

# CSc72010

## Proofs for asynchronous systems

### Examples

#### **Leader**

```
type Status = enumeration of leader, unknown

automaton Process (pos : Int, prev : Int, next : Int, rank: Int)
  signature
    input RECEIVE(m: Int , const prev, const pos : Int)
    output SEND (m: Int , const pos , const next : Int)
    output leader(const rank)

  states
%   pending : Mset [Int] := {rank},
    pending : Int := rank,
    status : Status := unknown,
    ready : Bool := true

  transitions
    input RECEIVE(m, j, i) where m > pending
      eff pending := m; ready := true

    input RECEIVE(m, j, i) where m < pending

    input RECEIVE(m, j, i) where m = rank
      eff status := leader

    output SEND (m, i, j)
      pre status ~= leader /\ m = pending /\ ready
      eff ready := false

    output leader(rank)
      pre status = leader
```

Representation of state: <4, T, unknown>, <7,F,leader>

That is, <pending, ready, status>

#### **Channel**

```
automaton Channel(i:Int , j:Int)
  signature
    input SEND (m: Int , const i : Int, const j : Int)
    output RECEIVE(m : Int, const i : Int ,const j : Int)

  states
    queue : Seq[Int] := {}

  transitions
    input SEND (m, i, j)
      eff queue := queue |- m

    output RECEIVE(m,i,j)
      pre head(queue) = m
      eff queue := tail(queue)
```

Representaton of state: {}, {4}, {4,7}

That is, *queue*

## Composition

automaton LCR  
components

```
P1: Process(1, 4, 2, 7);
P2: Process(2, 1, 3, 3);
P3: Process(3, 2, 4, 2);
P4: Process(4, 3, 1, 5);
C1: Channel(1, 2);
C2: Channel(2, 3);
C3: Channel(3, 4);
C4: Channel(4, 1)
```

## Proof Techniques

### Types of properties

#### Invariants

Properties of states that are true in all reachable states. Usually proof by induction. Step granularity is finer than rounds, so proofs are harder.

Here are invariants for leader:

$p.\text{pending} \geq p.\text{rank}$  for all processes  $p$

Another:

$c_{ij}.\text{queue}[k] \geq p_i.\text{rank}$  for all  $i, j, k$

Exercise – think about the inductive proof of these

How to state properties (section 14.1.2):

- 1) As an automaton – example, the Channel example that we are using can be used to specify the traces allowed by a reliable FIFO channel (so it is both a specification of such a channel and an example)
- 2) As axioms:  
There is a function *cause* mapping each receive to a send and satisfying
  - a. For every receive event  $\pi$ ,  $\pi$  and  $\text{cause}(\pi)$  contain the same message argument.
  - b. *cause* is surjective (onto).
  - c. *cause* is injective (one-to-one).
  - d. *cause* preserves order, that is, there do not exist receive events  $\pi_1$  and  $\pi_2$  with  $\pi_1$  preceding  $\pi_2$  in  $\beta$  and  $\text{cause}(\pi_2)$  preceding  $\text{cause}(\pi_1)$  in  $\beta$ .

Interpretation:

- a. Every message received was first sent
- b. Every send corresponds to a receive (reliable)
- c. Different receives correspond to different sends (no duplicates).
- d. This is the FIFO property.

## Trace properties

Any property of the external behavior sequences of automata. Formally, a trace property  $P$  is an external signature together with a set of sequences of actions in the signature, i.e., the allowable sequences:

$\langle \text{sig}(P), \text{traces}(P) \rangle$

An automaton  $A$  satisfies a trace property  $P$  if it has the same external signature and  $\text{traces}(A) \subseteq \text{traces}(P)$  or maybe  $\text{fairtraces}(A) \subseteq \text{traces}(P)$ .

All of the problems we will consider in asynchronous systems can be formulated as trace properties. Also, we'll usually be concerned with fairness, i.e., we'll be making statements about trace properties that hold for fair traces.

A trace property for leader:

In every trace, there is at most one leader(x) action:

Define the property by:

$$\begin{aligned} \text{sig}(P) &= \{ \text{SEND}(m,i,j), \text{RECEIVE}(m,i,j), \text{leader}(r) \mid m \in M, i,j \in \mathcal{N}, r \in \mathcal{N} \} \\ \text{traces}(P) &= \{ \tau_1, \tau_2, \dots \mid \forall i(\tau_i \in \text{sig}(P)) \wedge \neg \exists i, j, r, s (\tau_i = \text{leader}(r) \wedge \tau_j = \text{leader}(s)) \} \end{aligned}$$

## Safety properties

A safety property says that bad things don't happen.

In other words,  $\text{traces}(P)$  is non-empty, prefix-closed, limit-closed:

1.  $\lambda$  is in  $P$
2. if  $\alpha$  is in  $P$ , all prefixes of  $\alpha$  are in  $P$
3. if all prefixes of  $\alpha$  are in  $P$ , then  $\alpha$  is in  $P$

Examples:

At most one process declares itself as leader. (The traces are those including only one leader action.)

DHCP: No two processes get the same IP address

ARP: No two processes respond to the same ARP message

HTTP: No response to a GET is returned before the response to an earlier GET.

RIP (or any other routing protocol): No loops in routing tables. The path taken by any packet is the shortest path.

Mutual exclusion: No two simultaneous grants of resources

How to prove a safety property:

1. Relate it to a state invariant
2. Prove the state invariant

## Liveness properties

A liveness property says that good things do happen.  $\text{traces}(P)$  includes only those traces with the good things in them:

Every finite sequence over  $\text{sig}(P)$  has an extension in  $\text{traces}(P)$ .

## Examples

A leader is eventually elected

$$\text{traces(P)} = \{ \tau_1, \tau_2, \dots \mid \forall i (\tau_i \in \text{sig(P)}) \wedge \exists i, r (\tau_i = \text{leader}(r)) \}$$

A process eventually gets an IP address

A browser eventually gets a page

Note: you will have to assume fairness to get these properties.

What we can say is that “Every *trace* of a correct leader election algorithm contains at most one leader action” and “Every *fair trace* of a correct leader election algorithm contains a leader action.”

## Surprising fact

Every trace property can be expressed as a combination of a liveness property and a safety property. (e.g., a trace of a correct leader election algorithm contains exactly one leader action)

If  $\langle \text{sig(P)}, \text{traces(P)} \rangle$  is a trace property then  
there exists a safety property  $\langle \text{sig(P)}, \text{traces(S)} \rangle$  and  
there exists a liveness property  $\langle \text{sig(P)}, \text{traces(L)} \rangle$   
such that  $\text{traces(P)} = \text{traces(S)} \cap \text{traces(L)}$ .

So a good way to organize a specification is as a sequence of safety properties followed by a sequence of liveness properties.

Proof of theorem:

Let  $\text{traces(S)} =$  the prefix and limit closure of  $\text{traces(P)}$ , making it a safety property, and obviously  $\text{traces(S)}$  contains  $\text{traces(P)}$

Let  $\text{traces(L)} = \text{traces(P)} \cup \{ \beta \mid \beta \text{ is a finite sequence and no extension of } \beta \text{ is in } \text{traces(P)} \}$ .

We show that L is a liveness property:

Take any finite sequence  $\beta$  of actions in  $\text{sig(P)}$ . Either some extension of  $\beta$  is in  $\text{traces(P)}$  or not. If it is, then it is “live”; if not,  $\beta$  (an extension of itself) is in  $\text{traces(L)}$ . Hence L is a liveness property.

Clearly,  $\text{traces(P)} \subseteq \text{traces(S)} \cap \text{traces(L)}$ .

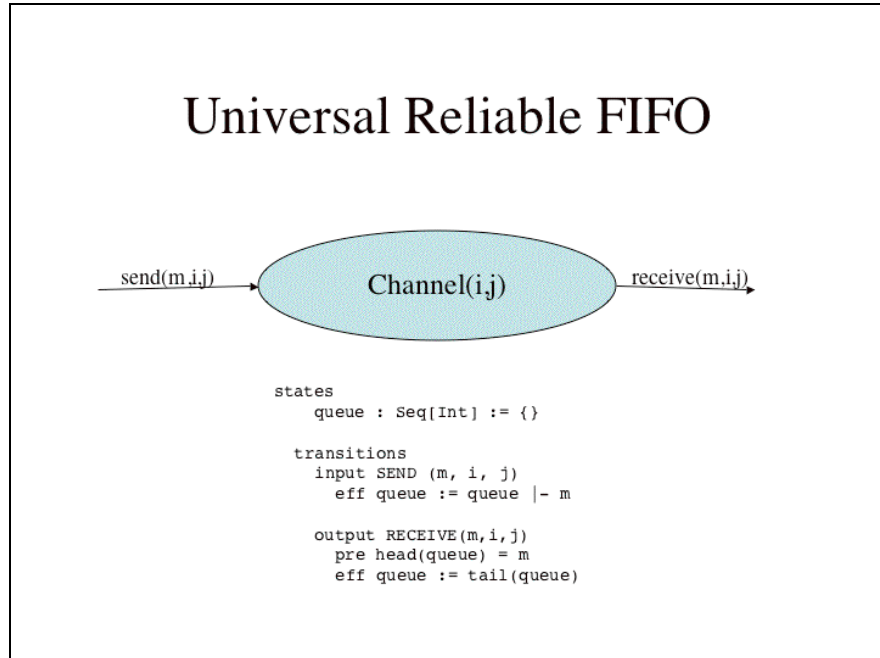
Let  $\beta \in (\text{traces(S)} \cap \text{traces(L)}) - \text{traces(P)}$ . This implies  $\beta \in \text{traces(L)} - \text{traces(P)}$ . So  $\beta$  must be a finite sequence with no extension in  $\text{traces(P)}$ . But also  $\beta \in \text{traces(S)}$ , and it is finite, so it must be a prefix of something in  $\text{traces(P)}$ . Contradiction.

## Hierarchical proofs

This is an important strategy for complex algorithms. We formulate the algorithm in a series of levels. For example, we could have a high-level centralized algorithm (easy to

prove, almost a specification), then do a simple but inefficient decentralized version, then do an optimized version.

High level version:



## Simulation Relations

Lower levels are harder to understand, so we relate them to higher levels with a simulation invariant rather than trying to deal with them directly. This is similar to the simulation relation for synchronous algorithms:

Run them side by side.

Define an invariant relating the states.

The invariant is called a *simulation relation* and is usually shown via induction.

Show that for each execution of the lower-level algorithm, there exists a related execution of the higher-level algorithm.

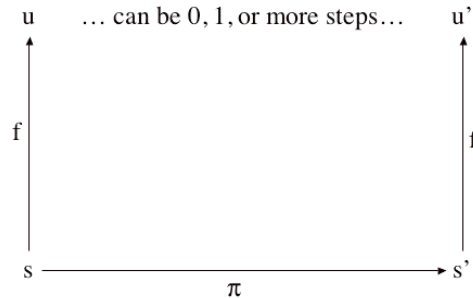
**Definition:** Assume A and B have the same external signature. Let f be a binary relation on  $\text{states}(A) \times \text{states}(B)$ .

Notation:  $(s,u) \in f$  or  $u \in f(s)$

Then f is a simulation relation from A to B provided that

1. If  $s \in \text{start}(A)$  then  $f(s) \cap \text{start}(B) \neq \emptyset$
2. If s,u are reachable states of A and B respectively, with  $u \in f(s)$ , and if  $(s, \pi, s')$  is a transition of A, then there is an execution fragment  $\alpha$  of B starting with u, and ending with  $u' \in f(s')$ , with  $\text{trace}(\alpha) = \text{trace}(\pi)$ .

# Simulation Relation



**Theorem.** If there's a simulation relation from A to B then  $\text{traces}(A) \subseteq \text{traces}(B)$ .

Proof: Take any execution of A, and iteratively construct the corresponding execution of B.

## Example proof

Example: Implementing the reliable FIFO channel with TCP (sort of)

## Altered spec

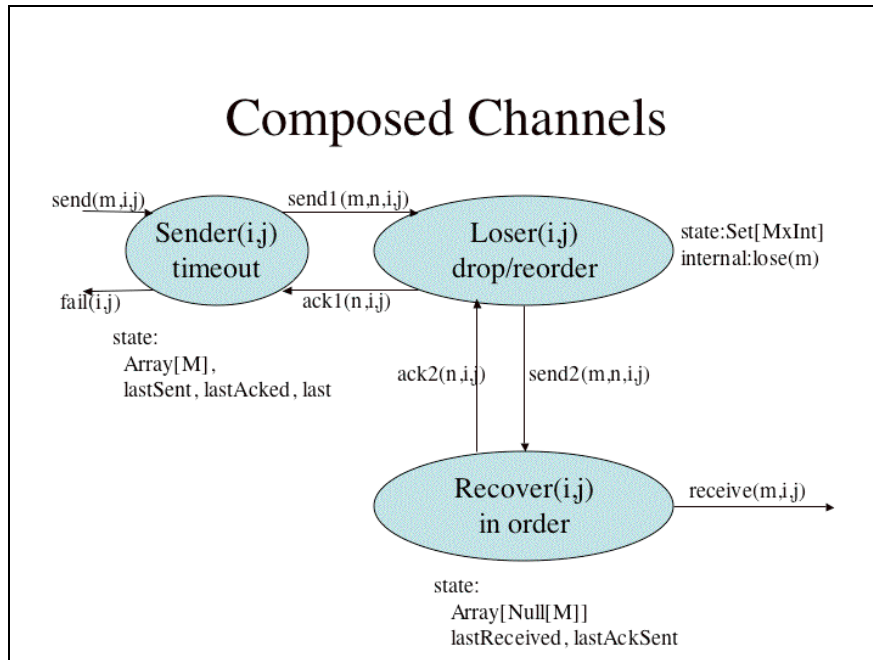
Add fail(i,j) transition to Reliable FIFO: all messages up to some point get sent in order. At some point there is a fail. No more messages can be sent after some point, but a subset of the messages in the queue can be received.

either there is a fail action or the cause is surjective.

if a send  $\text{cause}(\pi)$  is in  $\text{range}(\text{cause})$  then so is every earlier send

## Implementation

This is an adaptation of TCP.



### Discussion:

$\text{Sender}(i, j)$  puts arriving messages in  $\text{Sender}(i, j).\text{Array}[\text{last}+1]$  and increments  $\text{last}$ . It sends messages in  $\text{Sender}(i, j).\text{Array}[\text{lastAcked}+1 \dots \text{last}]$  repeatedly until it receives  $\text{ack1}(n, i, j)$ . Then it sets  $\text{lastAcked}$  to  $\max(\text{lastAcked}, n)$ .

$\text{Loser}(i, j)$  puts messages in  $\text{Loser}(i, j).\text{Set}$  when  $\text{send1}(m, i, j)$  happens. The precondition for  $\text{send2}(m, i, j)$  and for  $\text{delete}(m, i, j)$  is that  $m \in \text{Loser}(i, j).\text{Set}$ . The effect of  $\text{delete}(m, i, j)$  is to remove the message from the set, but  $\text{send2}(m, i, j)$  causes no change. This means that  $\text{Loser}(i, j)$  delivers each message to  $\text{Recover}(i, j)$  0 or more times.

$\text{Recover}(i, j)$  puts each message in its proper position in  $\text{Recover}(i, j).\text{Array}$ .  $\text{Recover}(i, j).\text{ack2}(n, i, j)$  is enabled when there are no holes in  $\text{Recover}(i, j).\text{Array}[1..n]$  and  $n > \text{lastAckSent}$ . Its effect is to update  $\text{lastAckSent}$  to  $n$ .

$\text{Recover}(i, j).\text{receive}(m, i, j)$  is enabled if  $m = \text{Recover}(i, j).\text{Array}[\text{lastReceived}+1]$  and  $\text{lastAckSent} > \text{lastReceived}$ .

### Simulation relation

We are comparing the composition of  $\text{Sender}$ ,  $\text{Loser}$ , and  $\text{Recover}$  with all actions hidden except  $\text{send}(m, i, j)$ ,  $\text{receive}(m, i, j)$ , and  $\text{fail}(i, j)$ . We need to show that the traces of this composition are contained in the traces of  $\text{Channel}(i, j)$ , modified by the fail action.

The important parts of the state are:

- $\text{Sender}(i, j).\text{Array}$
- $\text{Recover}(i, j).\text{Array}$
- $\text{lastReceived}$

lastAked

Facts:

$\text{Recover}(i, j).Array \subseteq \text{Sender}(i, j).Array[1..lastSent]$  (with holes)

$\text{Recover}(i, j).Array[1..lastReceived] \supseteq \text{Sender}(i, j).Array[1..lastAked]$

$\text{Recover}(i, j).Array[1..lastAckSent]$  is the messages already delivered plus those that are ready to be delivered

The correspondence is that  $\text{Channel}(i, j).Queue$  corresponds to any state with all messages in  $\text{Channel}(i, j).Queue$  sent and no messages in  $\text{Channel}(i, j).Queue$  received (yet) – that is,

$\text{Channel}(i, j).Queue$  corresponds to all states such that there exists  $n$  with  
 $\text{Send}(i, j).Array[n..last] = \text{Channel}(i, j).Queue$  and  
 $n = \text{Recover}(i, j).lastReceived + 1$

We must show that the simulation relation is preserved by all actions. Consider actions of the composition:

`send` adds a message to  $\text{Channel}(i, j).Queue$  and to  $\text{Sender}(i, j).Array[last+1]$  – this preserves the simulation relation

`send1`, `send2` don't change any of the relevant parts of the state

`receive` increases  $\text{Recover}(i, j).lastReceived$  and  $\text{Recover}(i, j).lastAckSent$ ; also removes an element from  $\text{Channel}(i, j).Queue$  – this preserves the simulation relation

`ack2` and `ack1` have no effect on relevant parts of state

`fail` can happen any time but after `fail`, `sender` quits doing anything - the only other thing that can happen is that messages get through the system and received. In `channel`, no more messages can be sent but some subset of the messages already in the channel can be delivered.