Measurement tool of one-way packet loss rates

Masato TSURU

TAO (Telecommunications Advancement Organization)

1. What to be measured (inferred)
2. Methodology
3. A stand-alone tool
4. A server-client style tool
5. Experiments on the Internet
6. Concluding remarks
7. Demo

APAN 2003 Conference in Fukuoka

What to be measured

- One-way packet loss rates
  - On a path segment (a portion of a path) from/to a user-host (a client) to/from a specified target-host (an app. server or a router),
  - Without any measurement on the target,
  - Where the “loss rate” is the probability of a packet being dropped on the path segment.
- To find the congested area along the end-to-end path.
- We developed (prototypes of) two types of tools:
  - A stand-alone tool running on the client.
  - A client-server style tool running on both the client and a proxy measurement server distributed on the Internet.

Methodology

How to infer O-W less rates on a path segment:

1. A trial: sending a very closely-spaced packet pair (P_a, P_b) along tree-structured paths P_a: $0 \to 3 \to 1$, and P_b: $0 \to 3 \to 2$
2. Dispatch a number of trials independently
3. Observe the arrival of each packet at node 1 and 2: (N, E_a, E_b, E_{ab})

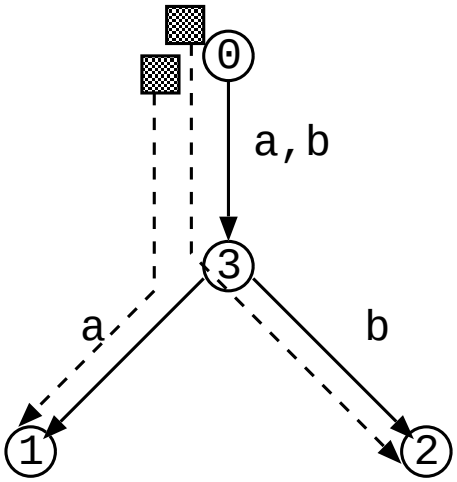

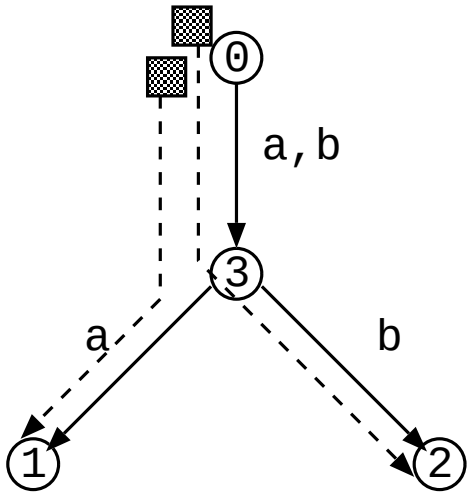

[Assumptions]

A1 Packets are not always dropped on each link.

A2 Packets are dropped on each link independently.

A3 Given the second P_b reaches node 3, the first P_a is very likely to reach node 3.

- A typical situation:
  - Losses on a link occur by queue overflows,
  - The overflow is caused by many independent, diverse traffic across the link,
  - The queue is managed as FIFO,
- Where “link” means a path segment.

"Loss Rates" on link $\underline{ab}$ and link $\underline{a}$ can be inferred:

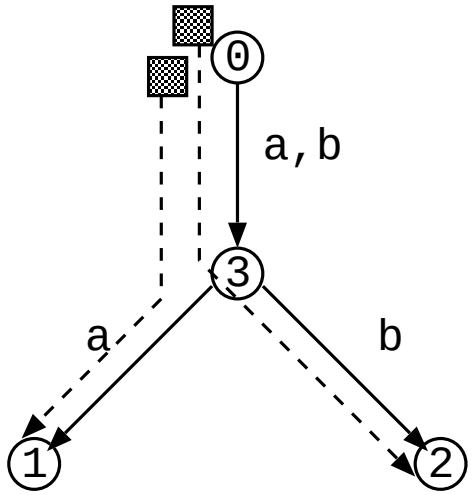

[Loss rates on link a]

$$\approx 1 - \frac{E_{ab}}{E_b}$$

[Loss rates on link ab]

$$\approx 1 - \frac{E_a}{N} \times \frac{E_b}{E_{ab}}$$

- If we see P_b reaches end-node 2, we also know P_a is likely to reach intermediate node 3.
- Thus, the conditional prob. of P_a reaching end-node 1 (given that P_b reaches node 2) approximates to the "no-loss" rate of link a.
- "No-loss rate" of link ab can be estimated by (no-loss rate of path a) $\div$ (no-loss rate of link a).

A stand-alone tool

- Running on client C (a user-host),
- Sending a closely-spaced packet pair of P_a and P_b traveling along the following paths.
- Inferring one-way loss rates on a path segment from/to a user-host C to/from a target host S.

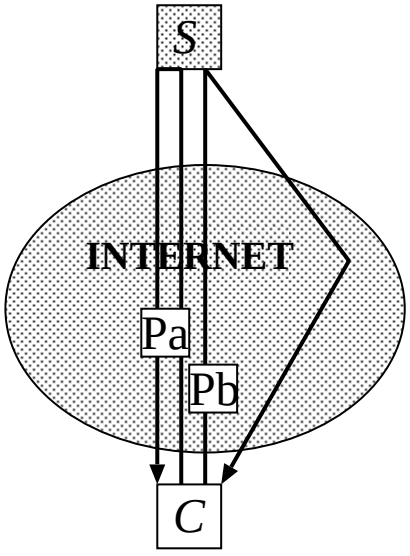

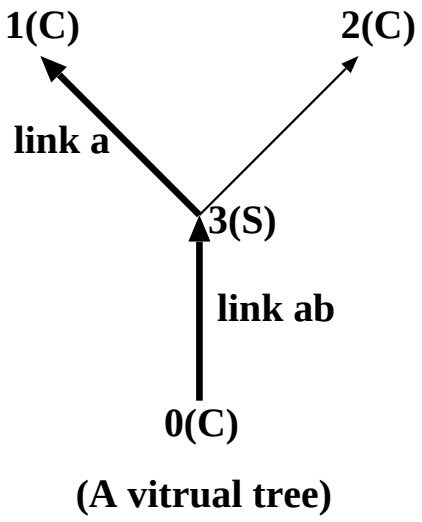

- UDP packet and its ICMP reply (e.g., Time-exceed) — optionally with LSRR

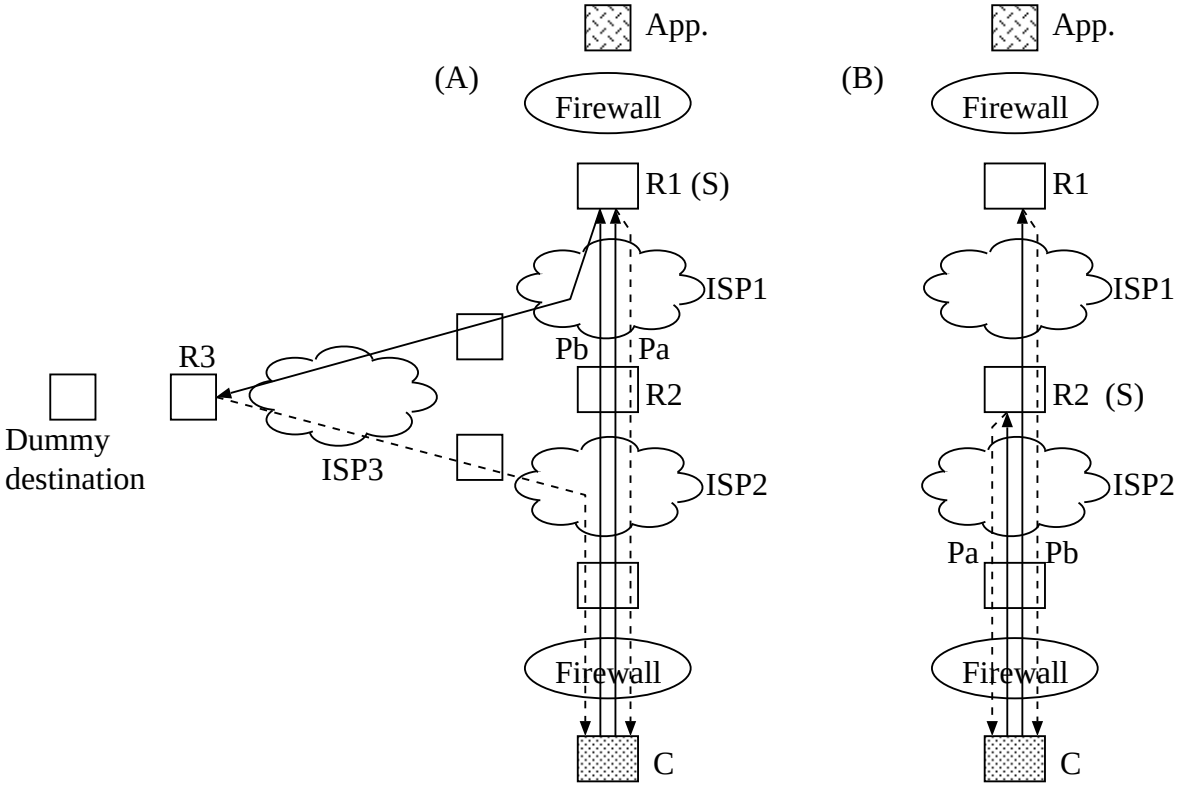

(A) Loss rates from/to C to/from a boundary router R_1 with use of LSRR.

(B) Loss rates from/to C to/from an intermediate router R_2 WITHOUT use of LSRR.

A client-server tool

- Running on both client C and a proxy measurement server M, assumed to be distributed in the Internet,
- Sending a closely-spaced packet pair of P_a and P_b traveling along the following paths.
- Inferring one-way loss rates on a path segment from a target host S to a user-host C.

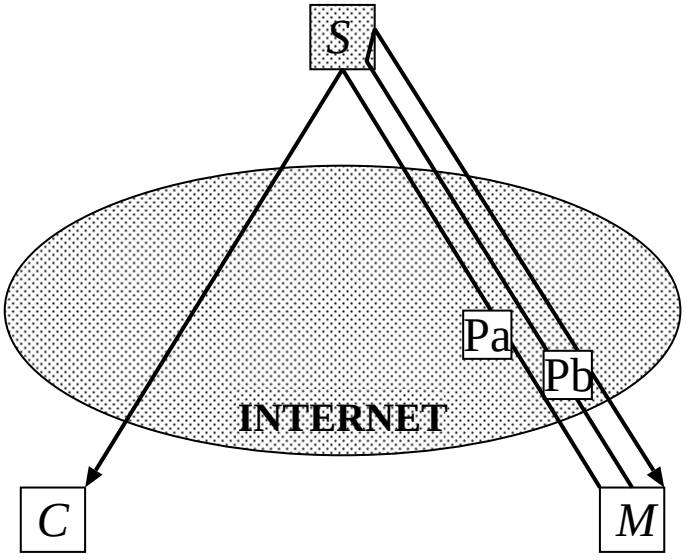

- UDP packet and its ICMP reply (e.g., Time-exceed)
- A forged (modified) IP source address in P_a

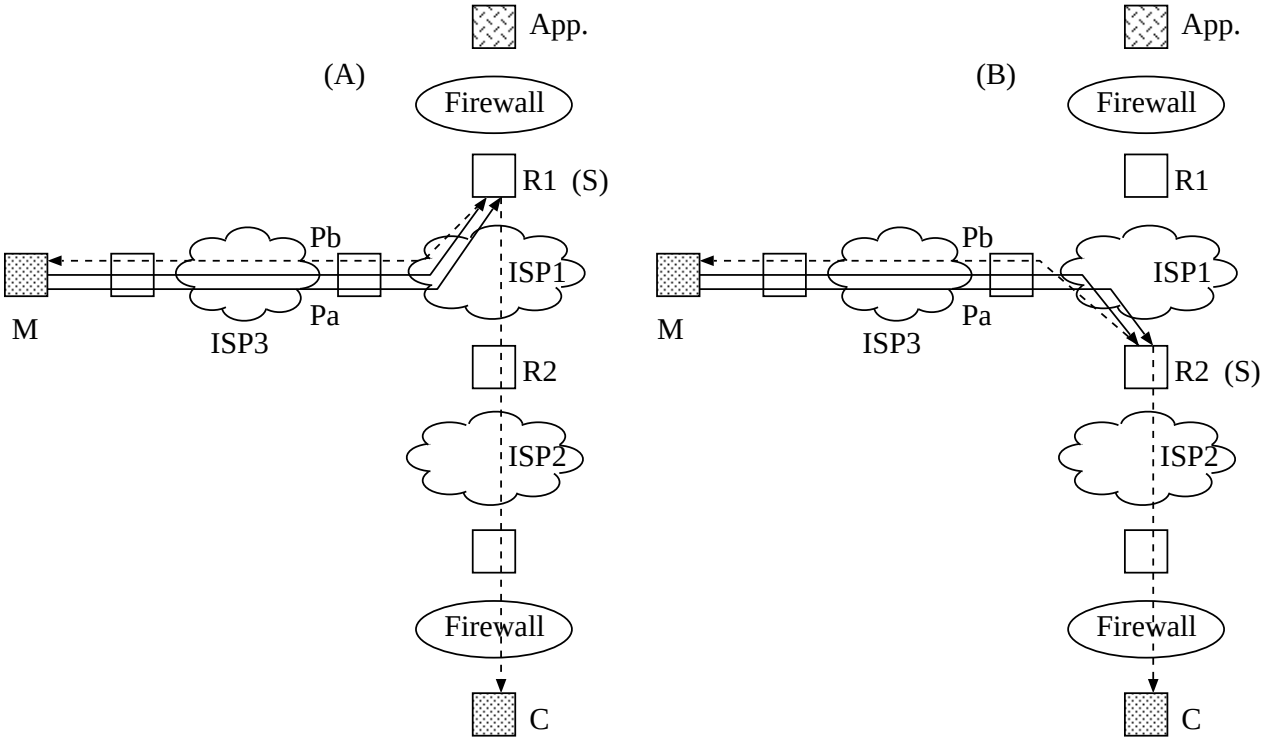

(A) Loss rates from a boundary R_1 to C

(B) Loss rates from an intermediate R_2 to C

Experiments

Three kinds of experiments on the Internet:

E1 Errors and convergence for the C-S type:

$\Rightarrow$ The inference errors of loss rates were $\leq 1\%$ within 3000 trials, when using randomly distributed inter-trial time.

E2 Robustness of the inference for the C-S type:

$\Rightarrow$ Two simultaneous instances of the C-S type using different proxy measurement servers returned nearly the same value.

E3 Consistency for the S-A type with the C-S type:

$\Rightarrow$ An instance of the S-A type using a router (as a reflector) far beyond the target and an instance of the C-S type running simultaneously returned nearly the same value.

E1: Inference error and convergence for the C-S

- On a test-bed consisting of four UNIX boxes distributed in Japan, over three months,
- We compared
  - the actual loss rates of probe packets, and
  - the inferred loss rates of probe packets using the C-S type.
- Using three of the four boxes as C, S, or M, we examined several different combinations.

Parameters	Tested values
Probe packet type	UDP-echo, UDP+ICMP
UDP packet size	$64, 256, 1400$ (bytes)
Inter-trial time distribution	fixed, uniform, exponential
mean inter-trial time	$0.1, 0.05$ (sec)
N (# of trials for inference)	$1000, 2000, 3000$

(1) The main result:

The inference error of loss rates can be $\leq 1\%$ within 3000 trials, when using randomly distributed inter-trial time.

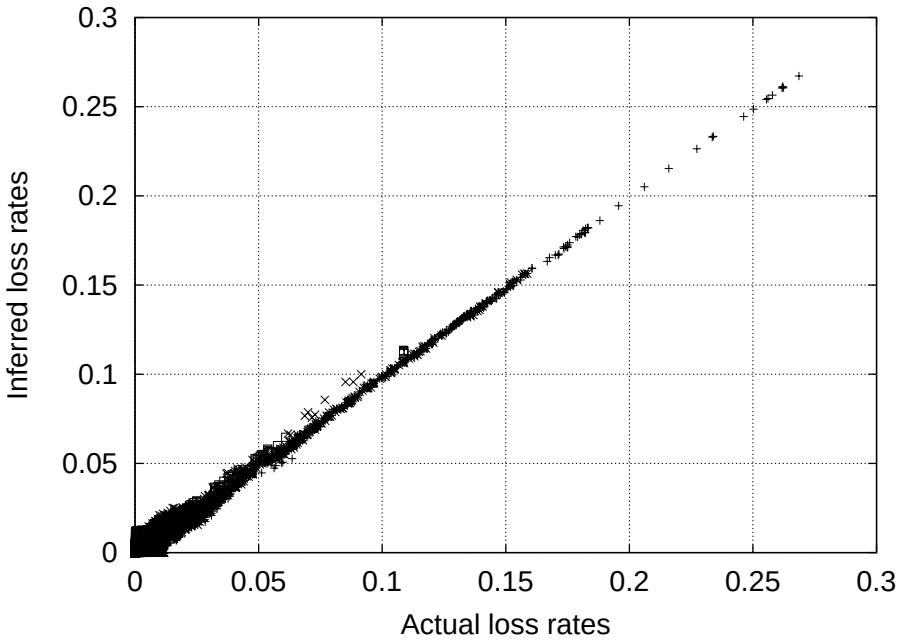

(2) An example of inference trackability:

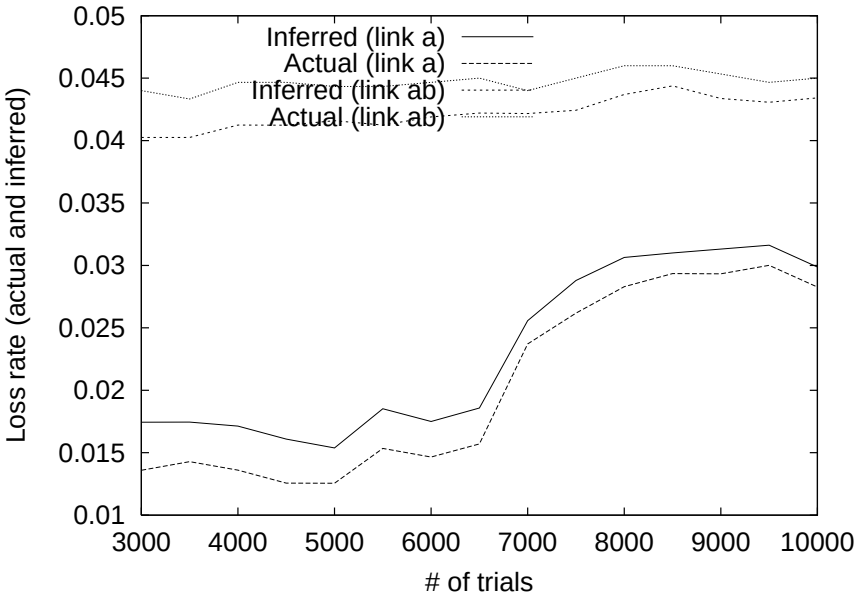

- We run 10000 successive trials with exponentially distributed, 0.1 sec mean inter-trial time.
- Mean loss rate over each 5 mins (3000 trials) are calculated using a moving measurement window having an overlap of 50 secs.
- The inferred values well track the actual values.

E2: Robustness of the inference for the C-S

- Two instances of the C-S type using different proxy measurement servers run simultaneously:

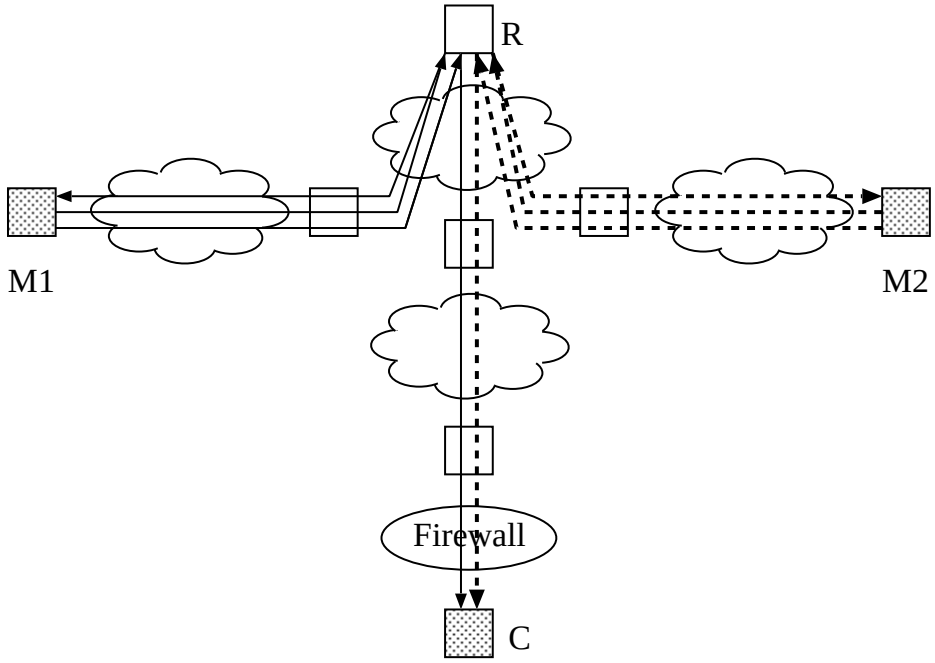

Two simultaneous instances can return nearly the same value.

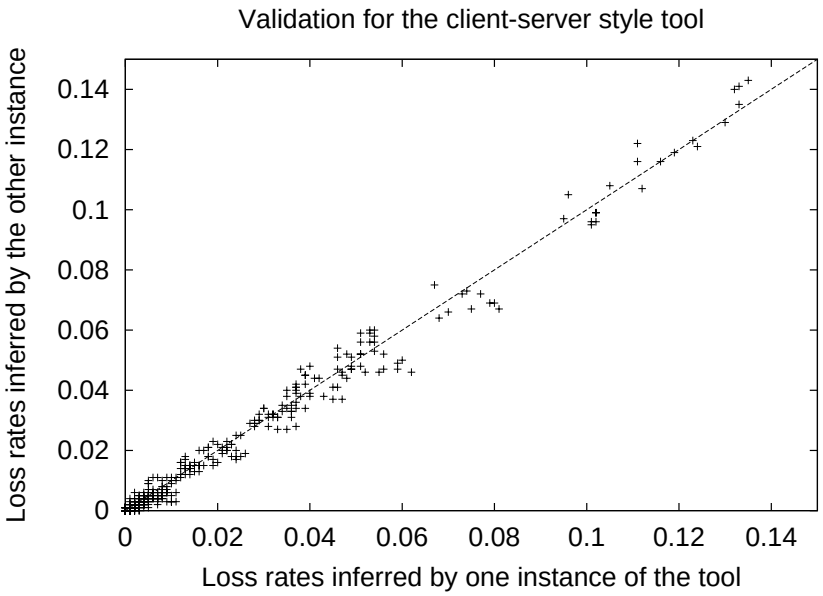

E3: Consistency for the S-A with the C-S

- Three instances run simultaneously:
  - Two instance of the S-A type:
    One using a router far (7-hops) beyond the target; and The other using a next-hop router
  - An instance of the C-S type

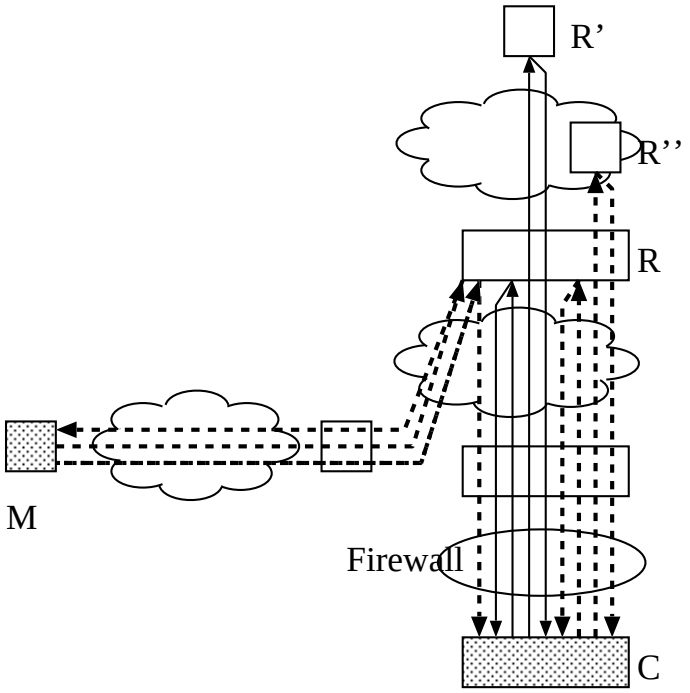

An instance of the S-A using a router far beyond the target can return nearly the same value as that returned by the instance of the C-S type.

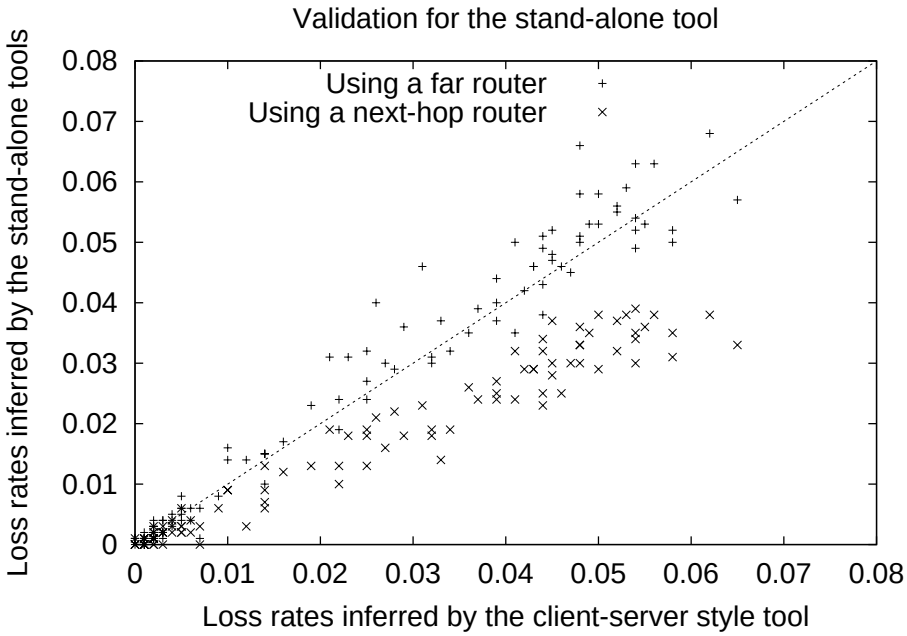

Concluding remarks

- New tools inferring one-way packet loss rates from/to the user-host to/from a target-host.
- For Stand-alone type running alone,
  - Use of LSRR is difficult in the current Internet.
  - If an appropriate host (beyond the target) can be used as a reflector, the tool can work well without use of LSRR.
- For Clinet-server type requiring cooperation with a proxy measurement server,
  - If use of modified IP source address is allowed for proxy measurement servers, the tool can work well with 1% error,
  - Although it can infer only downstream-direction one-way loss rates.
    $\Rightarrow$ Upstream-direction one-way loss rates can be inferred by being combined with round-trip loss rates estimations.

Appendix

Security Issues

- We should consider security issues on the current Internet.
  - The existance of firewalls
  - The prohibition of IP source route opt. (LSRR) and modified source IP address (MSA)
- Firewall (FW)
  - An app. server as S is often placed behind a FW.
    $\Rightarrow$ A boundary router can be S, instead of the server.
  - Client C may also be placed behind a FW.
    $\Rightarrow$ A dummy packet sent from C may make the FW accept some other measurement packets sent to C.
  - Proxy meas. server M can be placed outside FW(s).

- IP source route opt. (LSRR) and modified source IP address (MSA)
  - LSRR is the very function for controlling the route of a packet to travel along tree-structured paths, but is blocked by many routers. $\Rightarrow$
    - Use some auto-reply function like ICMP.
    - For (I), if another router exists beyond the target router, it can be a reflector of P_b to be returned to C.
    - For (II) and (III), sending P_a or P_b with MSA can let S return the packet to another node rather than the sender.
  - However, MSA may be blocked when it enters the Internet, e.g., at an ingress router.
    - it may be difficult for user-host C to send MSA packets in (II)
    - proxy meas. server S can be placed so as to be able to send MSAs in (III).