

Software Engineering using Formal Methods

Modeling Distributed Systems

Wolfgang Ahrendt

16 September 2014

This Lecture

You know you have a distributed system when the crash of a computer you've never heard of stops you from getting any work done. –Leslie Lamport

Using PROMELA channels for modeling distributed systems

Modeling Distributed Systems

Distributed systems consist of

- ▶ **nodes**
- ▶ interacting via **communication channels**
- ▶ **protocols** dictate how nodes communicate with each other

Distributed systems are very complex

Models of distributed systems abstract away from details of networks/protocols/nodes

In PROMELA:

- ▶ **nodes** modeled by **PROMELA processes**
- ▶ **communication channels** modeled by **PROMELA channels**
- ▶ protocols modeled by algorithm distributed over processes

Channels in PROMELA

In PROMELA, channels are first class citizens

Data type `chan` with two operations for **sending** and **receiving**

A variable of channel type is declared by initializer:

```
chan name = [capacity] of {type1, ..., typen}
```

name name of channel variable

capacity non-negative integer constant

*type*_{*i*} PROMELA data types

Example:

```
chan ch = [2] of { mtype, byte, bool }
```

Meaning of Channels

chan *name* = [*capacity*] of {*type*₁, ..., *type*_{*n*}}

Creates channel, stored in *name*

Messages communicated via channel are *n*-tuples $\in \textit{type}_1 \times \dots \times \textit{type}_n$

Can buffer up to *capacity* messages, if *capacity* ≥ 1

\Rightarrow "buffered channel"

The channel has *no* buffer if *capacity* = 0

\Rightarrow "rendezvous channel"

Meaning of Channels

Example:

```
chan ch = [2] of { mtype, byte, bool }
```

Creates channel, stored in `ch`

Messages communicated via `ch` are 3-tuples $\in \text{mtype} \times \text{byte} \times \text{bool}$

Given, e.g., `mtype = {red, yellow, green}`,
an example message on `ch` can be: `green, 20, false`

`ch` is a *buffered channel*, buffering up to 2 messages

Sending and Receiving

send statement has the form:

name ! expr₁, ... , expr_n

- ▶ *name*: channel variable
- ▶ *expr₁, ... , expr_n*: sequence of expressions, where number and types match message type
- ▶ sends *values* of *expr₁, ... , expr_n* as *one* message
- ▶ example: `ch ! green, 20, false`

receive statement has the form:

name ? var₁, ... , var_n

- ▶ *name*: channel variable
- ▶ *var₁, ... , var_n*: sequence of variables, where number and types match message type
- ▶ *assigns* values of message to *var₁, ... , var_n*
- ▶ example: `ch ? color, time, flash`

Client-Server

```
chan request = [0] of { byte };
```

```
active proctype Client0() {  
  request ! 0  
}
```

```
active proctype Client1() {  
  request ! 1  
}
```

...

Client0 and Client1 send messages 0 and 1 to request
order of sending is nondeterministic

Client-Server

```
chan request = [0] of { byte };  
  
...  
  
active proctype Server() {  
    byte num;  
    do  
        :: request ? num;  
        printf("serving client %d\n", num)  
    od  
}
```

Server loops on:

- ▶ receiving first message from request, storing value in num
- ▶ printing

Executability of receive Statement (non-buffered)

```
request ? num
```

executable only when another process offers send on channel request

⇒ receive statement frequently used as guard in `if/do`-statements

```
do
  :: request ? num ->
    printf("serving client %d\n", num)
od
```

("->" equivalent to ";", but highlights guard role of request ? num)

Rendezvous Channels

```
chan ch = [0] of { byte, byte };

/* global to make visible in SpinSpider */
byte hour, minute;

active proctype Sender() {
    printf("ready\n");
    ch ! 11, 45;
    printf("Sent\n")
}

active proctype Receiver() {
    printf("steady\n");
    ch ? hour, minute;
    printf("Received\n")
}
```

Which interleavings can occur? \Rightarrow ask SPINSPIDER

Rendezvous are Synchronous

On a rendezvous channel:

transfer of message from sender to receiver is **synchronous**,
i.e., **one single operation**

Sender		Receiver
⋮		⋮
(11,45)	→	(hour,minute)
⋮		⋮

Rendezvous are Synchronous

Either:

1. Location counter of sender process at send ("!"): *"offer to engage in rendezvous"*
2. Location counter of receiver process at receive ("?"): *"rendezvous can be accepted"*

or the other way round:

1. Location counter of receiver process at receive ("?"): *"offer to engage in rendezvous"*
2. Location counter of sender process at send ("!"): *"rendezvous can be accepted"*

in any cases:

location counter of **both** processes is incremented at once

only place where PROMELA processes execute synchronously

Reconsider Client Server

```
chan request = [0] of { byte };

active proctype Server() {
  byte num;
  do :: request ? num ->
    printf("serving client %d\n", num)
  od
}

active proctype Client0() {
  request ! 0
}

active proctype Client1() {
  request ! 1
}
```

so far **no reply** to clients

Reply Channels

```
chan request = [0] of { byte };
chan reply = [0] of { bool };

active proctype Server() {
  byte num;
  do :: request ? num ->
    printf("serving client %d\n", num);
    reply ! true
  od
}

active proctype Client0() {
  request ! 0;  reply ? _
}

active proctype Client1() {
  request ! 1;  reply ? _
}
```

(anonymous variable “_” used if interested in receipt, not content)

Reply Channels - Single Server

```
mtype = { nice, rude };
chan request = [0] of { mtype };
chan reply = [0] of { mtype };

active proctype Server() {
    mtype msg;
    do :: request ? msg; reply ! msg
    od
}

active proctype NiceClient() {
    mtype msg;
    request ! nice; reply ? msg;
    assert(msg == nice)
}

active proctype RudeClient() {
    mtype msg;
    request ! rude; reply ? msg
}
```

Is the assertion valid? Ask SPIN.

Several Servers

More realistic with several servers:

```
active [2] proctype Server() {
  mtype msg;
  do :: request ? msg; reply ! msg
  od
}
active proctype NiceClient() {
  mtype msg;
  request ! nice; reply ? msg;
  assert(msg == nice)
}
active proctype RudeClient() {
  mtype msg;
  request ! rude; reply ? msg
}
```

And here? Analyse with SPIN.

Sending Channels via Channels

To fix the protocol:

clients declare local reply channel + send it to server

Sending Channels via Channels

```
mtype = { nice, rude };
chan request = [0] of { mtype, chan };

active [2] proctype Server() {
  mtype msg; chan ch;
  do :: request ? msg, ch;
    ch ! msg
  od
}

active proctype NiceClient() {
  chan reply = [0] of { mtype }; mtype msg;
  request ! nice, reply; reply ? msg;
  assert( msg == nice )
}

active proctype RudeClient() {
  chan reply = [0] of { mtype }; mtype msg;
  request ! rude, reply; reply ? msg
}
```

verify with SPIN

Scope of Channels

channels are typically declared global

global channel

- ▶ usual case
- ▶ all processes can send and/or receive messages

local channel

- ▶ rarely used
- ▶ dies with its process
- ▶ can be useful to model security issues

example:

local channel could be passed
through a global channel

Sending Process IDs

used *fixed constants* used for identification (here nice, rude)

- ▶ inflexible
- ▶ doesn't scale

Alternative:

processes send their own, unique **process ID**, `_pid`, as part of message

experiment with rendezvous3.pml

example, clients code:

```
chan reply = [0] of { byte, byte };
request ! _pid, reply;
reply ? serverID, clientID;

assert( clientID == _pid )
```

Limitations of Rendezvous Channels

- ▶ rendezvous too restrictive for many applications
- ▶ servers and clients block each other too much
- ▶ difficult to manage uneven workload
(online shop: dozens of webservers serve thousands of clients)

Buffered Channel

buffered channels queue messages;
requests/services no not immediately block clients/servers

example:

```
chan ch = [3] of { mtype, byte, bool }
```

Buffered Channels

buffered channels, with capacity *cap*

- ▶ can hold up to *cap* messages
- ▶ are a FIFO (first-in-first-out) data structure:
always the 'oldest' message in channel is retrieved by a receive
- ▶ (normal) receive statement reads **and** removes message
- ▶ Sending and Receiving to/from buffered channels is asynchronous, i.e. interleaved

Executability of Buffered Channel operations

given channel ch , with capacity cap , currently containing n messages

receive statement $ch ? msg$

is executable iff ch is not empty, i.e., $n > 0$

send statement $ch ! msg$

is executable iff there is still 'space' in the message queue,
i.e., $n < cap$

An non-executable receive or send statement will **block** until it is executable again

(With option `-m`, SPIN has a different send semantics: attempt to send to a full channel does not block, but the message gets lost.)

Checking Channel for Full/Empty

this can save from unnecessary blocking:

given channel `ch`:

`full(ch)` checks whether `ch` is full

`nfull(ch)` checks whether `ch` is not full

`empty(ch)` checks whether `ch` is empty

`nempty(ch)` checks whether `ch` is not empty

illegal to negate those

avoid combining with `else`

Copy Message without Removing

Assume `ch` to be a buffered channel.

`ch ? color, time, flash`

- ▶ assigns values from the message to `color`, `time`, `flash`
- ▶ removes message from `ch`

`ch ? <color, time, flash>`

- ▶ assign values from the message to `color`, `time`, `flash`
- ▶ leaves message in `ch`

Dispatching Messages

Recurring task: Dispatch action depending on message

```
mtype = {hi, bye};
chan ch = [0] of {mtype};

active proctype Server () {
    mtype msg;
read:
    ch ? msg;
    do
        :: msg == hi -> printf("Hello.\n"); goto read
        :: msg == bye -> printf("See you.\n"); break
    od
}
...
```

There is a better way!

Pattern Matching

Receive statement allows also values as arguments:

$$ch ? exp_1, \dots, exp_n$$

- ▶ exp_1, \dots, exp_n any(!) expressions of correct type
- ▶ statement is **executable**, iff message msg_1, \dots, msg_n in channel ch **matches** arguments, i.e. if
 - ▶ exp_i is a variable, then any value of msg_i (of correct type) matches and is assigned if statement is executed
 - ▶ exp_i is a value, e.g. 23, msg_i must have same value

Pattern Matching Examples

Assume

```
chan ch = [0] of {int, int};  
int id = 5;
```

Does `ch ? 0, id` match message

- ▶ `[0, 5] ?` ✓ `[0, 7] ?` ✓ `[1, 7] ?` ✗
- ▶ Value of `id` afterwards?

To match the **value** stored in a variable `var` use `eval(var)`

Does `ch ? 0, eval(id)` match message

- ▶ `[0, 5] ?` ✓ `[0, 7] ?` ✗ `[1, 7] ?` ✗
- ▶ Value of `id` afterwards?

Dispatching Messages Revisited

Recurring task: Dispatch action depending on message type.

```
mtype = {hi, bye};  
chan ch = [0] of {mtype};  
  
active proctype Server () {  
    int i;  
    do  
        :: ch ? hi    -> printf("Hello.\n")  
        :: ch ? bye  -> printf("See you.\n"); break  
    od  
}  
...
```

Dispatching Messages Revisited

Random receive ?? (for buffered channels)

- ▶ Executable if matching message exists in channel.
- ▶ If executed, **first** matching message removed from channel.

```
mtype = {hi, bye};  
chan ch = [3] of {mtype};  
  
active proctype Server () {  
    int i;  
    do  
        :: ch ?? bye -> printf("See you.\n"); break  
        :: else      -> printf("Hello.\n")  
    od  
}  
...
```


Nicer Message Formatting

PROMELA provides an alternative, but equivalent syntax for

```
ch ! exp1, exp2, exp3
```

namely

```
ch ! exp1(exp2, exp3)
```

Increases readability for certain applications, e.g. protocol modelling:

```
ch!send(msg,id) vs. ch!send,msg,id
```

```
ch!ack(id) vs. ch!ack,id
```

And finally

Buffered channels are part of the state!

State space gets much bigger using buffered channels

Use with care (and with small buffers).