

Yet Another Reo Semantics: Reasoning about Speculative Execution

Hans-Dieter A. Hiep

Vrije Universiteit Amsterdam (MSc student)
Centrum Wiskunde & Informatica (Intern)

September 25, 2018

Overview

- ▶ Master's thesis:
Yet Another Reo Semantics: Reasoning about Speculative Execution

Overview

- ▶ Master's thesis:
Yet Another Reo Semantics: Reasoning about Speculative Execution
- ▶ Outline:
 1. Background and Challenges

Overview

- ▶ Master's thesis:
*Yet Another **Reo** Semantics: Reasoning about Speculative Execution*
- ▶ Outline:
 1. Background and Challenges
 2. Syntax
typed language for compositions and components

Overview

- ▶ Master's thesis:
*Yet Another Reo **Semantics**: Reasoning about Speculative Execution*
- ▶ Outline:
 1. Background and Challenges
 2. Syntax
typed language for compositions and components
 3. Semantics
logical framework with data types and data streams

Overview

- ▶ Master's thesis:
Yet Another Reo Semantics: Reasoning about Speculative Execution
- ▶ Outline:
 1. Background and Challenges
 2. Syntax
typed language for compositions and components
 3. Semantics
logical framework with data types and data streams
 4. Logical Analysis
independence, progress, linearity, and causality

Overview

- ▶ Master's thesis:
*Yet Another Reo Semantics: Reasoning about **Speculative Execution***
- ▶ Outline:
 1. Background and Challenges
 2. Syntax
typed language for compositions and components
 3. Semantics
logical framework with data types and data streams
 4. Logical Analysis
independence, progress, linearity, and causality
- ▶ Throughline:
Speculative Execution

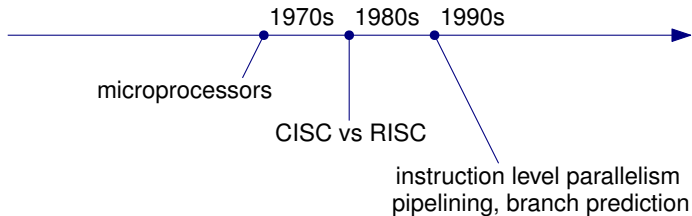
Overview

- ▶ Master's thesis:
Yet Another Reo Semantics: Reasoning about Speculative Execution
- ▶ Outline:
 1. Background and Challenges
 2. Syntax
typed language for compositions and components
 3. Semantics
logical framework with data types and data streams
 4. Logical Analysis
independence, progress, linearity, and causality
- ▶ Throughline:
Speculative Execution

Background and Challenges

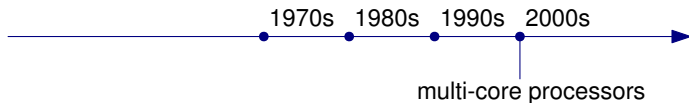
- ▶ Increasing demand for computation and communication
- ▶ Challenge to meet demand:
 - ▶ Increase **throughput**
 - ▶ Reduce **energy costs**
- ▶ Scalable architecture

Timeline



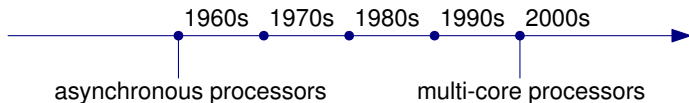
- ▶ RISC allows more complex optimizations, e.g. **speculative execution**

Timeline



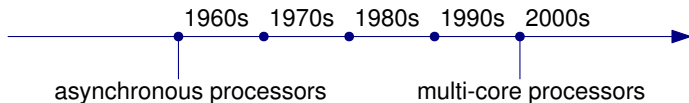
- ▶ Multiple processors on single die improve **throughput**
- ▶ Bounding clock speeds reduces **energy costs**

Timeline



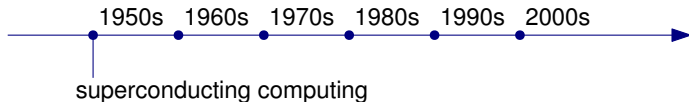
- ▶ Challenge of asynchronous processors is **concurrency**
- ▶ Concurrency resurrected by multi-core

Timeline



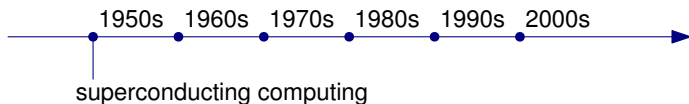
- ▶ Delay insensitivity improve **throughput**
- ▶ Asynchronous processors reduce **energy costs**

Timeline



- ▶ Cool your circuit to 4 degree Kelvin
- ▶ Superconductivity: zero electrical resistance on wires
- ▶ Josephson junction
- ▶ IARPA's Cryogenic Computing Complexity (C3) works towards switching speeds of ~1ns

Timeline



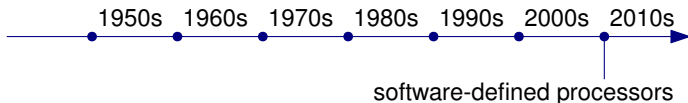
Intuition:

- ▶ Electrons flowing by nano pulses
- ▶ Computation is controlling the flow
- ▶ “The actual usefulness of computation lies in getting rid of the haystack, leaving only the needle.”

cf. Two puzzles about computation, S. Abramsky, 2014

cf. Irreversibility and heat generation in the c. process, R. Landauer, 1961

Timeline



- ▶ **compile to hardware**
- ▶ Costs of tape-out of integrated circuits has dropped:
hundred 1x1mm chips @ 28nm for only \$14,000
- ▶ Hennessy and Patterson's 2017 Turing Award Lecture

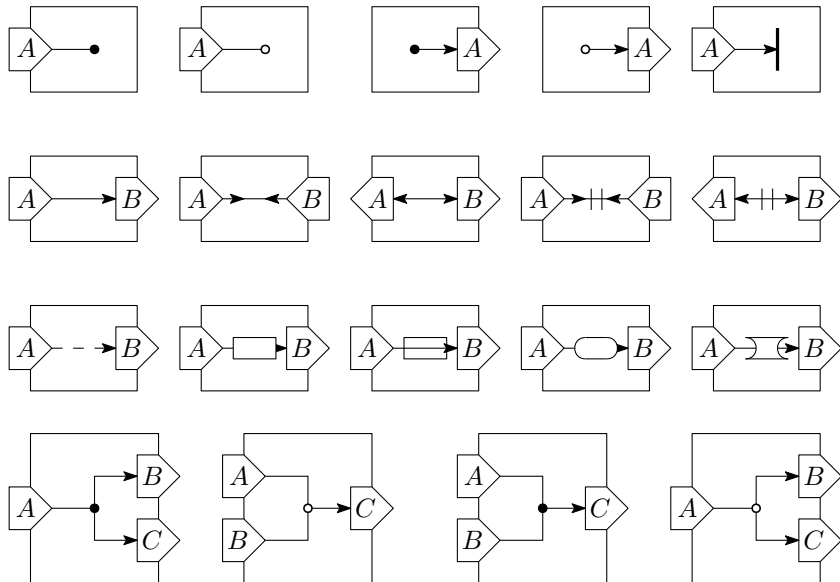
Background and Challenges

- ▶ Scalable architecture
 - ▶ 1 CPU $\approx 10^9$ transistors
 - ▶ 1 Internet $\approx 10^9$ CPUs
- ▶ Modular software correctness
- ▶ Polymorphic software scaling
- ▶ Reo

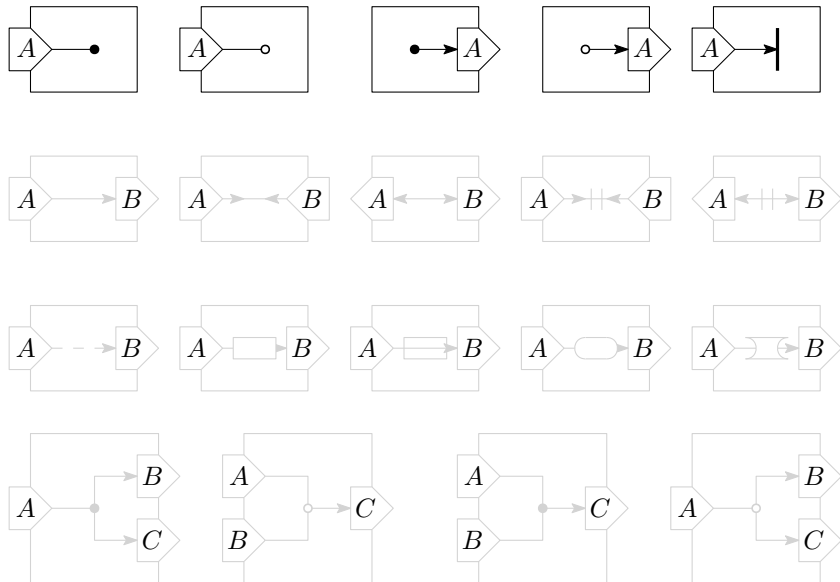
Overview

- ▶ Master's thesis:
Yet Another Reo Semantics: Reasoning about Speculative Execution
- ▶ Outline:
 1. Background and Challenges
 2. *Syntax*
typed language for compositions and components
 3. Semantics
logical framework with data types and data streams
 4. Logical Analysis
independence, progress, linearity, and causality
- ▶ Throughline:
Speculative Execution

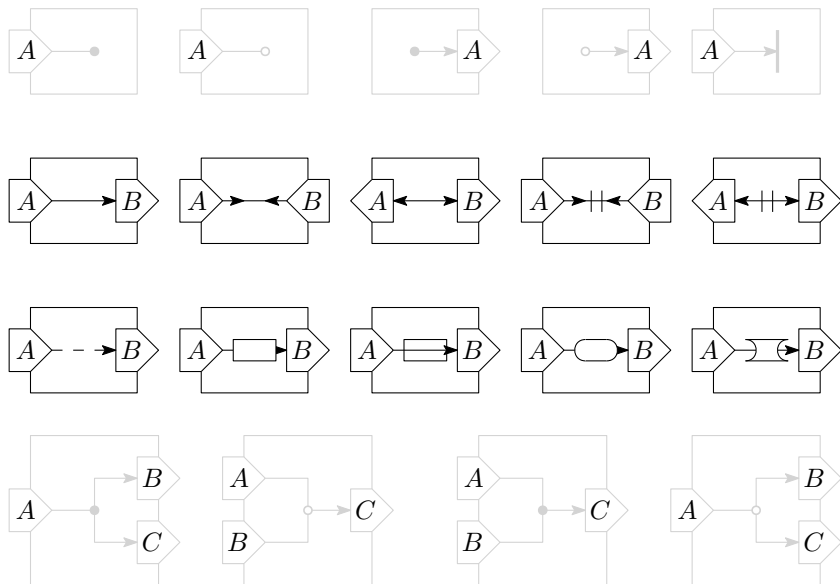
Standard Components



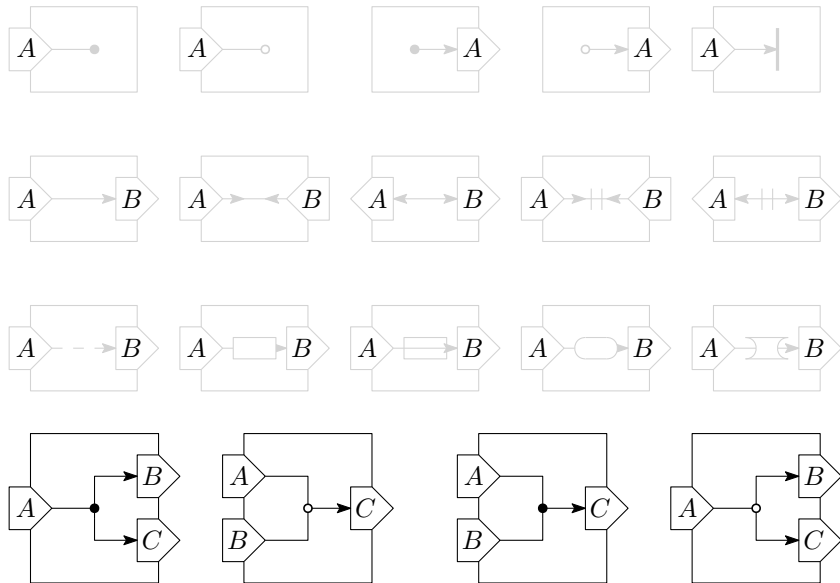
Unary Components



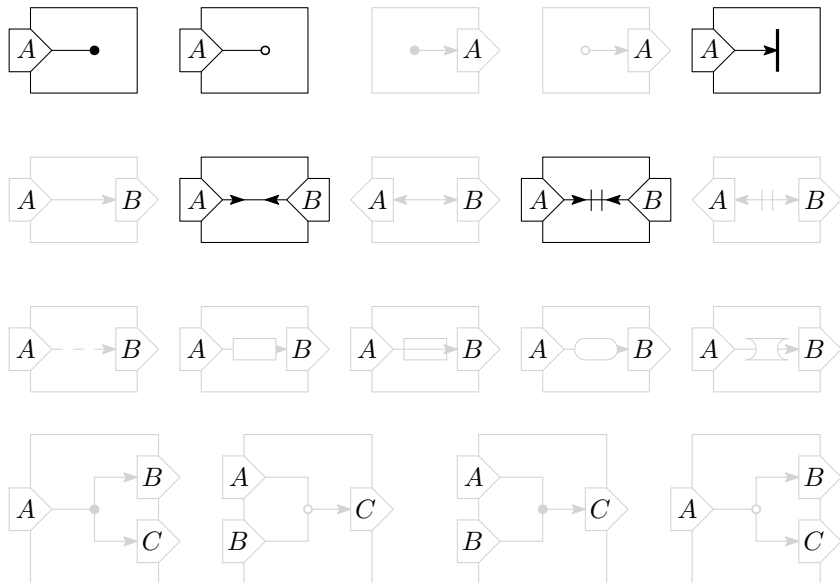
Binary Components



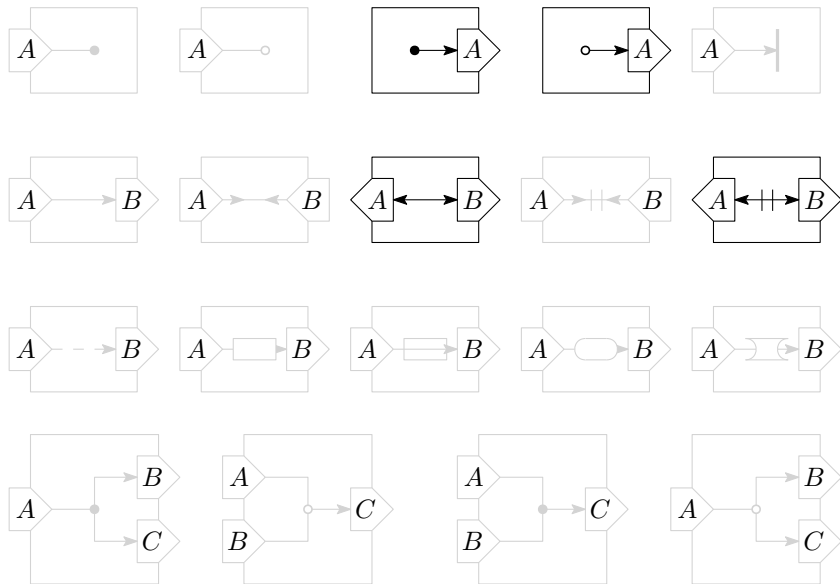
Ternary Components



Input-only Components



Output-only Components

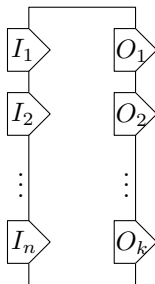


Interface

Definition

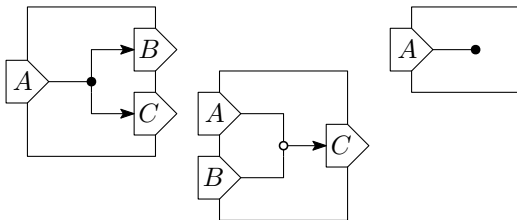
An *interface* $\langle I_1 : \alpha_1, \dots, I_n : \alpha_n \mid O_1 : \beta_1, \dots, O_k : \beta_k \rangle$ consists of:

1. A sequence of **input** ports I_1, \dots, I_n
2. A sequence of **output** ports O_1, \dots, O_k
3. A type assignment $I_j : \alpha_j$ and $O_j : \beta_j$ for data types α_j, β_j

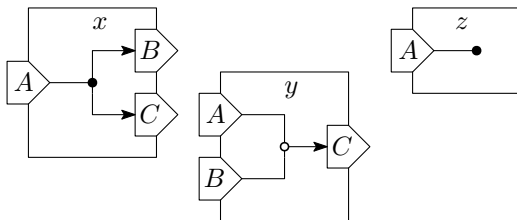


Given interface U , by U^\perp we denote its **dual**.

Instances and References

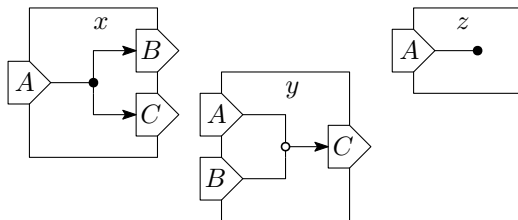


Instances and References



Component instances x, y, z, \dots

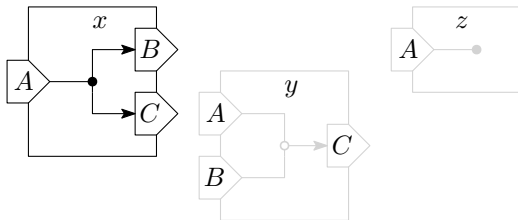
Instances and References



Component instances x, y, z, \dots

Qualified $x.A$ and **unqualified** A .

Composition

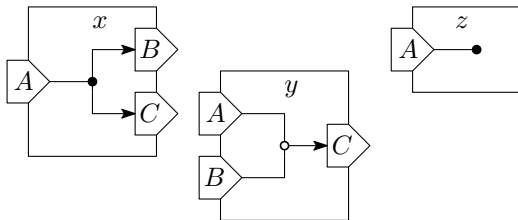


Definition

A *composition* is either:

- ▶ an instance

Composition

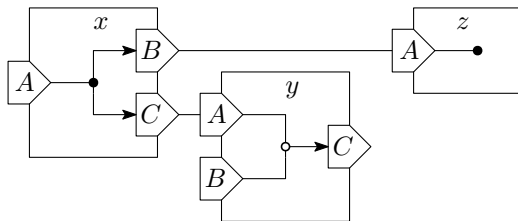


Definition

A *composition* is either:

- ▶ an instance
- ▶ two compositions adjoined

Composition

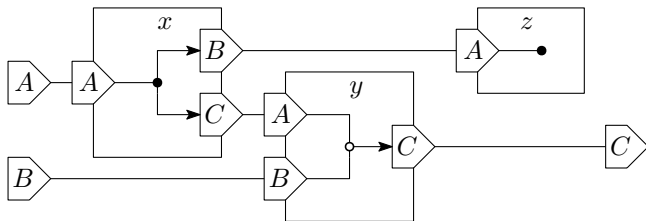


Definition

A composition is either:

- ▶ an instance
- ▶ two compositions adjoined
- ▶ identification of two references of a composition

Composition



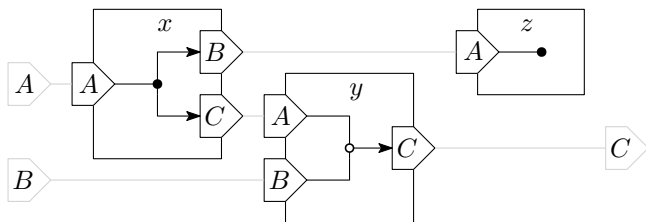
Definition

A *composition* is either:

- ▶ an instance
- ▶ two compositions adjoined
- ▶ identification of two references of a composition

Well-formed if no unqualified references

Component

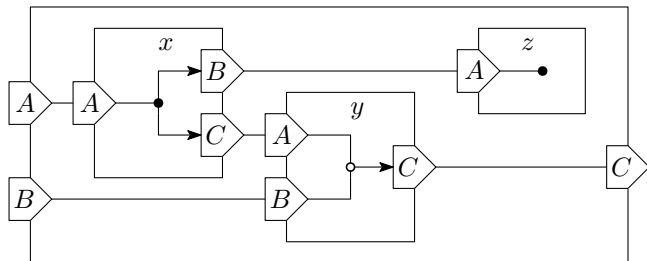


Definition (to be continued)

A *component* is either:

- ▶ primitive

Component



Definition (to be continued)

A *component* is either:

- ▶ primitive
- ▶ composite

Results

1. Formal graphical language for Reo
2. Designed normalization and type checker
3. Implemented prototype (see garbage can)

cf. A survey of graphical languages for monoidal categories,
P. Selinger, 2009

cf. Sequent calculus: a logic and a language for computation and duality,
P. Downen, 2017

Overview

- ▶ Master's thesis:
Yet Another Reo Semantics: Reasoning about Speculative Execution
- ▶ Outline:
 1. Background and Challenges
 2. Syntax
typed language for compositions and components
 3. Semantics
logical framework with data types and data streams
 4. Logical Analysis
independence, progress, linearity, and causality
- ▶ Throughline:
Speculative Execution

Data domains

Data types α, β, \dots

- ▶ contain 'null' value $*$
- ▶ are (infinitely) enumerable
- ▶ are regularly structured (...)

Examples

Signal = $\{*, 0\}$ and Nat = $\{*, 0, 1, 2, \dots\}$

Stream domains

Given data type α, β, \dots

Data streams $(\mathbb{N} \rightarrow \alpha), (\mathbb{N} \rightarrow \beta), \dots$

- ▶ streams are not enumerable

Examples

σ is a data stream over Signal

head $\sigma(0)$ is $*$ or 0

tail σ' is stream derivative

(cf. On Streams and Coinduction, J.J.M.M. Rutten, 2002)

Multi-sorted logic

Sorts:

- ▶ each data type is a distinct sort
- ▶ each data stream is a distinct sort
- ▶ there is a natural sort \mathbb{N}

Non-logical symbols:

- ▶ $*_{\alpha}$ null constant
- ▶ d_{α} data constant
- ▶ \perp proposition
- ▶ $=$ equality predicate
- ▶ $0, 1, 2, \dots$ and $+, -, \times, \leq$ for naturals
- ▶ at_{α} with arity $\langle (\mathbb{N} \rightarrow \alpha), \mathbb{N}, \alpha \rangle$

Multi-sorted logic

Terms: standard

- ▶ variables x^s
- ▶ constants
- ▶ function symbols

Formulas: standard

- ▶ propositions
- ▶ predicates of terms
- ▶ standard logical symbols $\neg, \wedge, \vee, \rightarrow$
- ▶ first-order quantification $\exists x^s, \forall x^s$ binders for x^s

We write $\text{at}(X, t)$ as $X(t)$, called *applications*.

Definition

A *coordination protocol* is a first-order formula such that all free variables X^s are of a data stream sort $s = (\mathbb{N} \rightarrow \alpha)$

Coordination protocols

Coordination protocols induce a set of infinite tables
(cf. Rule-Based Form for Stream Constraints, K. Dokter, 2018)

Examples

Let X be a port of data type `Signal`.

Consider $\phi = \forall t.(X(t) = * \vee X(t) = 0)$

Solution: assignment of X to a data stream that sats. formula

$$\mathcal{L}(\phi) = \left\{ \begin{array}{c|c} \hline t & X \\ \hline 0 & * \\ 1 & * \\ 2 & * \\ \vdots & \vdots \\ \hline \end{array} , \dots, \begin{array}{c|c} \hline t & X \\ \hline 0 & * \\ 1 & 0 \\ 2 & * \\ \vdots & \vdots \\ \hline \end{array} , \dots, \begin{array}{c|c} \hline t & X \\ \hline 0 & 0 \\ 1 & 0 \\ 2 & 0 \\ \vdots & \vdots \\ \hline \end{array} , \dots \right.$$

Coordination protocols

Fact: coordination protocols only 'interact' for shared variables

Examples (No interaction)

Consider $\phi = \forall t.(X(t) = *)$ and $\psi = \forall t.(Y(t) = 0)$

$$\mathcal{L}(\phi) = \overline{\begin{array}{cc} t & X \\ \hline 0 & * \\ 1 & * \\ \vdots & \vdots \end{array}} \quad \text{and} \quad \mathcal{L}(\psi) = \overline{\begin{array}{cc} t & Y \\ \hline 0 & 0 \\ 1 & 0 \\ \vdots & \vdots \end{array}}$$

Consider

$$\mathcal{L}(\phi \wedge \psi) = \overline{\begin{array}{ccc} t & X & Y \\ \hline 0 & * & 0 \\ 1 & * & 0 \\ \vdots & \vdots & \vdots \end{array}}$$

Coordination protocols

Fact: coordination protocols only 'interact' for shared variables

Examples (Interaction)

Consider $\phi = \forall t.(X(t) = *) \vee \forall t.(X(t) = 0)$ and $\psi = \exists t.(X(t) = 0)$

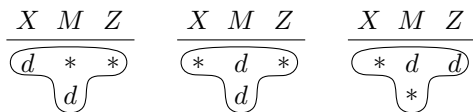
$$\mathcal{L}(\phi) = \left\{ \begin{array}{c|c} \overline{t} & \overline{X} \\ \hline 0 & * \\ 1 & * \\ \vdots & \vdots \end{array} , \begin{array}{c|c} \overline{t} & \overline{X} \\ \hline 0 & 0 \\ 1 & 0 \\ \vdots & \vdots \end{array} \right\} \text{ and } \mathcal{L}(\psi) = \left\{ \begin{array}{c|c} \overline{t} & \overline{X} \\ \hline 0 & * \\ 1 & 0 \\ \vdots & \vdots \end{array} , \begin{array}{c|c} \overline{t} & \overline{X} \\ \hline 0 & 0 \\ 1 & * \\ \vdots & \vdots \end{array} \right\}$$

Consider

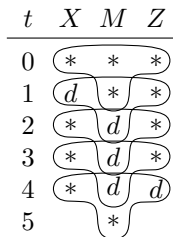
$$\mathcal{L}(\phi \wedge \psi) = \left\{ \begin{array}{c|c} \overline{t} & \overline{X} \\ \hline 0 & 0 \\ 1 & 0 \\ \vdots & \vdots \end{array} \right\}$$

Frame conditions

Insight: modeling of buffers using frame conditions.



These are overlapping as follows:



Frame conditions

Result: definition of buffer

$$\frac{X \quad M \quad Z}{\begin{array}{ccc} d & * & * \\ \text{---} & & \\ & d & \end{array}} \quad \frac{X \quad M \quad Z}{\begin{array}{ccc} * & d & * \\ \text{---} & & \\ & d & \end{array}} \quad \frac{X \quad M \quad Z}{\begin{array}{ccc} * & d & d \\ \text{---} & & \\ & * & \end{array}}$$

Examples

$$\forall t. ((\quad Z(t) = * \quad \wedge M(t) = * \wedge M(t+1) = X(t)) \vee \\ (X(t) = * \wedge Z(t) = * \quad \wedge M(t) \neq * \wedge M(t+1) = M(t)) \vee \\ (X(t) = * \wedge Z(t) = M(t) \wedge M(t) \neq * \wedge M(t+1) = *))$$

Frame conditions

Problem: column M is not a port. Can we get rid of it?

$$\begin{array}{c} X \quad Z \\ \hline * \quad * \end{array} \quad \begin{array}{c} X \quad Z \\ \hline \begin{array}{c} \boxed{\begin{array}{cc} d & * \\ * & * \\ * & d \end{array}} \end{array} \quad \circlearrowright$$

Examples

$$\forall t. (Z(t) = * \wedge X(t) = * \vee$$

$$(Z(t) = * \wedge \exists j. t < j \wedge X(j) = * \wedge Z(j) = X(t) \wedge$$

$$\forall i. t < i \wedge i < j \rightarrow X(i) = * \wedge Z(i) = *) \vee$$

$$(X(t) = * \wedge \exists j. j < t \wedge X(j) = Z(t) \wedge Z(j) = * \wedge$$

$$\forall i. j < i \wedge i < t \rightarrow X(i) = * \wedge Z(i) = *))$$

Components revisited

Definition

A *component* $U\phi$ is an interface U and coordination protocol ϕ .

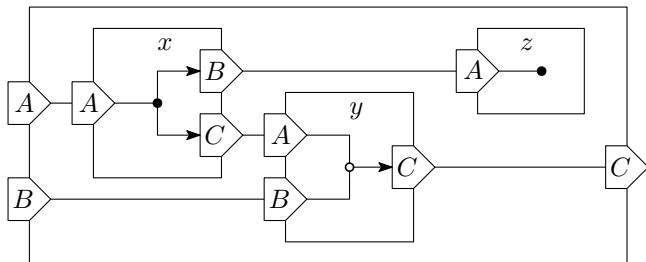
We lift composite constructions to this definition.

Primitive components given as $U\phi$.

Composite components by induction:

- ▶ Adjoined components $U\phi$ and $V\psi$:
take $\phi \wedge \psi$
- ▶ Identification of references X and Y on $U\phi$:
take $\phi \wedge \forall t.(X(t) = Y(t))$

Components revisited



$x :: \langle A : \alpha \mid B : \alpha, C : \alpha \rangle$

$\forall t. (A(t) = B(t) \wedge B(t) = C(t))$

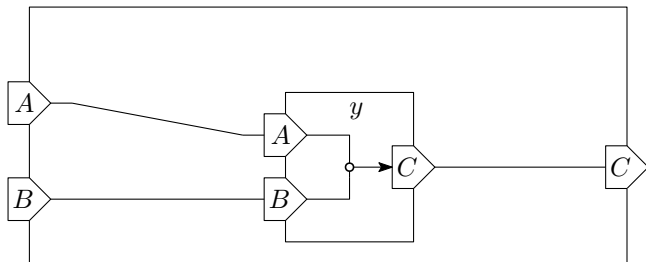
$y :: \langle A : \alpha, B : \alpha \mid C : \alpha \rangle$

$\forall t. ((A(t) = C(t) \wedge B(t) = *) \vee (B(t) = C(t) \wedge A(t) = *))$

$z :: \langle A : \alpha \mid \rangle$

$\forall t. (A(t) = * \vee A(t) \neq *)$

Components revisited



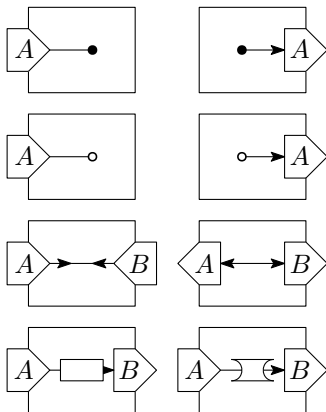
$\langle A : \alpha, B : \alpha \mid C : \alpha \rangle$

$\forall t. ((A(t) = C(t) \wedge B(t) = *) \vee$
 $(B(t) = C(t) \wedge A(t) = *))$

Duals

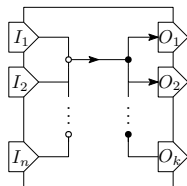
Definition

A dual of a component $U\phi$ is $U^\perp\phi$.



Results

1. Found a suitable logical framework with data types and data streams and defined primitives
2. Two different kinds of nodes (Reo has the one below)



1. Prophet is dual of buffer

cf. Rule-based form for stream constraints, K. Dokter, 2018

cf. The existence of refinement mappings, M. Abadi & L. Lamport, 1988

Overview

- ▶ Master's thesis:
Yet Another Reo Semantics: Reasoning about Speculative Execution
- ▶ Outline:
 1. Background and Challenges
 2. Syntax
typed language for compositions and components
 3. Semantics
logical framework with data types and data streams
 4. Logical Analysis
independence, progress, linearity, and causality
- ▶ Throughline:
Speculative Execution

Speculative Execution

A *speculation* predicts a future state.

- ▶ A *true speculation* if future is predicted.
- ▶ A *false speculation* if not.

Two methods:

- ▶ Parallel isolated worlds, discarding false worlds
- ▶ Branch prediction + backtracking

Interesting properties

- ▶ Independence
(delay insensitivity)
- ▶ Linearity
(superconducting/reversibility)
- ▶ Causality
(speculative execution)
- ▶ Instantaneous
(stateless)
- ▶ Synchronous and Asynchronous
(clocked and unclocked)
- ▶ Termination and Progress
(safety and liveness)

Independence

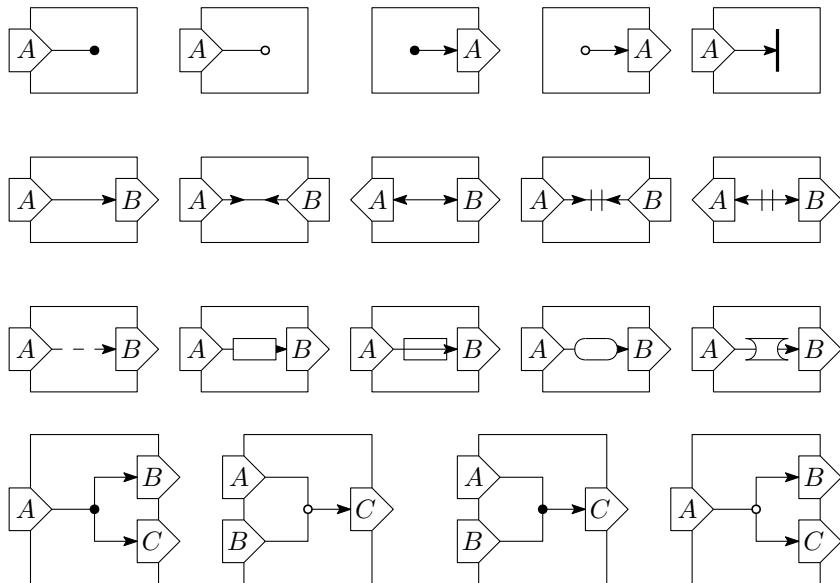
Semantic property

Independence is a closure condition on semantics.

$$\mathcal{L}(\phi) = \left\{ \overline{\begin{array}{cc} t & X \\ 0 & * \\ 1 & * \\ \vdots & \vdots \end{array}}, \dots, \overline{\begin{array}{cc} t & X \\ 0 & 0 \\ 1 & * \\ \vdots & \vdots \end{array}}, \dots \right.$$

- ▶ May remove rows with all *
 - ▶ May insert row with all * before row without all *
1. No real-time clocks (row depending on t).
 2. Independence is compositional

Independence



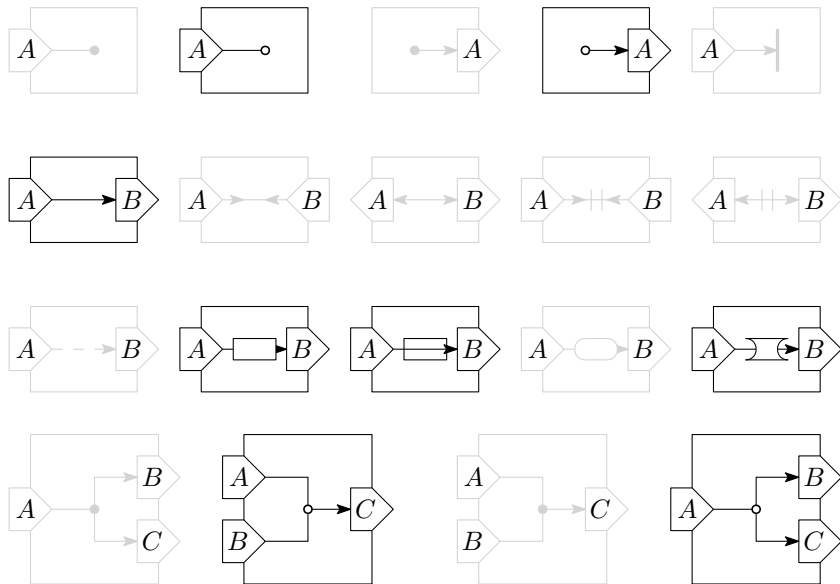
Linearity

Semantic property

Bijection between *input events*, *output events*

| t | X | Y |
|----------|----------|-----|
| 0 | * | d |
| 1 | d | * |
| 2 | e | e |
| 3 | f | * |
| 4 | * | f |
| \vdots | \vdots | |

Linearity



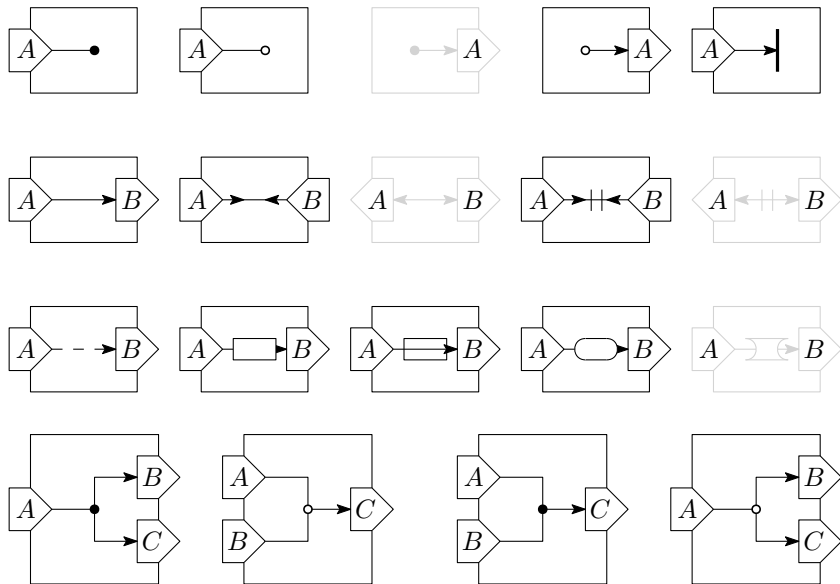
Causality

Semantic property

Every output event at time t related with input event at time $\leq t$.

| t | X | Y |
|----------|----------|-----|
| 0 | * | * |
| 1 | d | * |
| 2 | * | d |
| 3 | f | d |
| 4 | * | f |
| \vdots | \vdots | |

Causality



Instantaneous

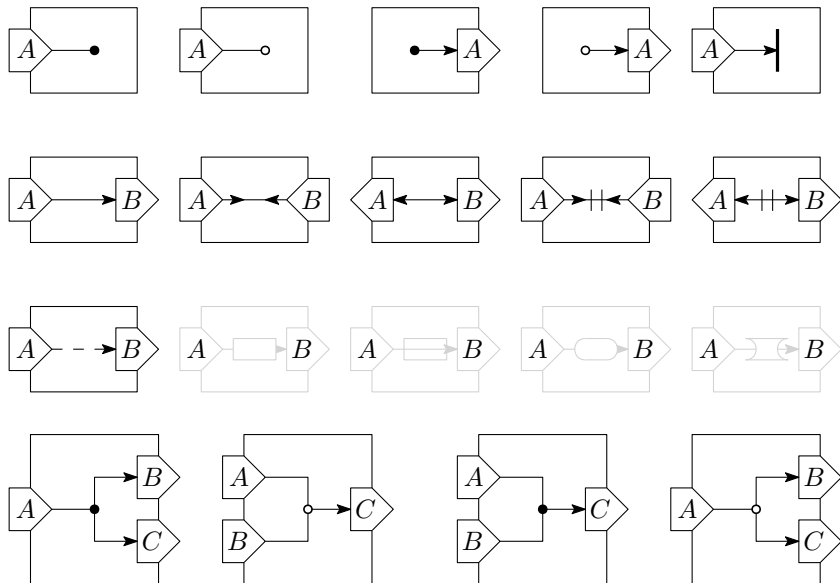
Syntactic property

Every formula of the shape:

$$\forall t. \phi$$

where ϕ only consists of applications with t .

Instantaneous



Synchronous

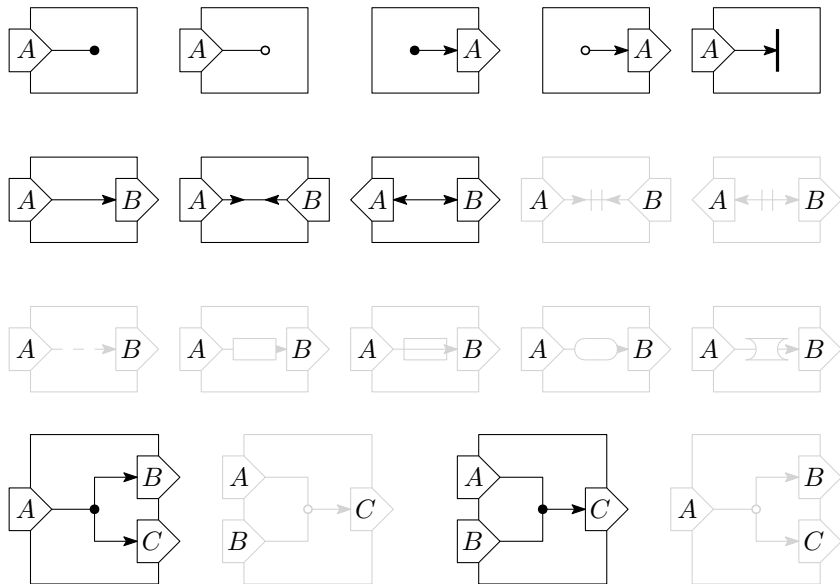
Logical property

A set of port is synchronous if all fire together, or none fire.

Let $P = \{X_1, \dots, X_n\}$

$$\forall t. ((X_1(t) = * \wedge \dots \wedge X_n(t) = *) \vee \\ (X_1(t) \neq * \wedge \dots \wedge X_n(t) \neq *))$$

Synchronous



Asynchronous

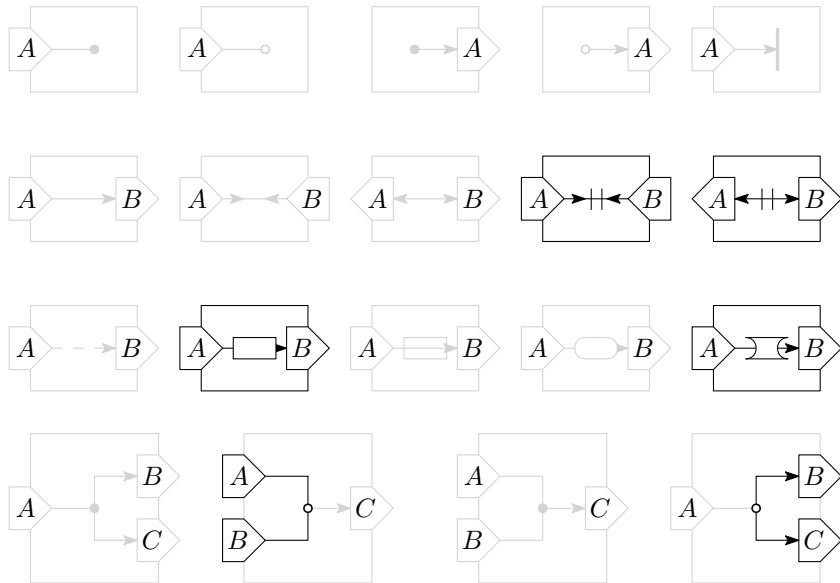
Logical property

A set of ports is asynchronous if at most one port fires.

Let $P = \{X_1, \dots, X_n\}$

$$\forall t. \bigwedge_{i \in \{1, \dots, n\}} \left(X_i(t) \neq * \rightarrow \bigwedge_{j \in \{1, \dots, n\} \setminus \{i\}} X_j(t) = * \right)$$

Asynchronous



Termination and Progress

Logical property

Progress:

Always, there is a port that eventually fires

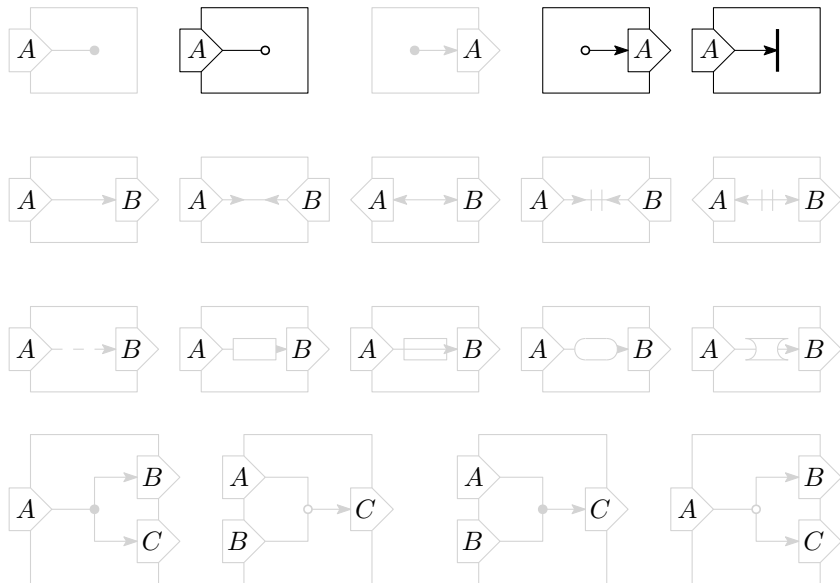
$$\forall t. (\exists s. (X_1(t+s) \neq * \vee \dots \vee X_n(t+s) \neq *))$$

Termination:

Eventually, all ports never fire anymore

$$\exists t. (\forall s. (X_1(t+s) = * \wedge \dots \wedge X_n(t+s) = *))$$

Termination and Progress



The End