in \sigma(\text{roles}(G, \sigma)) \quad G \xrightarrow{c/t \cdot 1'} G' \quad G', \text{filter}(1', \sigma(\text{active}(G'))) \vdash_{\sigma} l \triangleright 1}{G, \epsilon \vdash_{\sigma} e \cdot l \triangleright t \cdot 1}$$
$$\frac{\vdash e : t \quad G, 1 \vdash_{\sigma} l \triangleright 1'}{G, t \cdot 1 \vdash_{\sigma} e \cdot l \triangleright t \cdot 1'}$$

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle l_i \rangle . G_i) = \bigcup_{i \in I} \{R_i\}$

l has effective type 1 wrt G and σ if $G, \epsilon \vdash_{\sigma} l \triangleright 1$ is provable; where

$$\frac{\vdash e : t \in \sigma(\text{roles}(G, \sigma)) \quad G \xrightarrow{c/t \cdot 1'} G' \quad G', \text{filter}(1', \sigma(\text{active}(G'))) \vdash_{\sigma} l \triangleright 1}{G, \epsilon \vdash_{\sigma} e \cdot l \triangleright t \cdot 1}$$
$$\frac{\vdash e : t \quad G, 1 \vdash_{\sigma} l \triangleright 1'}{G, t \cdot 1 \vdash_{\sigma} e \cdot l \triangleright t \cdot 1'}$$
$$\frac{}{G, 1 \vdash_{\sigma} \epsilon \triangleright \epsilon}$$

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle l_i \rangle . G_i) = \bigcup_{i \in I} \{R_i\}$

l has effective type 1 wrt G and σ if $G, \epsilon \vdash_{\sigma} l \triangleright 1$ is provable; where

$$\frac{\vdash e : t \in \sigma(\text{roles}(G, \sigma)) \quad G \xrightarrow{c/t \cdot 1'} G' \quad G', \text{filter}(1', \sigma(\text{active}(G'))) \vdash_{\sigma} l \triangleright 1}{G, \epsilon \vdash_{\sigma} e \cdot l \triangleright t \cdot 1}$$
$$\frac{\vdash e : t \quad G, 1 \vdash_{\sigma} l \triangleright 1'}{G, t \cdot 1 \vdash_{\sigma} e \cdot l \triangleright t \cdot 1'}$$
$$\frac{G, 1 \vdash_{\sigma} l \triangleright 1' \quad \text{none of the other rules applies}}{G, 1 \vdash_{\sigma} e \cdot l \triangleright 1'}$$
$$\frac{}{G, 1 \vdash_{\sigma} \epsilon \triangleright \epsilon}$$

Effective types

Let $\text{active}(\sum_{i \in I} c_i @ R_i \langle 1_i \rangle . G_i) = \bigcup_{i \in I} \{R_i\}$

l has effective type 1 wrt G and σ if $G, \epsilon \vdash_{\sigma} l \triangleright 1$ is provable; where

$$\frac{\frac{\frac{\vdash e : t \in \sigma(\text{roles}(G, \sigma)) \quad G \xrightarrow{c/t \cdot 1'} G' \quad G', \text{filter}(1', \sigma(\text{active}(G'))) \vdash_{\sigma} l \triangleright 1}{G, \epsilon \vdash_{\sigma} e \cdot l \triangleright t \cdot 1}}{\frac{\vdash e : t \quad G, 1 \vdash_{\sigma} l \triangleright 1'}{G, t \cdot 1 \vdash_{\sigma} e \cdot l \triangleright t \cdot 1'}} \quad \frac{}{G, 1 \vdash_{\sigma} \epsilon \triangleright \epsilon}}{\frac{G, 1 \vdash_{\sigma} l \triangleright 1' \quad \text{none of the other rules applies}}{G, 1 \vdash_{\sigma} e \cdot l \triangleright 1'}}$$

Exercise

For the swarm protocol G on slide 37, find a condition on σ so that

$$G, \epsilon \vdash_{\sigma} l_{\text{auc}} \triangleright \text{Requested} . \text{Bid} . \text{BidderID} . \text{Selected} . \text{PassengerID}$$

Implementations

Write $l \equiv_{G,\sigma} l'$ when l and l' have the same effective type wrt G and σ .

A swarm (S, ϵ) is eventually faithful to G and σ if $(S, \epsilon) \Longrightarrow (S, l)$ then there is $(G, \epsilon) \Longrightarrow (G, l')$ with $l \equiv_{G,\sigma} l'$

Implementations

Write $l \equiv_{G,\sigma} l'$ when l and l' have the same effective type wrt G and σ .

A swarm (S, ϵ) is eventually faithful to G and σ if $(S, \epsilon) \Longrightarrow (S, l)$ then there is $(G, \epsilon) \Longrightarrow (G, l')$ with $l \equiv_{G,\sigma} l'$

A (σ, G) -realisation is a swarm (S, ϵ) of size n such that, for each $1 \leq i \leq n$, there exists a role $R \in \text{roles}(G, \sigma)$ such that $S(i) = G \downarrow_R^\sigma$

Implementations & projections

Write $\ell \equiv_{\mathbf{G},\sigma} \ell'$ when ℓ and ℓ' have the same effective type wrt \mathbf{G} and σ .


A swarm (\mathbf{S}, ϵ) is eventually faithful to \mathbf{G} and σ if $(\mathbf{S}, \epsilon) \Longrightarrow (\mathbf{S}, \ell)$ then there is $(\mathbf{G}, \epsilon) \Longrightarrow (\mathbf{G}, \ell')$ with $\ell \equiv_{\mathbf{G},\sigma} \ell'$

A (σ, \mathbf{G}) -realisation is a swarm (\mathbf{S}, ϵ) of size n such that, for each $1 \leq i \leq n$, there exists a role $\mathbf{R} \in \text{roles}(\mathbf{G}, \sigma)$ such that $\mathbf{S}(i) = \mathbf{G} \downarrow_{\mathbf{R}}^{\sigma}$

Lemma (Projections of well-formed protocols are eventually faithful)


If \mathbf{G} is a σ -WF protocol and $(\delta(\mathbf{G} \downarrow_{\mathbf{R}}^{\sigma}, \ell)) \downarrow_{\mathbf{c}/1}$ then there exists $\ell' \equiv_{\mathbf{G},\sigma} \ell$ such that $(\mathbf{G}, \epsilon) \Longrightarrow (\mathbf{G}, \ell')$ and $\delta(\mathbf{G}, \ell') \xrightarrow{\mathbf{c}/1} \mathbf{G}'$

On correct realisations



*A set of runs is consistent
when its elements are
pair-wise consistent*

On correct realisations




A set of runs is consistent when its elements are pair-wise consistent

Notation

For $(G, \epsilon) \xrightarrow{c_1 / l_1} (G, l_1) \xrightarrow{c_2 / l_2} \dots \xrightarrow{c_n / l_n} (G, \overbrace{l_1 \cdot l_2 \cdot \dots \cdot l_n}^{=l})$
 let $l^{(j)} = l_j \cdot \dots \cdot l_1$

On correct realisations



A set of runs is consistent when its elements are pair-wise consistent

Notation

For $(G, \epsilon) \xrightarrow{c_1 / l_1} (G, l_1) \xrightarrow{c_2 / l_2} \dots \xrightarrow{c_n / l_n} (G, \overbrace{l_1 \cdot l_2 \cdot \dots \cdot l_n}^{=l})$
 let $l^{(j)} = l_j \cdot \dots \cdot l_1$

Admissibility

A log l is admissible for a σ -WF protocol G if there are consistent runs $\{(G, \epsilon) \implies (G, l_i)\}_{1 \leq i \leq k}$ and a log $l' \in (\boxtimes_{1 \leq i \leq k} l_i)$ such that $l = \bigcup_{1 \leq i \leq k} l_i$ and

$$l' \equiv_{G, \sigma} l \quad \text{and} \quad l_i^{(j)} \sqsubseteq l \quad \text{for all } 1 \leq i \leq k$$

Hereafter, G denotes a σ -WF protocol

Results

Lemma (Well-formedness generates any admissible log)

If ℓ is admissible for G then there is a log ℓ' such that $(G, \epsilon) \Longrightarrow (G, \ell')$ and $\ell \equiv_{G, \sigma} \ell'$

Lemma (Admissibility is preserved)

Let ℓ_1 and $\ell_2 \subseteq \ell_1$ be admissible logs for G . If $(G, \ell_2) \xrightarrow{c/1} (G, \ell_2 \cdot \ell_3)$ and $\ell \in \ell_1 \bowtie (\ell_2 \cdot \ell_3)$ then ℓ is admissible for G

Theorem (Well-formed protocols generate only admissible logs)

If $(S, \epsilon) \Longrightarrow (S', \ell)$ for (S, ϵ) realisation of G then ℓ is admissible for G

Corollary

Every realisation of G is eventually faithful wrt G and σ

On complete realisations

Complete realisations

A (σ, \mathbf{G}) -realisation (\mathbf{S}, ϵ) of size n is complete if for all $\mathbf{R} \in \text{roles}(\mathbf{G}, \sigma)$ there exists $1 \leq i \leq n$ such that $\mathbf{S}(i) = \mathbf{G} \downarrow_{\mathbf{R}}^{\sigma}$

Lemma (Projections reflect swarm protocols)

If $(\mathbf{G}, \epsilon) \implies (\mathbf{G}, \ell)$ then $\delta(\mathbf{G} \downarrow_{\mathbf{R}}^{\sigma}, \ell) = \delta(\mathbf{G}, \ell) \downarrow_{\mathbf{R}}^{\sigma}$ for all $\mathbf{R} \in \text{roles}(\mathbf{G}, \sigma)$

Theorem (Complete realisations reflect the protocol)

Let (\mathbf{S}, ϵ) be a complete realisation of \mathbf{G} . If $(\mathbf{G}, \epsilon) \implies (\mathbf{G}, \ell)$ then there is a swarm \mathbf{S}' such that $(\mathbf{S}, \epsilon) \implies (\mathbf{S}', \ell)$

Plan of the talk

A motivating case study

Our formalisation

Our typing discipline

Tool support

Open issues

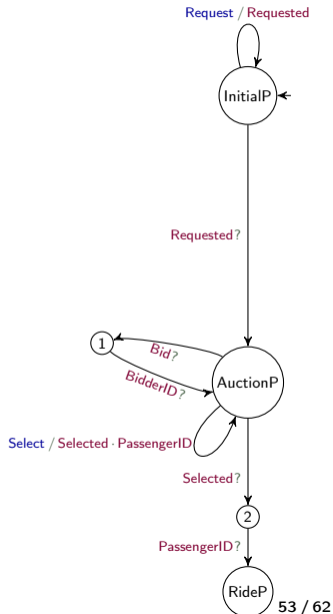
– Tooling –

```

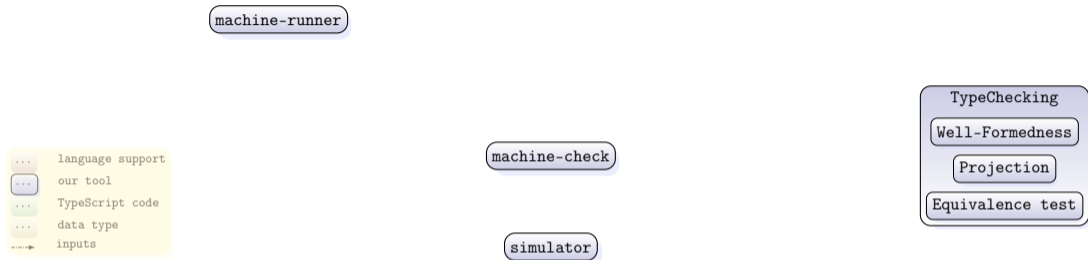
// analogous for other events; "type" property matches type name (checked by tool)
type Requested = { type: 'Requested'; pickup: string; dest: string }
type Events = Requested | Bid | BidderID | Selected | ...

/** Initial state for role P */
@proto('taxiRide') // decorator injects inferred protocol into runtime
export class InitialP extends State<Events> {
  constructor(public id: string) { super() }
  execRequest(pickup: string, dest: string) {
    return this.events({ type: 'Requested', pickup, dest })
  }
  onRequest(ev: Requested) {
    return new AuctionP(this.id, ev.pickup, ev.dest, [])
  }
}
@proto('taxiRide')
export class AuctionP extends State<Events> {
  constructor(public id: string, public pickup: string, public dest: string,
    public bids: BidData[]) { super() }
  onBid(ev1: Bid, ev2: BidderID) {
    const [ price, time ] = ev1
    this.bids.push({ price, time, bidderID: ev2.id })
    return this
  }
  execSelect(taxiId: string) {
    return this.events({ type: 'Selected', taxiID },
      { type: 'PassengerID', id: this.id })
  }
  onSelect(ev: Selected, id: PassengerID) {
    return new RideP(this.id, ev.taxiID)
  }
}
@proto('taxiRide')
export class RideP extends State<Events> { ... }

```




Architecture



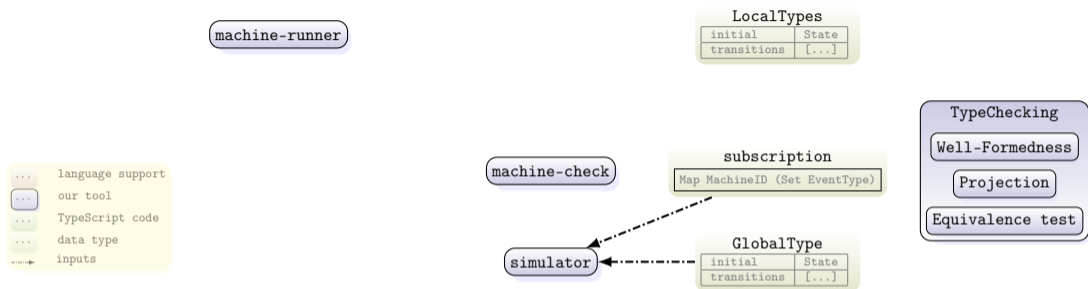
- `TypeChecking` implements the functionalities of our typing discipline
- `simulator` simulates the semantics of swarm realisations
- `machine-check` and `machine-runner` integrate our framework in the Actyx platform

Architecture



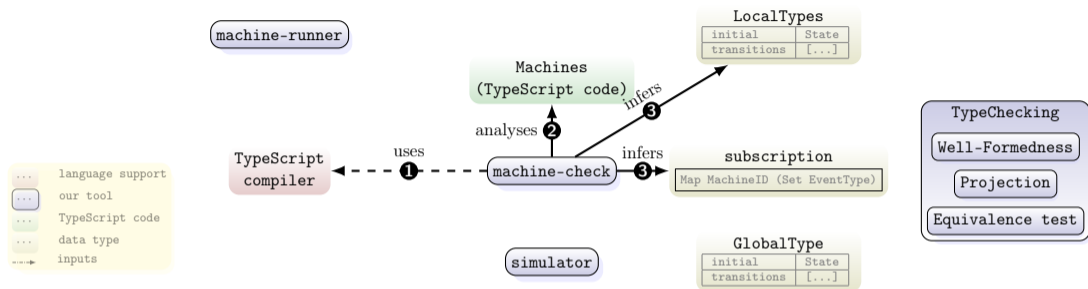
- TypeChecking implements the functionalities of our typing discipline
- simulator simulates the semantics of swarm realisations
- machine-check and machine-runner integrate our framework in the Actyx platform

Architecture



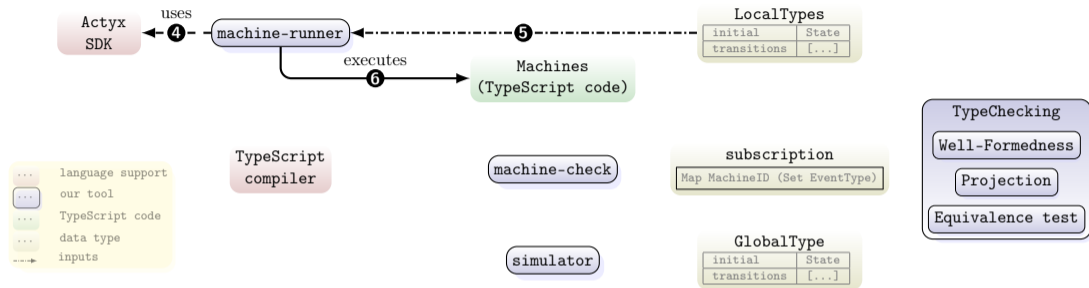
- TypeChecking implements the functionalities of our typing discipline
- simulator simulates the semantics of swarm realisations
- machine-check and machine-runner integrate our framework in the Actyx platform

Architecture



- TypeChecking implements the functionalities of our typing discipline
- simulator simulates the semantics of swarm realisations
- machine-check and machine-runner integrate our framework in the Actyx platform

Architecture



- TypeChecking implements the functionalities of our typing discipline
- simulator simulates the semantics of swarm realisations
- machine-check and machine-runner integrate our framework in the Actyx platform

If you want to play with our prototype?

Have a look at

- our ECOOP artifact paper (not online yet; extended version at <https://arxiv.org/abs/2305.04848>)
- code at <https://doi.org/10.5281/zenodo.7737188>
- An ISSTA tool paper from Actyx (<https://arxiv.org/abs/2306.09068>)

Plan of the talk

A motivating case study

Our formalisation

Our typing discipline

Tool support

Open issues

– Epilogue –

To be continued....

There are a number of future directions to explore:

To be continued....

There are a number of future directions to explore:

Identify weaker conditions for well-formedness

To be continued....

There are a number of future directions to explore:

Identify weaker conditions for well-formedness

“Efficiency”

To be continued....

There are a number of future directions to explore:

- Identify weaker conditions for well-formedness

- “Efficiency”

- Subscriptions are hard to determine

To be continued....

There are a number of future directions to explore:

- Identify weaker conditions for well-formedness

- “Efficiency”

- Subscriptions are hard to determine

- Relax some of our assumptions

To be continued....

There are a number of future directions to explore:

- Identify weaker conditions for well-formedness

- “Efficiency”

- Subscriptions are hard to determine

- Relax some of our assumptions

 - Compensations

 - Unreliable propagation

To be continued....

There are a number of future directions to explore:

- Identify weaker conditions for well-formedness

- “Efficiency”

- Subscriptions are hard to determine

- Relax some of our assumptions

 - Compensations

 - Unreliable propagation

 - Adversarial contexts

To be continued....

There are a number of future directions to explore:

Identify weaker conditions for well-formedness

“Efficiency”

Subscriptions are hard to determine

Relax some of our assumptions

Compensations

Unreliable propagation

Adversarial contexts

.....

An interesting paradigm grounded on principles for local-first software

Summary

An interesting paradigm grounded on principles for local-first software

We defined an operational semantics that captures the platform of Actyx AG

Summary

An interesting paradigm grounded on principles for local-first software

We defined an operational semantics that captures the platform of Actyx AG

We introduced behavioural types to specify and verify eventual consistency

Summary

An interesting paradigm grounded on principles for local-first software

We defined an operational semantics that captures the platform of Actyx AG

We introduced behavioural types to specify and verify eventual consistency

The key idea is to trade consistency for availability: temporary inconsistency are tolerated provided that they can be resolved at some point

Thank you!

– Solutions –

Solutions to exercises

- Slide 22: $\delta(\text{InitialP}, \ell \cdot \text{Requested}) = \text{AuctionP}$
- Slide 26: $\text{src}(e) \neq \text{Alice}$
- Slide 28: $(a \cdot b \cdot c) \bowtie (b \cdot d \cdot e) = \{a \cdot b \cdot c \cdot d \cdot e, a \cdot b \cdot d \cdot c \cdot e, a \cdot b \cdot d \cdot e \cdot c\}$
- Slide 29: Because [prop] won't apply since e is not a sublog of the local log of B
- Slide 41: The solution of the first exercise is in our ECOOP paper. For the second exercise, the idea is not bad because with such subscription the protocol is not well-formed (work out why)
- Slide 45: Apply the operational semantics of swarms
- Slide 46: $\sigma(P) \ni \text{Requested}, \text{BidderID}, \text{Selected}, \text{PassengerID}$