Evaluation of LISP+ALT performance

LISP WG, IETF-75, Stockholm

Loránd Jakab, Albert Cabellos, Florin Coraş, Jordi Domingo
Universitat Politècnica de Catalunya

Damien Saucez, Olivier Bonaventure
Université catholique de Louvain

Motivation

- How would a LISP ITR perform in the current Internet?
- Performance of ALT?
- Current testbed too small to get some approximate performance numbers
- ALT has to be deployed in a scalable and efficient manner
- We propose the CoreSim simulator to get an idea of global ALT performance

A 3-layer ALT hierarchy

- No description in the ALT draft and no consensus on the mailing list about how ALT will be deployed on global scale
- L1 – fully meshed root layer
- L2 – /8 aggregation
- L3 – Map-Server
- L3 = current BGP
- No peering on L2

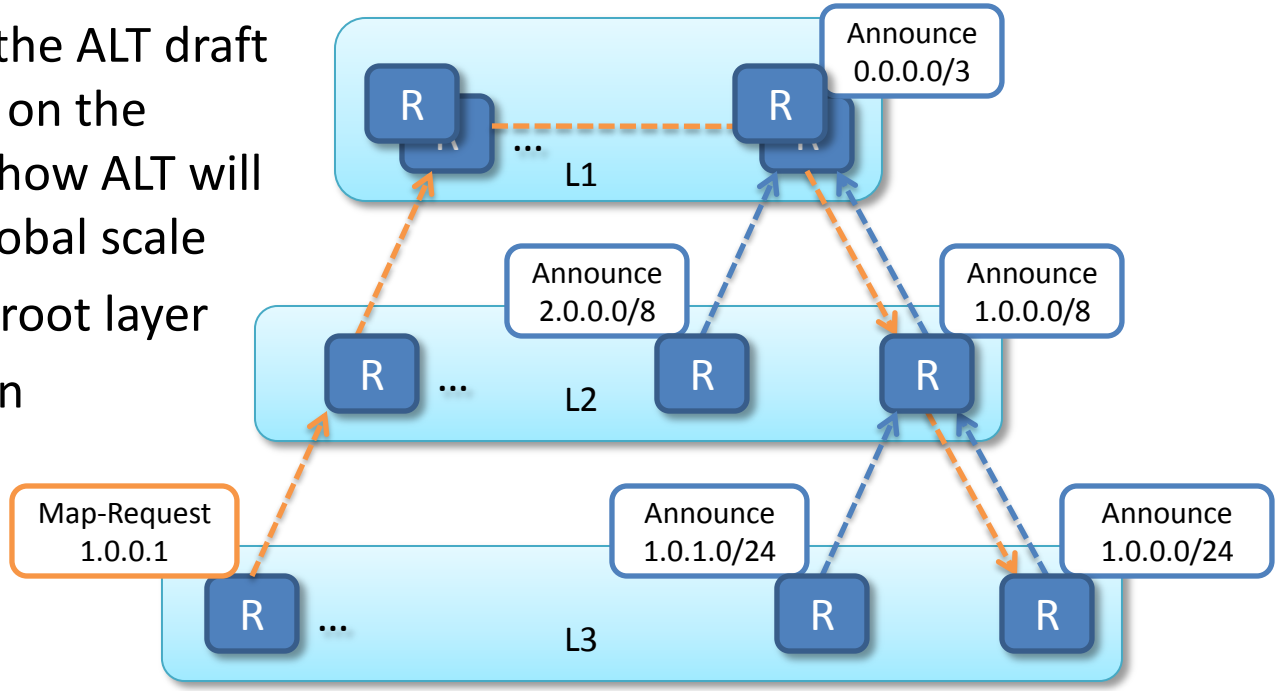

Topology

- Using the iPlane infrastructure (U. Washington):
  - DFZ prefix list
    - We filtered longer prefixes included in shorter
    - We have 112.233 prefixes after filtering
  - AS connectivity
  - Latency between arbitrary IPs
    - We observed about 65-80% coverage
- Apply to the 3-layer ALT

CoreSim

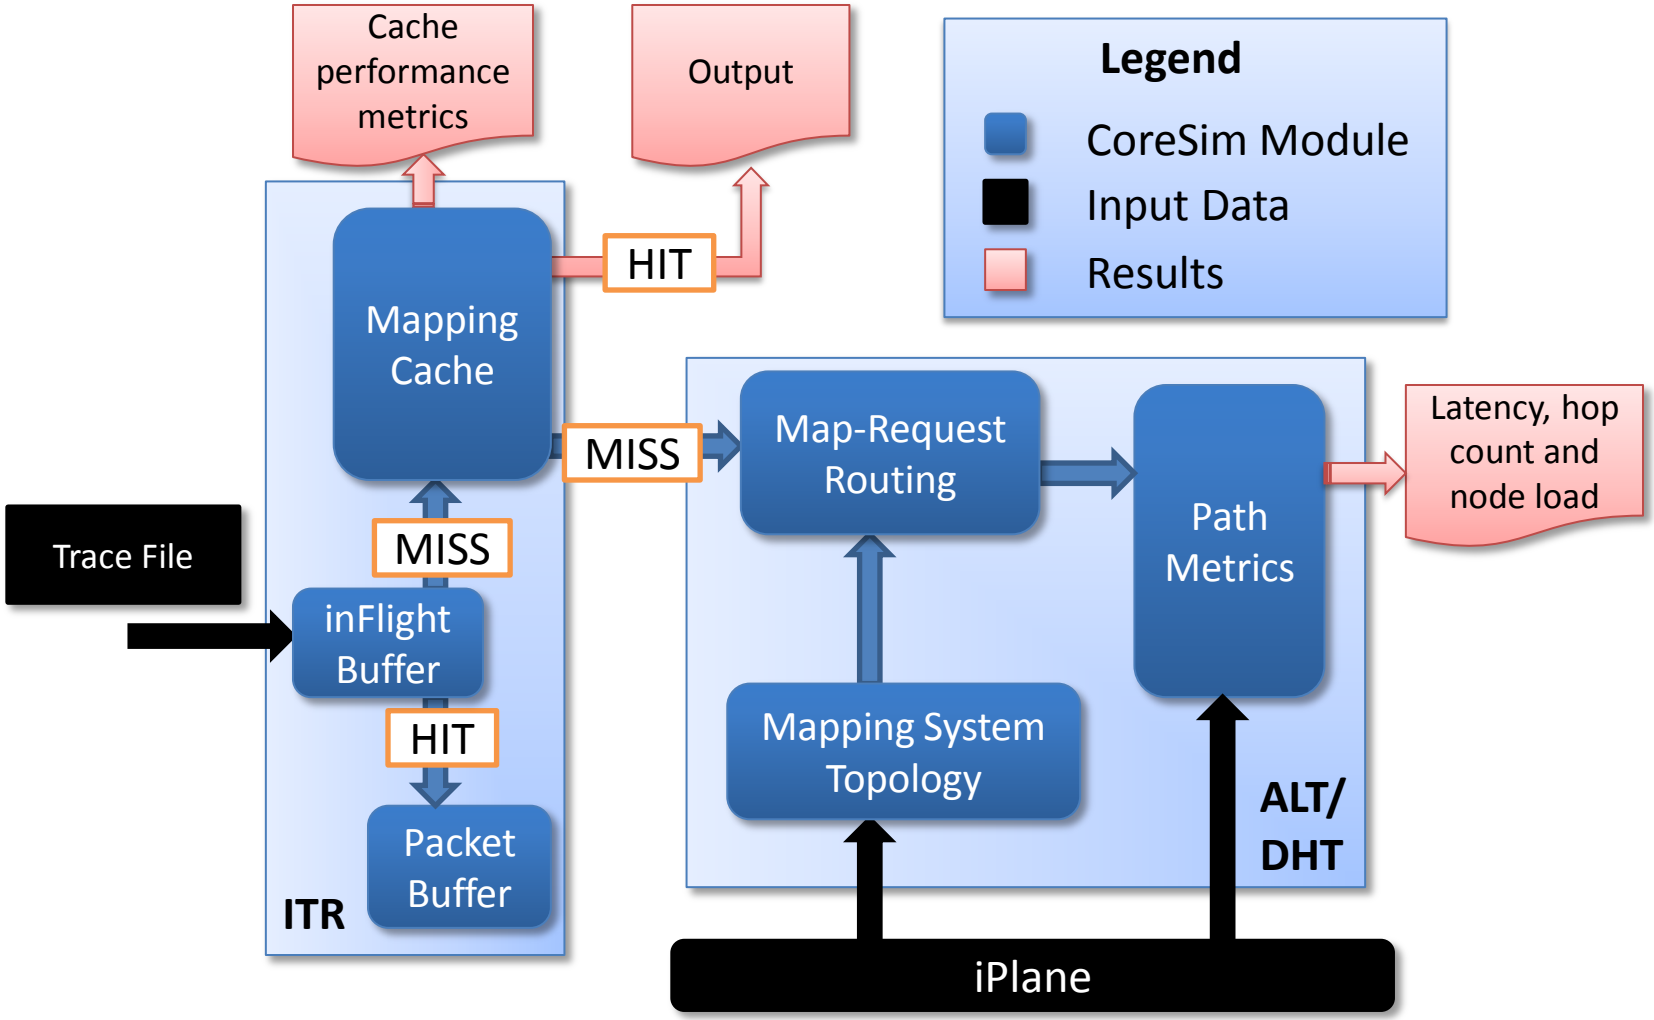
Cache performance metrics
Output
Legend
CoreSim Module
Input Data
Results
Mapping Cache
HIT
MISS
Map-Request Routing
Path Metrics
Latency, hop count and node load
Trace File
MISS
inFlight Buffer
HIT
Packet Buffer
ITR
Mapping System Topology
ALT/ DHT
iPlane

Traces

- **24h** egress traffic @ **UCL** border router, Louvain (03/23/2009)
  - 752 GB / 1200 M packets = 69 Mbps avg. BW
  - 4.3 million IPs / 123,804 BGP prefixes

- **4h** egress traffic @ **UPC** border router, Barcelona (05/26/2009)
  - 463 GB /1200 M packets = 289 Mbps avg. BW
  - 4.3 million IPs / 111,492 BGP prefixes

Simulation Results

- About 10 days on Core 2 Xeon for each trace / MS combo
- Map-Request RTT:

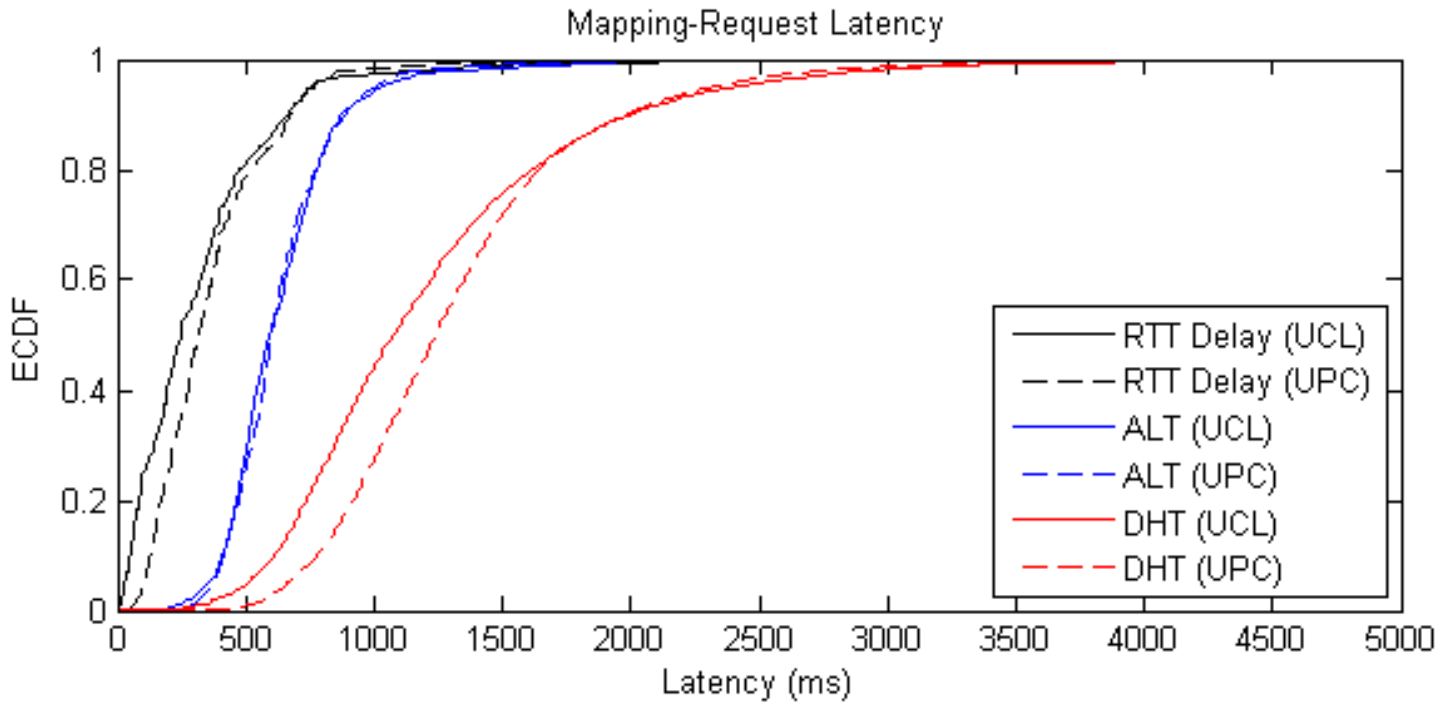

Simulation Results (cont.)

- Hop count:
  - 95% of the time is 6 hops for ALT: to the root and down to L3
- Load:
  - Very non-homogeneous in ALT, due to uneven IPs/prefix distribution
  - In DHT has an interesting property: the first prefix after a large unallocated space has significantly more load

Dropping vs. Buffering

- How big a buffer do we need for “normal” traffic?
- Cache hit ratio of 99.5% for our traces
- Simulator replays trace, does not emulate connection setup → **worst case values**
- Median values of buffer occupancy:
  - ALT: **86 packets / 65 KB**
  - DHT: 136 packets / 114 KB
- Traffic anomalies (malicious or benign) cause important spikes: maximum value: **70 MB** !!!

Future Work

- Evaluate other possible ALT deployment scenarios?
- Different EID distribution
- Cache eviction algorithms
- Other traces
  - E.g. : content providers (vs. educational networks)
  - Simulator is open source, feedback and results with your data is welcome

Draft ?

- ALT deployment recommendations draft?

http://www.cba.upc.edu/lisp

Buffer Occupancy (bytes)