A Kernel-based Communication Fault Injector for Dependability Testing of Distributed Systems

Roberto Jung Drebes, Gabriela Jacques-Silva,
Joana Fonseca da Trindade, Taisy Silva Weber

Universidade Federal do Rio Grande do Sul, Brazil

drebes@inf.ufrgs.br

PADTAD 3 - Fault Injection-based Testing Session

Motivation

- Specification and project errors, and inevitable hardware faults can lead a system to have fatal consequences.
- Developers guarantee dependable behavior using techniques of fault tolerance.
- Fault tolerance implementation should be validated:
  - The recovery mechanisms should mask faults;
  - Unmasked faults should take the system through a known, expected, fault process.

Motivation (cont.)

- Fault Injection
  - goal: experimental validation
  - definition:
    the controlled introduction of faults from a given scenario into a test system to observe how it behaves under the presence of faults.
    (speed up of hard to reach situations)

ComFIRM

Communication Fault Injection through operating system Resources Modification

- explores the extensibility of the Linux kernel through modules, minimizing the probe effect (interference) in the kernel itself.
- aims to inject communication faults to validate fault tolerance aspects of network protocols and distributed systems.

ComFIRM (cont.)

- uses the Netfilter architecture and other high-level features. It is architecture independent and portable to any device running recent versions of Linux: servers, workstations, embedded devices.
- accepts configuration of the experiment during runtime. The description of fault scenarios and fault activation is done through simple rules written to virtual files in the proc filesystem.

ComFIRM (cont.)

- Messages can be selected by content, flow or counters and flags.
- Manipulation defines actions that change some message field, drop, delay (late delivery) or duplicate messages.
- There are also actions that manipulate counters, timers, variables and user level warnings.

ComFIRM (cont.)

- Selection and manipulation instructions are combined in rules, that are evaluated for every incoming and outgoing packet. A ruleset is a set of all the individual rules.

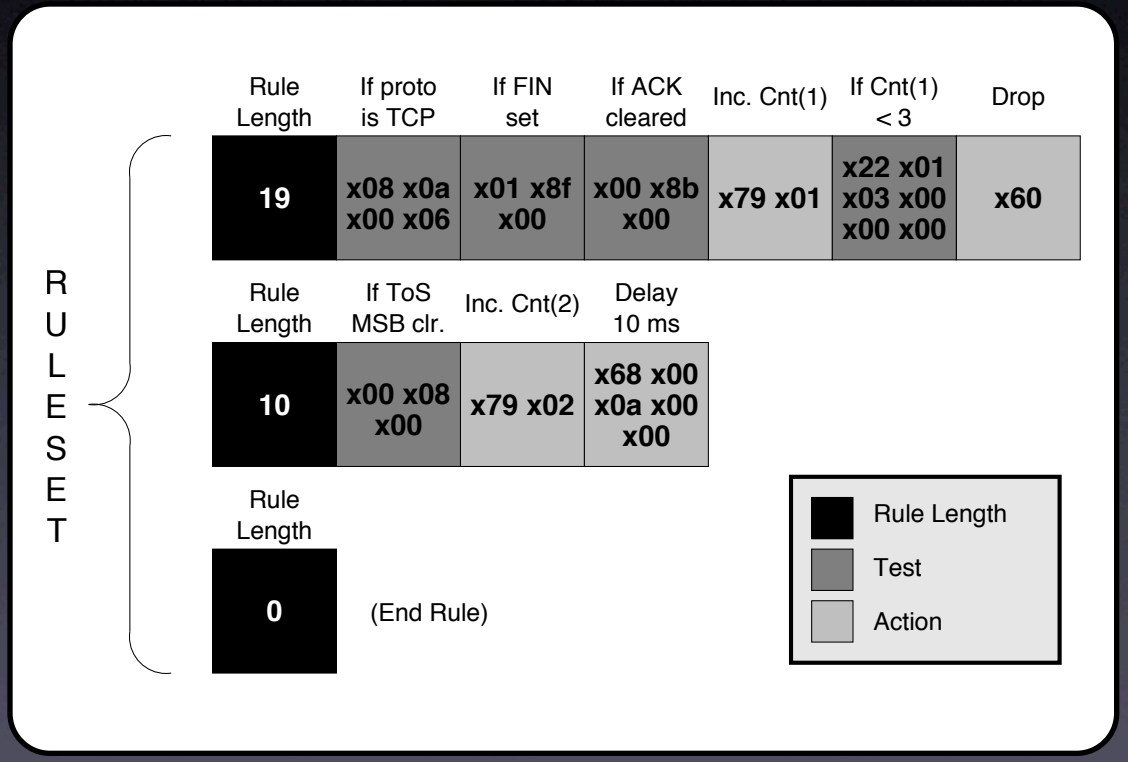

Demonstration 1

- Testbed: JGroups, an open source Java toolkit.
  - provides reliable multicast communication for distributed systems and applications;
  - allows its protocol stack to be adapted to different development requirements and/or network characteristics, through the inclusion of specific protocol layers;
  - provides a group membership service (GMS).

Demonstration I (cont.)

- JGroups protocol stack
  - multicast at its base (UDP);
  - a loss-less transmission layer, through the addition of a NAKACK (negative acknowledgement) protocol
  - a group membership service (GMS protocol) to provide membership change notifications (i.e. node join, leave or crash detection)

Demonstration I (cont.)

- The GMS layer was tested:
  - without a fault detection mechanism,
  - with a fault detection mechanism based on a heartbeat protocol (FD) and
  - with a fault detection mechanism based on sockets (FD SOCK).

Demonstration I (cont.)

- JGroups application: joins a group and notifies members when the membership changes.
- ComFIRM causes loss of conectivity.
- Aim is not to provide a complete example of a dependability evaluation, but to exemplify how ComFIRM can be applied to such kind of evaluation.
- Link crash can be emulated through many ways. ComFIRM's instructions allow the selection of specific messages, which helps in a more precise setting of the experiments.

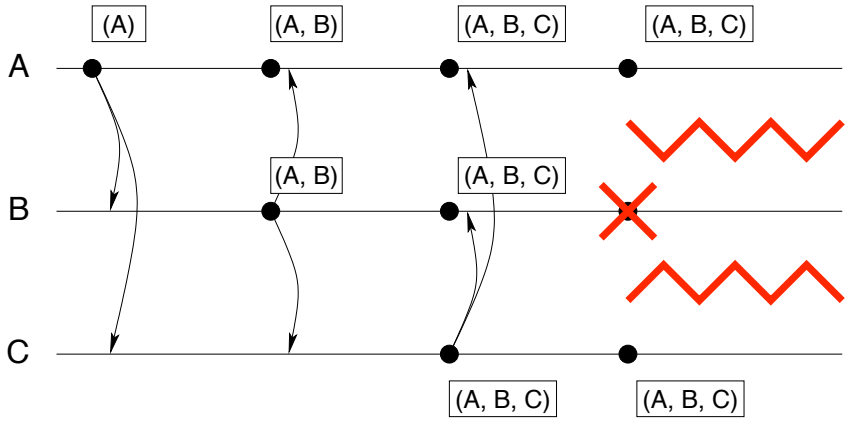
(A)
(A, B)
(A, B, C)
(A, B, C)
A
B
C
(A, B)
(A, B, C)
(A, B, C)
(A, B, C)

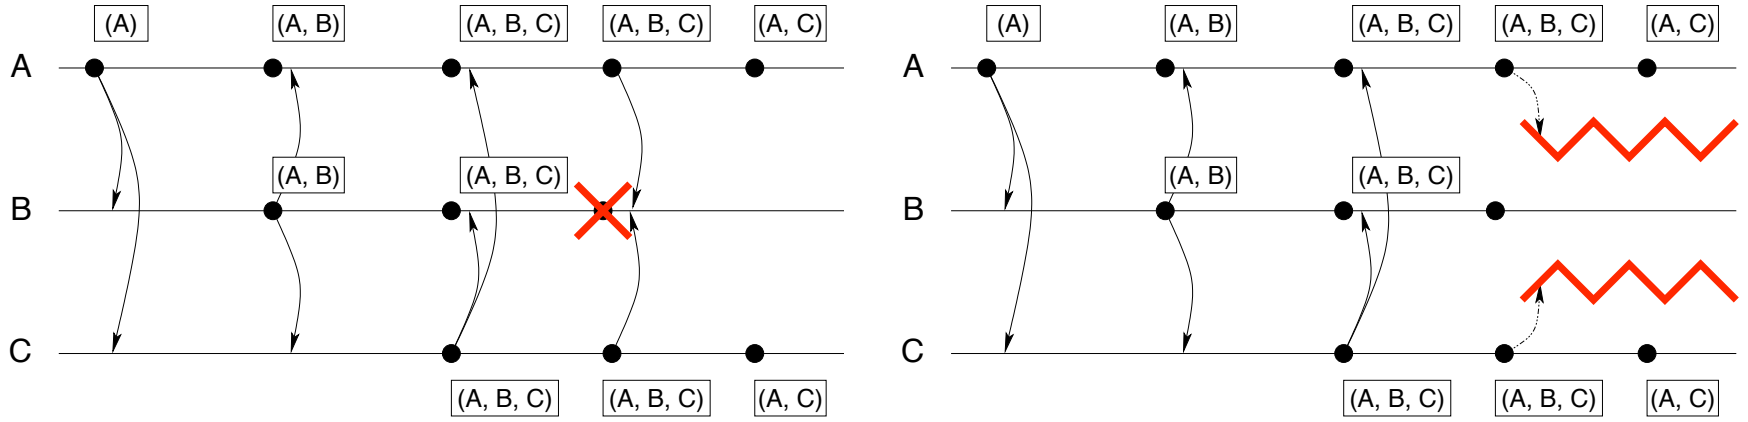
(A)
(A, B)
(A, B, C)
(A, B, C)
(A, C)
A
B
C
(A, B)
(A, B, C)
(A, B, C)
(A, B, C)
(A, C)
(A)
(A, B)
(A, B, C)
(A, B, C)
(A, C)
A
B
C
(A, B)
(A, B, C)
(A, B, C)
(A, B, C)
(A, C)

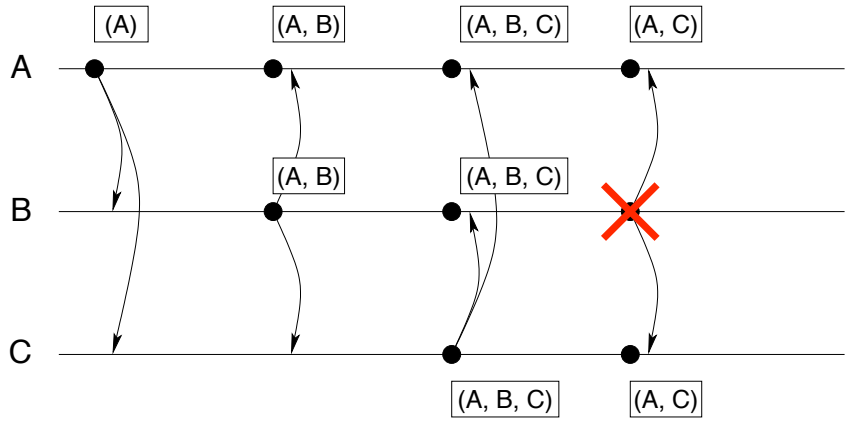
(A)
(A, B)
(A, B, C)
(A, C)
A
B
C
(A, B)
(A, B, C)
(A, B, C)
(A, C)

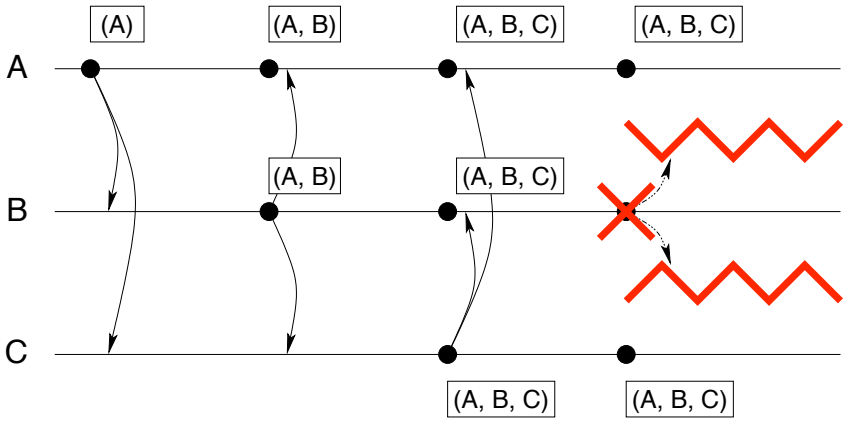
(A)
(A, B)
(A, B, C)
(A, B, C)
A
B
C
(A, B)
(A, B, C)
(A, B, C)
(A, B, C)

Demonstration 2

- Testbed: a standard MPEG4 streaming video session, which uses RTP over UDP.
  (Darwin Streaming Server 5.0.1, QuickTime Player 6.5.2)
- Server: Linux host running ComFIRM
  Client: remote machine, connected through a switched 100 Mbps Ethernet.
- Fault scenario: statistically drop 10% of messages whose transport level protocol is UDP and the application level protocol is RTP.

Demonstration 2 (cont.)

- Average video bitrate: 962 kbps
- We captured statistical data about the stream, starting 10 minutes into the session to eliminate transient effects like buffering.
- The bar graphs represent the state of the transmission during the latest 120 second interval.

Demonstration 2 (cont.)

- While packet loss is not homogeneous, its average value is close to the configured 10% rate. This is expected, since we are using a statistical loss rate.
- Qualitative results can be seen by the large amount of video artifacts observed during playback.

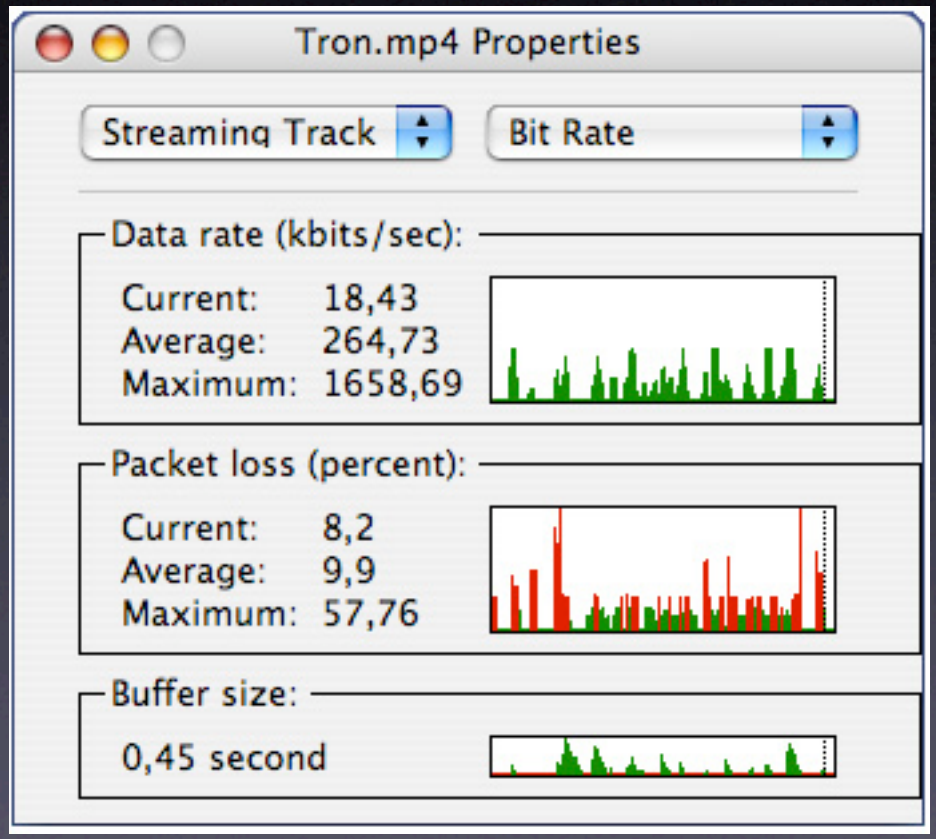

Demonstration 2 (cont.)

- The server/client pair can mask faults only when the buffer is filled. This is appropriate to burst faults, but not to intermittent faults.
- In this case, faults are manifested even if the available bitrate after faults is larger than the average video bitrate.
- This happens because retransmissions are never made, even if there are still data in the buffer.

Final Remarks

- Fault injection is a powerful technique to evaluate protocol and application behaviour under faults, measuring the efficiency of detection, correction and error recovery mechanisms of a system before it is put into effective operation.
- ComFIRM is a FI tool which uses the Netfilter framework. This gives the tool full access to the incoming and outgoing message flows in a clean and non-intrusive way.

Final Remarks (cont.)

- ComFIRM instructions allow messages to be inspected and selected in a deterministic or statistical way and provide actions which mimic the behavior of real faults.
- The tool is fully operational.

Current and Future work

- ComFIRM started being adapted to work with IPv6. New instructions were necessary so that the flow of the fault rule could be based on the message contents.
- Another tool, FIRMAMENT, is being developed, which employs the concept of faultlets, fault decision programs which are executed under a virtual machine located in the communication subsystem.

Current and Future work (cont.)

- Regarding the coordination of multi-node experiments, ComFIRM (and FIRMAMENT) are being integrated to the FIONA framework, for centralized configuration of experiments.
- This is performed through the translation of fault scenarios specifications, made by specific backends responsible for this tranlsation and communications on the nodes.

ComFIRM

communication fault injection
through operating system resource modification

Roberto Jung Drebes drebes@inf.ufrgs.br

Gabriela Jacques-Silva gjsilva@inf.ufrgs.br

Joana Trindade jmftrindade@inf.ufrgs.br

Taisy Silva Weber taisy@inf.ufrgs.br