

# Automata for choreographies

Emilio Tuosto



Department of Computer Science

DI/NOVA LINCS

July 27, 2022

## Take-home message

- Choreographies for communicating systems
- Two automata-based models
  - point-to-point communications
  - event-notification coordination
- An emerging pattern
  - fix a communication model
  - find suitable global and local specs
  - define well-formedness
  - get correct realisations by projection

– Act I –

Formal Choreographies, informally

(joint work with Roberto Guanciale)

# Top-down model-driven development

Choreography = Global spec + Local spec

## 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**"*

# Top-down model-driven development

Choreography = Global spec + Local spec

## 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*”




# Top-down model-driven development

Choreography = Global spec + Local spec

## 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*”

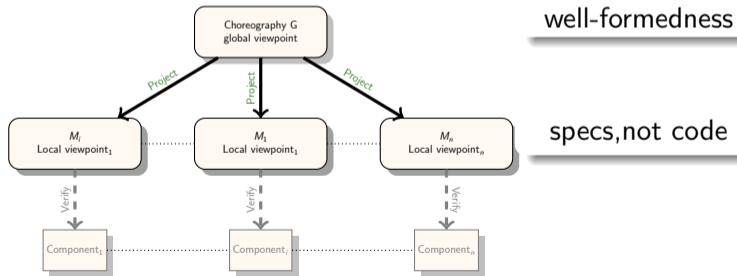


# Top-down model-driven development

Choreography = Global spec + Local spec

## 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*”



# A model of global specs

$G, G'$	$::=$	(o)	empty
		$A \rightarrow B : m$	interaction
		$G; G'$	sequential
		$G \mid G'$	parallel
		$G + G'$	branch
		$G^*$	iteration







# A model of global specs



$G, G' ::=$

- (o)
- $A \rightarrow B : m$
- $G; G'$
- $G \mid G'$
- $G + G'$
- $G^*$

empty  
interaction  
sequential  
parallel  
branch  
iteration




## Some examples




# A model of local specs

Communicating Finite-State Machines [Brand&Zafiropulo 1983]



A






B

Global specs can be projected (i.e., compiled) on CFSMs

# A model of local specs


Communicating Finite-State Machines [Brand&Zafiropulo 1983]



Global specs can be projected (i.e., compiled) on CFSMs

# A model of local specs


Communicating Finite-State Machines [Brand&Zafiropulo 1983]



Global specs can be projected (i.e., compiled) on CFSMs

# A model of local specs


Communicating Finite-State Machines [Brand&Zafiropulo 1983]



Global specs can be projected (i.e., compiled) on CFSMs

# A model of local specs

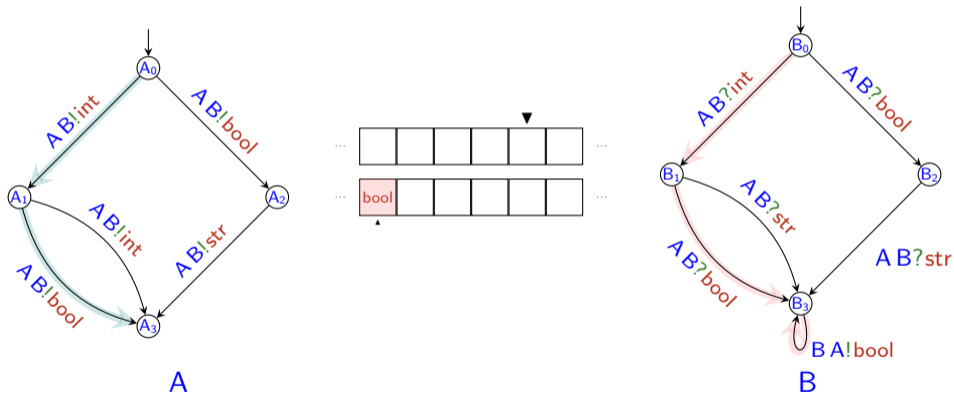
Communicating Finite-State Machines [Brand&Zafiropulo 1983]



Global specs can be projected (i.e., compiled) on CFSMs

# A model of local specs

Communicating Finite-State Machines [Brand&Zafiropulo 1983]

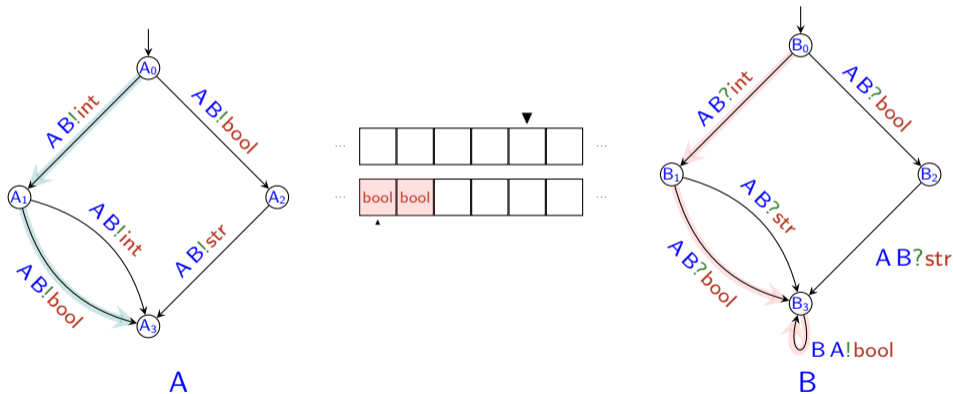


Global specs can be projected (i.e., compiled) on CFSMs



# A model of local specs


Communicating Finite-State Machines [Brand&Zafiropulo 1983]



Global specs can be projected (i.e., compiled) on CFSMs

# A model of local specs

Communicating Finite-State Machines [Brand&Zafiropulo 1983]



Global specs can be projected (i.e., compiled) on CFSMs

An obvious (fundamental) question

Given a global specification, is it  
**realisable** distributively?

## An obvious (fundamental) question

Given a global specification, is it **realisable** distributively?

Put simply...



A global spec  $G$  is **realizable** if there is a **deadlock-free**<sup>a</sup> system of CFSMs whose traces “**have some relation with**”  $G$ .

---

<sup>a</sup>A system  $S$  is *deadlock-free* if none of its reachable configurations  $s$  is a deadlock, that is  $s \nrightarrow$  and either some buffers are not empty or some CFSMs have transitions from their state in  $s$ .



# Class test

## Revisiting our examples



# Class test

## Revisiting our examples



## A (main) source of problems: Well-branchedness

### Distributed consensus

A distributed choice  $G_1 + G_2 + \dots$  is **well-branched** if

- there is **one active** participant
- any non-active participant is **passive**

## A (main) source of problems: Well-branchedness

### Distributed consensus

A distributed choice  $G_1 + G_2 + \dots$  is **well-branched** if

- there is **one active** participant
- any non-active participant is **passive**

**Def.** **A** is **active** when it **locally** decides which branch to take in a choice

**Def.** **B** is **passive** when

- either **B** behaves uniformly in **each branch**
- or **B** “unambiguously understands” which branch **A** opted for from some inputs



# A (main) source of problems: Well-branchedness

## Distributed consensus

A distributed choice  $G_1 + G_2 + \dots$  is **well-branched** if

- there is **one active** participant
- any non-active participant is **passive**

**Def.**  $A$  is **active** when it **locally** decides which branch to take in a choice

**Def.**  $B$  is **passive** when

- either  $B$  behaves uniformly in **each branch**
- or  $B$  “unambiguously understands” which branch  $A$  opted for from some inputs

## Well-branchedness

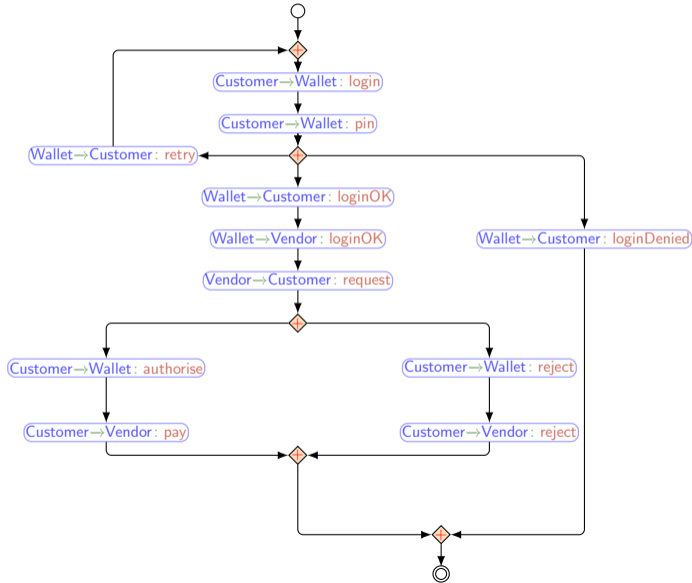
When the above holds true for each choice, the choreography is **well-branched**. This enables **correctness-by-design**.

– Act II –

# Choreography Automata

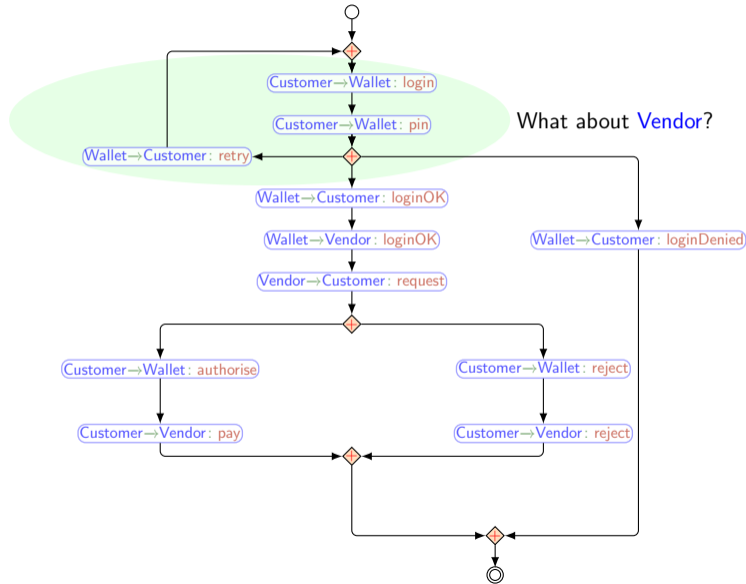
(joint work with Franco Barbanera, Ivan Lanese)

# The online-wallet protocol



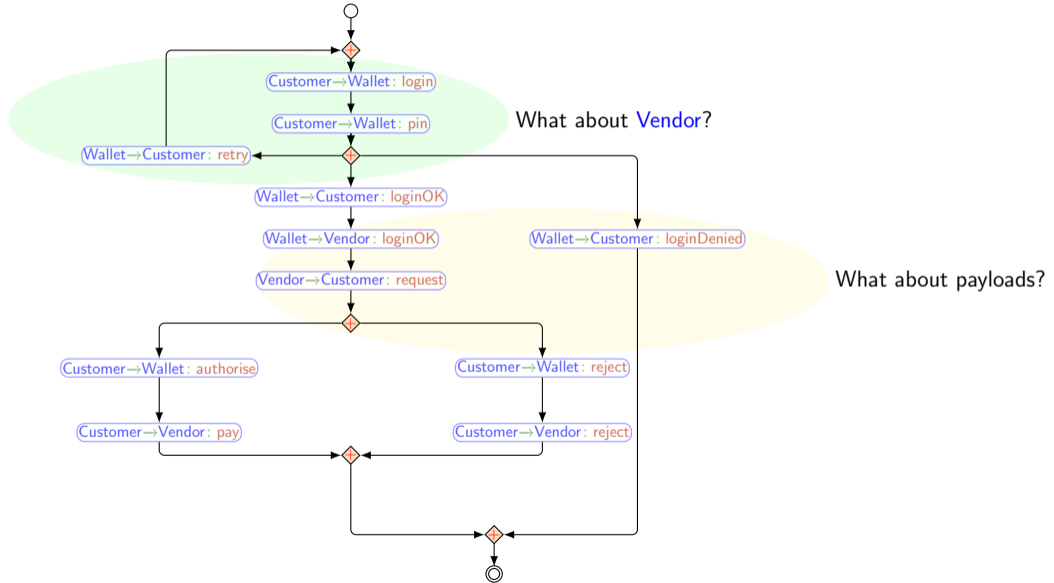
# The online-wallet protocol

...some modelling problems




# The online-wallet protocol

...some modelling problems




# Our global & local specs

## Choreography automata: Interaction, globally




# Our global & local specs

## Intermediate automata: from interactions to communications

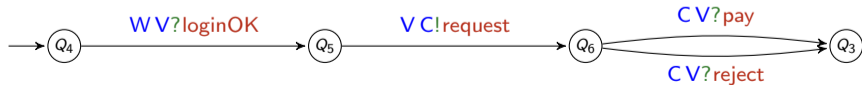


# Our global & local specs

## Intermediate automata: from interactions to communications



## CFSMs locally: determinise the intermediate automaton






# Projections preserve semantics

**Theorem.** Choreography automata are bisimilar to their projections

$\implies$  traces equivalence


# Flexibility by example

## Selective participation in OLW



# Flexibility by example


## Selective participation in OLW



- at  $q_2$  **Wallet** and **Customer** aware from the very beginning

# Flexibility by example


## Selective participation in OLW



- at  $q_2$  **Wallet** and **Customer** aware from the very beginning
  - **Vendor** involved on one branch only, but that's fine: **Wallet** is aware

# Flexibility by example


## Selective participation in OLW



- at  $q_2$  **Wallet** and **Customer** aware from the very beginning
  - **Vendor** involved on one branch only, but that's fine: **Wallet** is aware
- at  $q_6$  **Wallet** and **Customer** aware from the very beginning

# Flexibility by example

## Selective participation in OLW



- at  $q_2$  **Wallet** and **Customer** aware from the very beginning
  - **Vendor** involved on one branch only, but that's fine: **Wallet** is aware
- at  $q_6$  **Wallet** and **Customer** aware from the very beginning
  - **Vendor** eventually informed by **Customer** on each branch


## Correctness by construction

**Theorem.** Projections of well-formed choreography automata are deadlock-free

**Theorem.** Projections of well-formed choreography automata are lock-free

# DbC vs. choreography automata


## Asserting (an excerpt of) OLW





# DbC vs. choreography automata

## Asserting (an excerpt of) OLW



## Consistency

- **history senesitiveness**: in  $q \xrightarrow{\lambda}_A q'$ ,  $A$  predicates on *known* variables
- **temporal satisfiability**: the conjunction of the predicates on a path is satisfiable
- well-formedness of the underlying choreography automaton

# Theorems

Projections are a bit more complicated than for choreography automata

On consistent asserted choreography automata

**Theorem.** Asserted choreography automata are **weakly** bisimilar to their projections

$\implies$  trace equivalence

**Theorem.** Projections of WF choreography automata are deadlock-free

# Theorems

Projections are a bit more complicated than for choreography automata

On consistent asserted choreography automata

**Theorem.** Asserted choreography automata are **weakly** bisimilar to their projections

⇒ trace equivalence

**Theorem.** Projections of WF choreography automata are deadlock-free

And more...cf. [ECOOP 2022]

A tool chain for


- **validating finitary Scribble protocols** via choreography automata
- TypeScript **web programming** via API generation

– Act III –

Local-first!

(joint work with Daniela Marottoli, Hernán Melgratti, Roland Kuhn)

# A completely different setting




## Desiderata

- different features
  - arbitrary (and variable) number of instances
  - **local-first principle!**
    - As rock climbers say: “Don’t Be Afraid To Fail. Be Afraid Not To Try.”
  - pub-sub (instead of point-to-point)
- different properties
  - progress despite unavailability
    - ⇒ inconsistent views
  - eventual-consistency instead of “old” properties (eg. session fidelity)


# Swarm protocols and machines by example

## Global




(log types omitted for readability)

## Local (projected)




# Swarm protocols and machines by example

## Global



(log types omitted for readability)


## Local (projected)



local log:  $r_1 \cdot r_2 \cdot b$


# Swarm protocols and machines by example

## Global



(log types omitted for readability)

## Local (projected)




local log:  $\cdot r_2 \cdot b$




# Swarm protocols and machines by example

## Global



(log types omitted for readability)


## Local (projected)



local log: · · b


# Swarm protocols and machines by example

## Global



(log types omitted for readability)

## Local (projected)



local log: . . and now select is enabled

# Semantics, intuitively

- Types “produce/consume” events
  - **swarm protocols**: how/when **roles produce** events
  - **machines**: how/when **instances consume** events “skipping” the ones irrelevant to them
- Deterministic types only
  - **swarm protocols**: log types of branches have no common non-trivial prefixes and command/role pairs are pairwise distinct
  - **machines**: event types of branches are pairwise distinct
- Non-deterministic events’ propagation

# Swarms

## Machines, local logs, and global log (...a mirage)

Events are univocally associated to the machines generating them.

**Def.** *swarm* = global log + map from unique identities to pairs machines/local logs

$$(S, l) = (M_1, l_1) \mid \dots \mid (M_n, l_n) \mid l$$

such that  $l_i \sqsubseteq l$  where,

$$l_i \sqsubseteq l \iff l_i = \begin{matrix} e_{i,1} \\ \vdots \\ e_{i,n} \end{matrix} \quad \begin{matrix} e_1 \\ \vdots \\ e_m \end{matrix} = l$$

i.e., there is an order-preserving and downward-total morphism from  $l_i$  into  $l$  on events of a same machine.

# Swarms

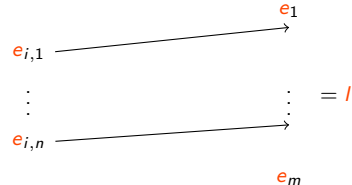
## Machines, local logs, and global log (...a mirage)

Events are univocally associated to the machines generating them.

**Def.** *swarm* = global log + map from unique identities to pairs machines/local logs

$$(S, l) = (M_1, l_1) \mid \dots \mid (M_n, l_n) \mid l$$

such that  $l_i \sqsubseteq l$  where,

$$l_i \sqsubseteq l \iff l_i = \begin{array}{c} \vdots \\ \vdots \\ \vdots \end{array} = l$$


i.e., there is an order-preserving and downward-total morphism from  $l_i$  into  $l$  on events of a same machine.


# Swarms

## Machines, local logs, and global log (...a mirage)

Events are univocally associated to the machines generating them.

**Def.** *swarm* = global log + map from unique identities to pairs machines/local logs

$$(S, l) = (M_1, l_1) \mid \dots \mid (M_n, l_n) \mid l$$



i.e., there is an order-preserving and downward-total morphism from  $l_i$  into  $l$  on events of a same machine.

# Swarms

## Machines, local logs, and global log (...a mirage)


Events are univocally associated to the machines generating them.

**Def.** *swarm* = global log + map from unique identities to pairs machines/local logs

$$(S, l) = (M_1, l_1) \mid \dots \mid (M_n, l_n) \mid l$$

such that  $l_i \sqsubseteq l$  where,

$$l_i \sqsubseteq l \iff$$



i.e., there is an order-preserving and downward-total morphism from  $l_i$  into  $l$  on events of a same machine.

## Swarms' semantics...intuitively

- **Events' generation**  
The local log of a machine is extended with the fresh events generated by (the execution of a command on) the machine
- **Events' propagation**  
Emitted events propagate asynchronously & non-deterministically



## Swarms' semantics: formally

[LOCAL]

$$\frac{\mathbf{S} : i \mapsto (\mathbf{M}, l) \quad (\mathbf{M}, l) \xrightarrow{c/l} (\mathbf{M}, l') \quad l'' \in l \bowtie \hat{l}}{(\mathbf{S}, \hat{l}) \xrightarrow{c/l} (\mathbf{S}[i \mapsto (\mathbf{M}, l')], l'')}$$

where  $l_1 \bowtie l_2 = \{l \mid l \subseteq l_1 \cup l_2 \wedge l_1 \sqsubseteq l \wedge l_2 \sqsubseteq l\}$

[PROP]

$$\frac{\mathbf{S} : i \mapsto (\mathbf{M}, l) \quad l \sqsubseteq l' \sqsubseteq \hat{l} \quad l \subset l'}{(\mathbf{S}, \hat{l}) \xrightarrow{\tau} (\mathbf{S}[i \mapsto (\mathbf{M}, l')], \hat{l})}$$

# Properties of our semantics

## Coherence

A swarm  $(M_1, I_1) \mid \dots \mid (M_n, I_n) \mid I$  is **coherent** if

$$\text{for all } i, I_i \subseteq I \quad \text{and} \quad I = \bigcup_{i \in \underline{n}} I_i$$

## Coherence preservation

[LOCAL] & [PROP] preserve coherence

## Eventual Consistency

If

$S = (M_1, I_1) \mid \dots \mid (M_n, I_n) \mid I$  is coherent

then

$$S \xrightarrow{\tau}^* (M_1, I) \mid \dots \mid (M_n, I) \mid I$$

# Realisation

It is hard to get it right (even **without** multi-instances or choices!)

A trivial protocol

Take the swarm protocol



Are

request / requested



and

bid / bid



a

realisation?

# Realisation

It is hard to get it right (even **without** multi-instances or choices!)

A trivial protocol

Take the swarm protocol



Are

request / requested



and

bid / bid



a (**correct**) realisation?

## Realisation

It is hard to get it right (even **without** multi-instances or choices!)

A trivial protocol

Take the swarm protocol



Are

request / requested



and

bid / bid



a (**correct**) realisation? What does that actually mean?

## Not so simple

A swarm correctly realises a swarm protocol if it generates only logs that the protocol can generate.

That's impossible due to events' skipping at local level but not at the global one.

## Not so simple

A swarm correctly realises a swarm protocol if it generates only logs that the protocol can generate.

That's impossible due to events' skipping at local level but not at the global one.

## A weaker condition

A swarm correctly realises a swarm protocol if it generates only logs that are **admissible** with some that the protocol can generate.

A log is **admissible** for a swarm protocol when its restriction to the events processed by the active machines is equivalent to a log of the protocol.

## Well-formedness of swarm protocols

Each log type  $\mathbf{l}$  of a branch should be

- causal consistent
  - each selector in (the continuation of)  $\mathbf{l}$  reacts to  $\mathbf{l}$
  - each role involved in the continuation of  $\mathbf{l}$  cannot react to more events on  $\mathbf{l}$  than selectors on the branch
- determined
  - each role in the continuation of  $\mathbf{l}$  reacts to  $\mathbf{l}[0]$
- confusion-free
  - an event type cannot occur in more than one branch



– Epilogue –

[

...

]

# Summing up

## Automata models for choreography

### Advantages

- increased flexibility
- good basis for (enhanced) tool support
- good also for practitioners

## Plans

- weakening well-formedness conditions
- studying more complex communication models (eg non-atomic propagation of events)

[

Thank you!

]