UiO : Faculty of Mathematics and Natural Sciences
University of Oslo

Performance Evaluation of In-network Packet Retransmissions using Markov Chains

Runa Barik, Michael Welzl, Peyman Teymoori, **Safiqul Islam**, Stein Gjessing

ICNC 2020
Kona, Hawaii

Overview

- Idea and contributions
- LOOP
  - applications
  - analysis
- Observations
- Future work

Idea

- The famous end to end paper [1] suggests that functions placed at low levels of a system may be redundant or of little value when compared with the cost of providing them at that low level
- One of the cornerstones of the Internet Design
  - Strictly avoiding placing application specific functions in the network.

[1] J .H. Saltzer, D. P. Reed, and D. D. Clark, “End-to-end Arguments in System Design,” *ACM Trans. Comput. Syst.*, vol. 2, no. 4, Nov. 1984.

Idea

- Retransmission of packets inside the network on a segment of an end-to-end path
  - especially when packets are frequently dropped on path segments with a very short round-trip time (RTT), and the end-to-end RTT is very long
- Similar techniques: Automatic Repeat Request (ARQ) in wireless links
  - but it is a result of the separate development of Internet standards and link layer standards

Our Contribution

- We consider only retransmission, not ordering
- Questions to answer: given the end-to-end RTT, the RTT on a path segment, and the packet loss probabilities of the various parts of an end-to-end path,
  - how long should a system ideally store packets for retries?
  - how large is the potential benefit from in-network retransmission?
- We never delay packets

Our Contribution

- We introduce a stochastic model that describes the relationship between packet drop probabilities in different path segments, RTT, and buffer size.
- We investigate the influence of a local loss recovery mechanism: caching packets in an intermediate node and loss detection before packets being arrived at the destination using the above model.
  - We derive the impact of each parameter on the system behavior.

Local Optimizations On a Path (LOOP)

- We utilize a cache and a loss detector at the start and the end of the path segment with high packet loss:
  - sender S emits packets with rate λ
  - cache C additionally stores a copy of packets
  - loss detector L confirms reception of each packet
  - C retransmits if timer rto_c expires
  - R acknowledges packets to S

p_i: packet loss probability in segment i

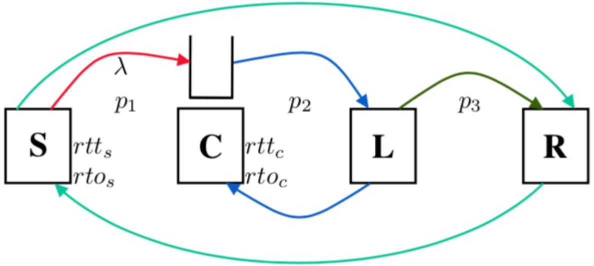

LOOP architecture

Applications of LOOP

- This closely matches IETF LOOPS
  - tunnel mode: C encapsulate packets in addition to caching
  - transparent mode: C caches packets together with a hash identifier that is calculated from immutable header fields
- Recursive Internetworking Architecture (RINA)
  - evaluating retransmissions in stacked layers with different scopes
- Information-Centric Networking (ICN)
  - R^2T: a close proposal in ICN

LOOP Analysis

- We use a Continuous-Time Markov Chain (CTMC) with finite states
- Cache can be modelled as an M/D/1/N queuing model
  - Poisson arrival
  - deterministic service time
  - N: buffer size

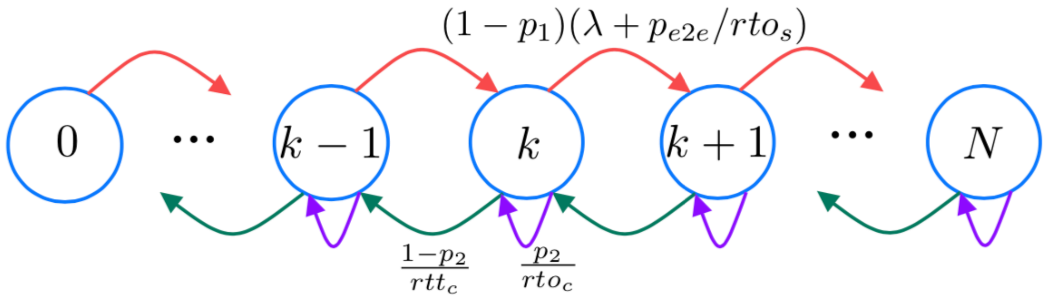

Solving the CTMC

- E2E packet drop prob.: $$p_{e2e} = p_1 + (1-p_1)p_b p_2 + (1-p_1)(1-p_b p_2)p_3$$

- Arrival rate at the cache: $$\lambda_1 = (1-p_1)(\lambda + \frac{p_{e2e}}{rto_s})$$

Service rate → $\mu = \frac{1-p_2}{rtt_c}$

$\rho = \frac{\lambda_1}{\mu}$ ← Load at C

- The M/M/1/N equivalent:

$$\pi_0 = \frac{1-\rho}{1-\rho^{N+1}} ,$$

$$\pi_k = \rho^k * \pi_0 , \quad k = 1, \cdots , N$$

$$E[X] = \frac{\rho}{1-\rho} - \frac{(N+1)\rho^{N+1}}{1-\rho^{N+1}} ,$$

$$E[\lambda_1] = \lambda_1(1-\pi_N) ,$$

$$E[T] = \frac{E[X]}{E[\lambda_1]} ,$$

$$p_b = \frac{(1-\rho)\rho^N}{1-\rho^{N+1}} ,$$

Cache Size

- As a Function of p_2 and p_b

$$N = \left\lceil \log_\rho \frac{p_b}{1 - \rho + \rho\, p_b} \right\rceil$$

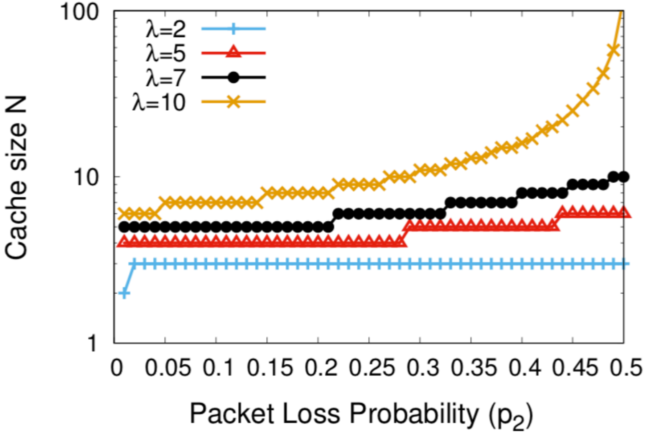

(a) The required cache size (N) at C (in packets)

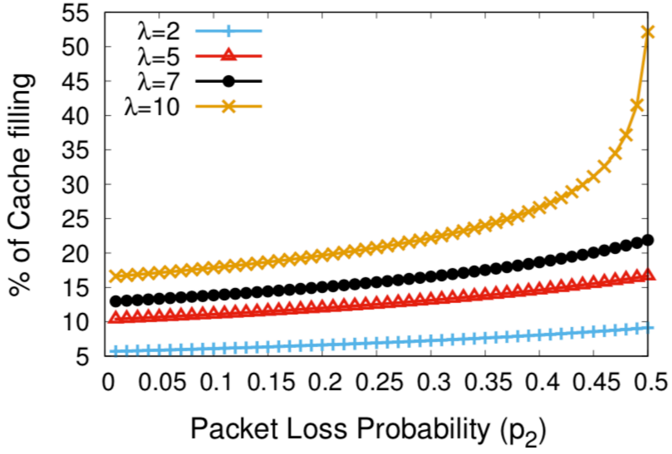

(b) The filling percentage of the cache ($100\frac{E[X]}{N}\%$) as a function of p_2

$p_1 = 0.01$, $p_3 = 0.01$, $rtt_s = 0.1$, $rto_s = 0.2$, $rtt_c = 0.05$, $rto_c = 0.1$, initial $p_b = 0.01$

Effects of rtt_c

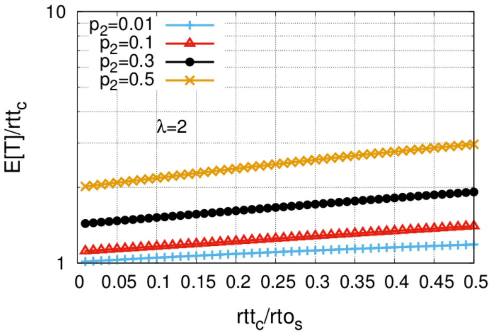

(a) $\lambda = 2$

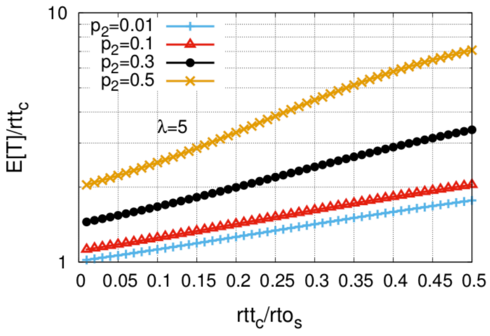

(b) $\lambda = 5$

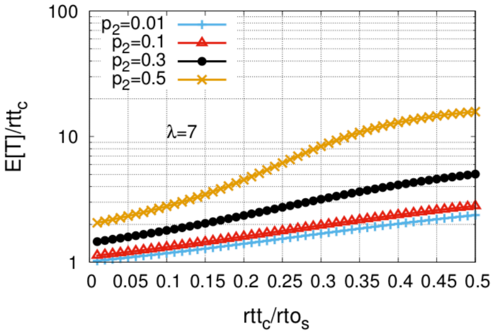

(c) $\lambda = 7$

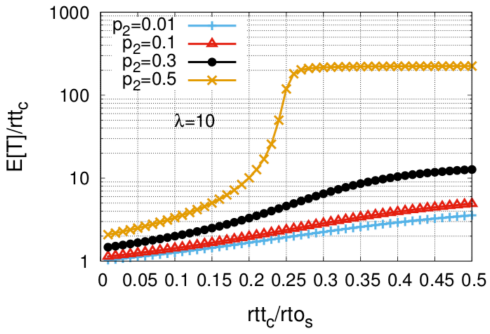

(d) $\lambda = 10$

$p_1 = 0.01$, $p_3 = 0.01$, $rtt_s = 0.1$, $rto_s = 0.2$, initial $p_b = 0.01$

Fixed Cache Size Independent of p_2 and p_b

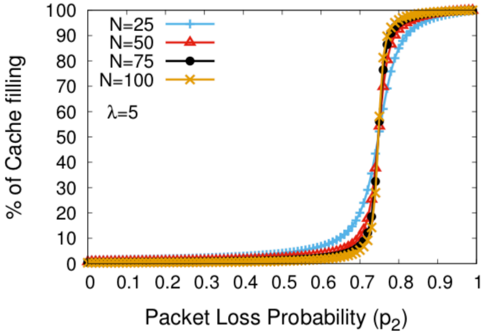

(a) $\lambda = 5$

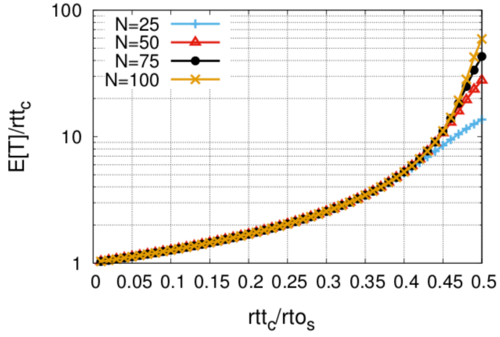

(a) Varying the cache size

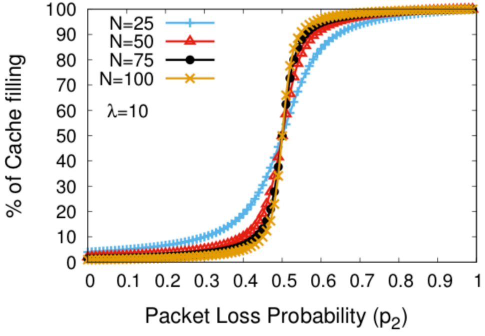

(b) $\lambda = 10$

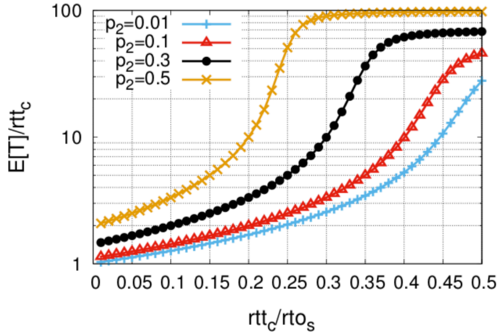

(b) Varying p_2

$p_1 = 0.01$, $p_3 = 0.01$, $rtt_s = 0.1$, $rto_s = 0.2$, $rtt_c = 0.05$, $rto_c = 0.1$, initial $p_b = 0.01$

End-System Analysis

- Buffer could be modelled as an M/D/1/∞ queue

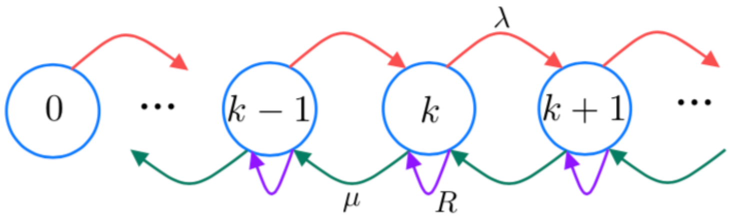

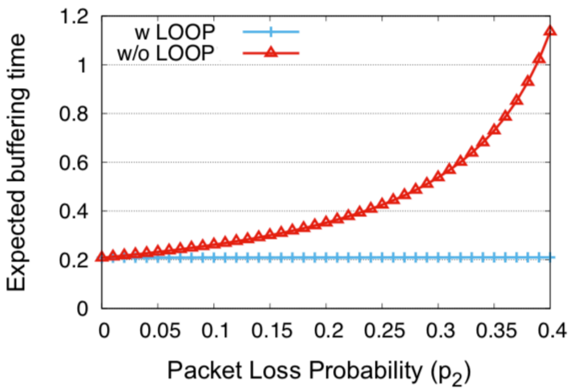

(a) Expected waiting time of packets at S

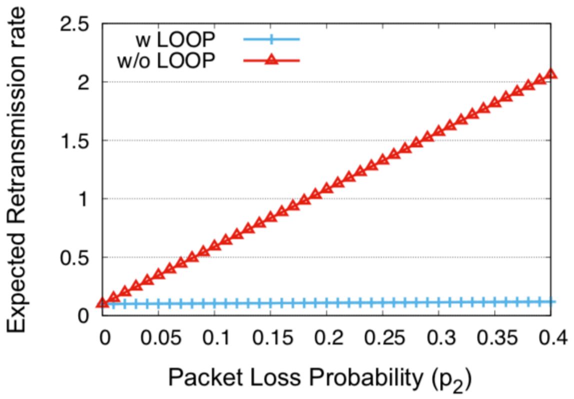

(b) Retransmission rate at S

$\lambda = 5$, $p_1 = 0.01$, $p_3 = 0.01$, $rtt_s = 0.1$, $rto_s = 0.2$, initial $p_b = 0.01$

Observations

- Non-linearity dependence between filling the cache and the packet loss probability
- Cache size has the least impact on its utilization:
  - irrespective of the size, it can be fully utilized by higher packet loss probabilities
- With LOOP: lower retransmission rate and expected caching time at S than without LOOP being deployed in the network

Future Work

- Trade-offs involved in re-sequencing
- Using NACKs
- Sequence of or even nested loops

Thank you!