

A Theory of System Fault Tolerance

Fossacs 06

Adrian Francalanza and Matthew Hennessy

`adrian.francalanza@um.edu.mt` `matthew.hennessy@sussex.ac.uk`

Aim of The Paper

- Formalise the notion of **Fault Tolerance** (in a distributed setting)
- Develop proof techniques to show fault-tolerance.


Talk Overview

- Fault Tolerance Intuitions
- Language
- Formal Definition
- Proof Techniques.

Talk Overview


- Fault Tolerance Intuitions
- Language
- Formal Definition
- Proof Techniques.

Fault Tolerance




processes
executing in
parallel and
interacting

Fault Tolerance



partitioned
across
container
units


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
is partial


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
preserved
up to 1 failure


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
preserved
up to 1 failure


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
preserved
up to 1 failure


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
preserved
up to 2 failures


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
preserved
up to 2 failures


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
preserved
up to 3 failures


Fault Tolerance



OBSERVER VIEW

Observed
behaviour
preserved
up to 3 failures

Fault Tolerance Analysis



Talk Overview

- Fault Tolerance Intuitions
- Language
- Formal Definition
- Proof Techniques.

The Language

Processes

$$\begin{aligned} P, Q ::= & u!\langle V \rangle.P & | & u?(X).P \\ & | \text{if } v=u \text{ then } P \text{ else } Q & | & * u?(X).P \\ & | (\nu n:T) P & | & \text{go } u.P \\ & | \mathbf{0} & | & P|Q \\ & | \text{ping } u.P \text{ else } Q & & \end{aligned}$$

Systems

$$\begin{aligned} M, N, O ::= & l[[P]] & | & N|M \\ & | (\nu n:T)N & & \end{aligned}$$

The Language

Assuming $\Gamma \vdash l : \mathbf{alive}$

(r-comm)

$$\frac{}{\Gamma \triangleright l[[a!\langle V \rangle.P]] \mid l[[a?(X).Q]] \longrightarrow \Gamma \triangleright l[[P]] \mid l[[Q\{V/X\}]]}$$

(r-go)

$$\frac{}{\Gamma \triangleright l[[\mathbf{go} \ k.P]] \longrightarrow \Gamma \triangleright k[[P]]} \quad \Gamma \vdash k : \mathbf{alive}$$

(r-ngo)

$$\frac{}{\Gamma \triangleright l[[\mathbf{go} \ k.P]] \longrightarrow \Gamma \triangleright k[[\mathbf{0}]]} \quad \Gamma \not\vdash k : \mathbf{alive}$$

The Language

Assuming $\Gamma \vdash l : \mathbf{alive}$

(r-ping)

$$\frac{}{\Gamma \triangleright l[[\text{ping } k.P \text{ else } Q]] \longrightarrow \Gamma \triangleright l[[P]]} \Gamma \vdash k : \mathbf{alive}$$

(r-nping)

$$\frac{}{\Gamma \triangleright l[[\text{ping } k.P \text{ else } Q]] \longrightarrow \Gamma \triangleright l[[Q]]} \Gamma \not\vdash k : \mathbf{alive}$$

Examples

$$\text{server}_1 \Leftarrow (\nu \text{ data}) \left(\begin{array}{l} l \llbracket \text{req?}(x, y). \text{ go } k_1. \text{ data!}\langle x, y, l \rangle \rrbracket \\ | k_1 \llbracket \text{data?}(x, y, z). \text{ go } z. y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{server}_1 \Leftarrow (\nu \text{ data}) \left(\begin{array}{l} l \llbracket \text{req?}(x, y). \text{ go } k_1. \text{ data!}\langle x, y, l \rangle \rrbracket \\ | k_1 \llbracket \text{data?}(x, y, z). \text{ go } z. y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{server}_1 \Leftarrow (\nu \text{ data}) \left(\begin{array}{l} l \llbracket \text{req?}(x, y). \text{ go } k_1. \text{ data!}\langle x, y, l \rangle \rrbracket \\ | k_1 \llbracket \text{data?}(x, y, z). \text{ go } z. y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{server}_2 \Leftarrow (\nu \text{ data}) \left(l \left[\left[\begin{array}{l} \text{go } k_1.\text{data!}\langle x, s, l \rangle \\ | \text{go } k_2.\text{data!}\langle x, s, l \rangle \\ | s?(x).y!\langle x \rangle \end{array} \right] \right] \right. \\ \left. \begin{array}{l} | k_1 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \\ | k_2 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{server}_2 \Leftarrow (\nu \text{data}) \left(l \left[\left[\begin{array}{l} \text{go } k_1.\text{data}!\langle x, s, l \rangle \\ | \text{go } k_2.\text{data}!\langle x, s, l \rangle \\ | s?(x).y!\langle x \rangle \end{array} \right] \right] \right. \\ \left. \begin{array}{l} | k_1 \llbracket \text{data}?(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \\ | k_2 \llbracket \text{data}?(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{server}_2 \Leftarrow (\nu \text{ data}) \left(\begin{array}{l} l \left[\left[\begin{array}{l} \text{go } k_1.\text{data!}\langle x, s, l \rangle \\ | \text{go } k_2.\text{data!}\langle x, s, l \rangle \\ | s?(x).y!\langle x \rangle \end{array} \right] \right] \\ | k_1 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \\ | k_2 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{server}_3 \Leftarrow (\nu \text{ data}) \left(l \left[\left[\begin{array}{l} \text{go } k_1.\text{data!}\langle x, s, l \rangle \\ | \text{go } k_2.\text{data!}\langle x, s, l \rangle \\ | \text{go } k_3.\text{data!}\langle x, s, l \rangle \\ | s?(x).y!\langle x \rangle \end{array} \right] \right] \right. \\ \left. \begin{array}{l} | k_1 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \\ | k_2 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \\ | k_3 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{server}_3 \Leftarrow (\nu \text{ data}) \left(l \left[\left[\begin{array}{l} \text{go } k_1.\text{data!}\langle x, s, l \rangle \\ | \text{go } k_2.\text{data!}\langle x, s, l \rangle \\ | \text{go } k_3.\text{data!}\langle x, s, l \rangle \\ | s?(x).y!\langle x \rangle \end{array} \right] \right] \right. \\ \left. \begin{array}{l} | k_1 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \\ | k_2 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \\ | k_3 \llbracket \text{data?}(x, y, z).\text{go } z.y!\langle f(x) \rangle \rrbracket \end{array} \right)$$


Examples

$$\text{sPing} \Leftarrow (v \text{ data}) \left(\begin{array}{l} l \left[\begin{array}{l} \text{serv?}(x, y).\text{ping } k_1.\text{go } k_1.\text{data!}\langle x, y, l \rangle \\ \text{else go } k_2.\text{data!}\langle x, y, l \rangle \end{array} \right] \\ | k_1 \llbracket \text{data?}(x, y, z).\text{go } z .y!\langle f(x) \rangle \rrbracket \\ | k_2 \llbracket \text{data?}(x, y, z).\text{go } z .y!\langle f(x) \rangle \rrbracket \end{array} \right)$$

Examples

$$\text{sPing} \Leftarrow (v \text{ data}) \left(\begin{array}{l} l \left[\begin{array}{l} \text{serv?}(x, y). \text{ping } k_1. \text{go } k_1. \text{data!}\langle x, y, l \rangle \\ \text{else go } k_2. \text{data!}\langle x, y, l \rangle \end{array} \right] \\ | k_1 \llbracket \text{data?}(x, y, z). \text{go } z. y! \langle f(x) \rangle \rrbracket \\ | k_2 \llbracket \text{data?}(x, y, z). \text{go } z. y! \langle f(x) \rangle \rrbracket \end{array} \right)$$


Server Fault Tolerance



$$\text{server}_1 = (v \text{ data}) \left(\begin{array}{l} l[\dots] \\ | k_1[\dots] \end{array} \right)$$

OBSERVER VIEW


Server Fault Tolerance



$$\text{server}_2 = (v \text{ data}) \begin{pmatrix} l[\dots] \\ | k_1[\dots] \\ | k_2[\dots] \end{pmatrix}$$

OBSERVER VIEW


Server Fault Tolerance



OBSERVER VIEW

$$\text{server}_2 = \begin{matrix} (v \text{ data}) \left(\begin{array}{l} l[\dots] \\ | k_1[\dots] \\ | k_2[\dots] \\ | k_3[\dots] \end{array} \right) \end{matrix}$$

Server Fault Tolerance



OBSERVER VIEW

$$\text{sPing} = (v \text{ data}) \begin{pmatrix} l[\dots] \\ | k_1[\dots] \\ | k_2[\dots] \end{pmatrix}$$

Talk Overview

- Fault Tolerance Intuitions
- Language
- **Formal Definition**
- Proof Techniques

Defining Fault Tolerance Preliminaries

- We partition Γ into two sets of live locations $\langle \mathcal{R}, \mathcal{U} \rangle$

Defining Fault Tolerance Preliminaries

- We partition Γ into two sets of live locations $\langle \mathcal{R}, \mathcal{U} \rangle$
Reliable: denoted by \mathcal{R} . They are **immortal!**

Defining Fault Tolerance Preliminaries

- We partition Γ into two sets of live locations $\langle \mathcal{R}, \mathcal{U} \rangle$
Reliable: denoted by \mathcal{R} . They are **immortal!**
Unreliable: denoted by \mathcal{U} . They **may fail!**

Defining Fault Tolerance Preliminaries

- We partition Γ into two sets of live locations $\langle \mathcal{R}, \mathcal{U} \rangle$
 - Reliable:** denoted by \mathcal{R} . They are **immortal!**
 - Unreliable:** denoted by \mathcal{U} . They **may fail!**
- We limit observations to reliable locations

Defining Fault Tolerance Preliminaries

- We partition Γ into two sets of live locations $\langle \mathcal{R}, \mathcal{U} \rangle$
Reliable: denoted by \mathcal{R} . They are **immortal!**
Unreliable: denoted by \mathcal{U} . They **may fail!**
- We limit observations to reliable locations
Contexts: for all $[-] \mid N$ we have $\mathbf{fl}(N) \subseteq \mathcal{R}$

Defining Fault Tolerance Preliminaries

- We partition Γ into two sets of live locations $\langle \mathcal{R}, \mathcal{U} \rangle$
Reliable: denoted by \mathcal{R} . They are **immortal!**
Unreliable: denoted by \mathcal{U} . They **may fail!**
- We limit observations to reliable locations
Contexts: for all $[-] | N$ we have $\mathbf{fl}(N) \subseteq \mathcal{R}$
Barbs: $\Gamma \triangleright M \Downarrow_{a@l}$ iff $\Gamma \triangleright M \longrightarrow^* \equiv \Gamma \triangleright (v\tilde{n}) M | l[[a!\langle V \rangle.P]]$
where $l, a \notin \tilde{n}$ and $l \in \mathcal{R}$

Defining Fault Tolerance Preliminaries

- We partition Γ into two sets of live locations $\langle \mathcal{R}, \mathcal{U} \rangle$
Reliable: denoted by \mathcal{R} . They are **immortal!**
Unreliable: denoted by \mathcal{U} . They **may fail!**
- We limit observations to reliable locations
Contexts: for all $[-] \mid N$ we have $\mathbf{fl}(N) \subseteq \mathcal{R}$
Barbs: $\Gamma \triangleright M \Downarrow_{a@l}$ iff $\Gamma \triangleright M \longrightarrow^* \equiv \Gamma \triangleright (v\tilde{n}) M \mid l[[a!\langle V \rangle.P]]$
where $l, a \notin \tilde{n}$ and $l \in \mathcal{R}$
- We define reduction barbed congruence \cong for configurations with the same reliable network \mathcal{R}

$$\langle \mathcal{R}, \mathcal{U} \rangle \triangleright M \cong \langle \mathcal{R}, \mathcal{U}' \rangle \triangleright N$$

Definition Fault Tolerance

Inducing Faults

Static: $\langle \mathcal{R}, \mathcal{U} \rangle - l = \langle \mathcal{R}, \mathcal{U} / \{l\} \rangle$
(r-kill)

Dynamic:
$$\frac{}{\Gamma \triangleright l[[\text{kill}]] \longrightarrow (\Gamma - l) \triangleright l[[\mathbf{0}]]}$$

Fault Contexts

Static: $F_S^n(\Gamma) = \Gamma - l_1 \dots - l_n$

Dynamic: $F_D^n(M) = M | l_1[[\text{kill}]] | \dots | l_n[[\text{kill}]]$

Definition Fault Tolerance

Static Fault Tolerance

$\Gamma \triangleright M$ is statically fault tolerant up to n faults if for any $F_S^n(-)$ we have

$$\Gamma \triangleright M \cong F_S^n(\Gamma) \triangleright M$$


Dynamic Fault Tolerance

$\Gamma \triangleright M$ is dynamically fault tolerant up to n faults if for any $F_D^n(-)$ we have

$$\Gamma \triangleright M \cong \Gamma \triangleright F_D^n(M)$$

Examples

Recall...



Examples

Good to show **negative** results. Assuming

$\Gamma = \langle \{l\}, \{k_1, k_2, k_3\} \rangle$:

- $\Gamma \triangleright \text{server}_1$ is **not** 1-statically fault tolerant because

$$\Gamma \triangleright \text{server}_1 \not\equiv \Gamma - k_1 \triangleright \text{server}_1$$

- $\Gamma \triangleright \text{server}_2$ is **not** 2-dynamically fault tolerant because

$$\Gamma \triangleright \text{server}_2 \not\equiv \Gamma \triangleright \text{server}_2 \mid k_1 \llbracket \text{kill} \rrbracket \mid k_2 \llbracket \text{kill} \rrbracket$$

- $\Gamma \triangleright \text{sPing}$ is **not** 1-dynamically fault tolerant because

$$\Gamma \triangleright \text{sPing} \not\equiv \Gamma \triangleright \text{sPing} \mid k_1 \llbracket \text{kill} \rrbracket$$

Examples

Hard to prove positive results:

It is difficult to prove that $\Gamma \triangleright \text{server}_2$ is 1-dynamic fault tolerant because:

1. \cong **quantifies over all valid contexts.**
2. Dynamic fault tolerance definition **quantifies over all fault contexts**, amongst which there is **considerable overlap.**
3. There are a number of **confluent** reductions that increase the burden of our analysis.

Talk Overview

- Fault Tolerance Intuitions
- Language
- Formal Definition
- **Proof Techniques**

Problems we need to address

Hard to prove positive results with our fault tolerance definition because:

1. \cong **quantifies over all valid contexts.**
2. Dynamic fault tolerance definition **quantifies over all fault contexts**, amongst which there is **considerable overlap**.
3. There are a number of **confluent** reductions that increase the burden of our analysis.

Solving Observer Quantification

- Define Its over configurations

(I-out)

$$\frac{}{\langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[a!\langle V \rangle.P]] \xrightarrow{l:a!\langle V \rangle} \langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[P]]} \quad l \in \mathcal{R}$$

Solving Observer Quantification

- Define Its over configurations

(I-out)

$$\frac{}{\langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[a!\langle V \rangle.P]] \xrightarrow{l:a!\langle V \rangle} \langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[P]]} \quad l \in \mathcal{R}$$

- Define bisimulation, \approx , for configurations based on Its

Solving Observer Quantification

- Define Its over configurations

(I-out)

$$\frac{}{\langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[a!\langle V \rangle.P]] \xrightarrow{l:a!\langle V \rangle} \langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[P]]} \quad l \in \mathcal{R}$$

- Define bisimulation, \approx , for configurations based on Its
- Prove **Soundness**:

$$\langle \mathcal{R}, \mathcal{U} \rangle \triangleright M \approx \langle \mathcal{R}, \mathcal{U}' \rangle \triangleright N$$

implies

$$\langle \mathcal{R}, \mathcal{U} \rangle \triangleright M \cong \langle \mathcal{R}, \mathcal{U}' \rangle \triangleright N$$

Problems we need to address

Hard to prove positive results with our fault tolerance definition because:

1. \cong **quantifies over all valid contexts.**
2. **Dynamic fault tolerance definition quantifies over all fault contexts, amongst which there is considerable overlap.**
3. There are a number of **confluent** reductions that increase the burden of our analysis.

...recall

Dynamic Fault Tolerance


$\Gamma \triangleright M$ is dynamically fault tolerant up to n faults if for any $F_D^n(-)$ we have

$$\Gamma \triangleright M \cong \Gamma \triangleright F_D^n(M)$$

Thus for every $F_D^n(-)$ we have to show


$$\Gamma \triangleright M \approx \Gamma \triangleright F_D^n(M)$$

Fault Context Quantification and Overlap (1)




Nodes are
bisimilar
tuples.
Edges are
transitions.

Fault Context Quantification and Overlap (1)




$$F_D^2 = \begin{cases} \text{kill}(k_1) \\ | \text{kill}(k_2) \\ | [-] \end{cases}$$

Fault Context Quantification and Overlap (1)



$$F_D^2 = \begin{cases} \text{kill}(k_1) \\ | \text{kill}(k_3) \\ | [-] \end{cases}$$

Fault Context Quantification and Overlap (1)



Merging
the two
relations

Fault Context Quantification and Overlap (2)

Define new transition that **counts** failures

(I-fail)

$$\frac{}{\langle \mathcal{R}, \mathcal{U} \rangle \triangleright N \xrightarrow{\text{fail}} \langle \mathcal{R}, \mathcal{U} \rangle - l \triangleright N} \quad l \in \mathcal{U}$$

Fault Context Quantification and Overlap (2)

Define **Fault Tolerant Simulation**, \leq_D^n , the largest *asymmetric* relation over configurations such that $\Gamma_1 \triangleright M_1 \leq_D^n \Gamma_2 \triangleright M_2$ implies

- $\Gamma_1 \triangleright M_1 \xrightarrow{\gamma} \Gamma'_1 \triangleright M'_1$ implies $\Gamma_2 \triangleright M_2 \xRightarrow{\widehat{\gamma}} \Gamma'_2 \triangleright M'_2$ such that $\Gamma'_1 \triangleright M'_1 \leq_D^n \Gamma'_2 \triangleright M'_2$
- $\Gamma_2 \triangleright M_2 \xrightarrow{\gamma} \Gamma'_2 \triangleright M'_2$ implies $\Gamma_1 \triangleright M_1 \xRightarrow{\widehat{\gamma}} \Gamma'_1 \triangleright M'_1$ such that $\Gamma'_1 \triangleright M'_1 \leq_D^n \Gamma'_2 \triangleright M'_2$
- **if $n > 0$** , $\Gamma_2 \triangleright M_2 \xrightarrow{\text{fail}} \Gamma'_2 \triangleright M'_2$ implies $\Gamma_1 \triangleright M_1 \xRightarrow{\widehat{\gamma}} \Gamma'_1 \triangleright M'_1$ such that $\Gamma'_1 \triangleright M'_1 \leq_D^{n-1} \Gamma'_2 \triangleright M'_2$

Fault Context Quantification and Overlap (2)

Give an alternative definition for Fault Tolerance up to n -dynamic faults.

$$\Gamma \triangleright M \leq_D^n \Gamma \triangleright M$$

Prove its **Soundness** with respect to the previous definition

$$\begin{aligned} &\Gamma_1 \triangleright M_1 \leq_D^n \Gamma_2 \triangleright M_2 \\ &\text{implies } \forall F_D^n(-) \\ &\Gamma_1 \triangleright M_1 \approx \Gamma_2 \triangleright F_D^n(M_2) \end{aligned}$$

Problems we need to address

Hard to prove positive results with our fault tolerance definition because:

1. \cong **quantifies over all valid contexts.**
2. Dynamic fault tolerance definition **quantifies over all fault contexts**, amongst which there is **considerable overlap.**
3. **There are a number of confluent reductions that increase the burden of our analysis.**

Confluent τ -transitions

Identify Confluent Moves

(b-eq)

$$\frac{}{\Gamma \triangleright l[[\text{if } u = u \text{ then } P \text{ else } Q]] \xrightarrow{\tau}_{\beta} \Gamma \triangleright l[[P]]}$$

(b-ngo)

$$\frac{}{\langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[\text{go } k.P]] \xrightarrow{\tau}_{\beta} \langle \mathcal{R}, \mathcal{U} \rangle \triangleright k[[\mathbf{0}]]} \quad k \notin \mathcal{R} \cup \mathcal{U}$$

Extend Equivalence Relation

((bs-dead))

$$\frac{}{\langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[P]] \equiv_f \langle \mathcal{R}, \mathcal{U} \rangle \triangleright l[[Q]]} \quad l \notin \mathcal{R} \cup \mathcal{U}$$

Confluent τ -transitions

Proving Confluence

$$\begin{array}{ccc} \Gamma \triangleright N & \xrightarrow{\tau} & \Gamma \triangleright M \\ \mu \downarrow & & \beta \\ \Gamma' \triangleright N' & & \end{array}$$

Confluent τ -transitions

Proving Confluence

$$\begin{array}{ccc} \Gamma \triangleright N & \xrightarrow[\beta]{\tau} & \Gamma \triangleright M \\ \mu \downarrow & & \mu \downarrow \\ \Gamma' \triangleright N' & \xrightarrow[\beta]{\tau} & \Gamma' \triangleright M' \end{array}$$

or

$$\begin{array}{ccc} \Gamma \triangleright N & \xrightarrow[\beta]{\tau} & \Gamma \triangleright M \\ \mu \downarrow & & \mu \downarrow \\ \Gamma' \triangleright N' & \equiv_f & \Gamma' \triangleright M' \end{array}$$

or $\mu = \tau$ and $\Gamma \triangleright M = \Gamma' \triangleright N'$

Confluent τ -transitions

Fault Tolerance up to β -moves

$\Gamma_1 \triangleright M_1 \leq_{\beta}^n \Gamma_2 \triangleright M_2$ implies

- $\Gamma_1 \triangleright M_1 \xrightarrow{\mu} \Gamma'_1 \triangleright M'_1$ implies $\Gamma_2 \triangleright M_2 \xRightarrow{\hat{\mu}} \Gamma'_2 \triangleright M'_2$ such that $\Gamma'_1 \triangleright M'_1 \mathcal{A}_l \circ \leq_{\beta}^n \circ \approx_{cnt} \Gamma'_2 \triangleright M'_2$
- $\Gamma_2 \triangleright M_2 \xrightarrow{\mu} \Gamma'_2 \triangleright M'_2$ implies $\Gamma_1 \triangleright M_1 \xRightarrow{\hat{\mu}} \Gamma'_1 \triangleright M'_1$ such that $\Gamma'_2 \triangleright M'_2 \mathcal{A}_l \circ \leq_{\beta}^n \circ \approx \Gamma'_1 \triangleright M'_1$
- If $n > 0$ then $\Gamma_2 \triangleright M_2 \xrightarrow{fail} \Gamma'_2 \triangleright M'_2$ implies $\Gamma_1 \triangleright M_1 \xRightarrow{\hat{\mu}} \Gamma'_1 \triangleright M'_1$ such that $\Gamma'_2 \triangleright M'_2 \leq_{\beta}^{n-1} \circ \approx \Gamma'_1 \triangleright M'_1$

where \mathcal{A}_l is the relation $\xRightarrow{\beta} \circ \equiv$

\approx_{cnt} is a bisimulation ranging over μ and the new counting action fail.

Confluent τ -transitions

Soundness of \leq_{β}^n

$$\Gamma_1 \triangleright M_1 \leq_{\beta}^n \Gamma_2 \triangleright M_2$$

implies


$$\Gamma_1 \triangleright M_1 \leq_D^n \Gamma_2 \triangleright M_2$$

Talk Summary

- Fault Tolerance Intuitions
- Language
- Formal Definition
- Proof Techniques

Main Result

***Induce
Faults***



***Observer
View***

Main Result

To show that $\Gamma \triangleright M$ is fault tolerant up to n faults we just have to give a witness fault tolerant simulation up to β -moves satisfying

$$\Gamma \triangleright M \leq_{\beta}^n \Gamma \triangleright M$$