An Empirical Mixture Model for Large-Scale RTT Measurements

Romain Fontugne[1,2] Johan Mazel[1,2] Kensuke Fukuda[1,3]

[1]National Institute of Informatics [2]JFLI [3]Sokendai

June 9, 2015

Introduction

RTT: Round Trip Time

- Time to send a packet and receive its acknowledgment
- Key indicator of network conditions
- Important for server selection, overlay network, geolocation...

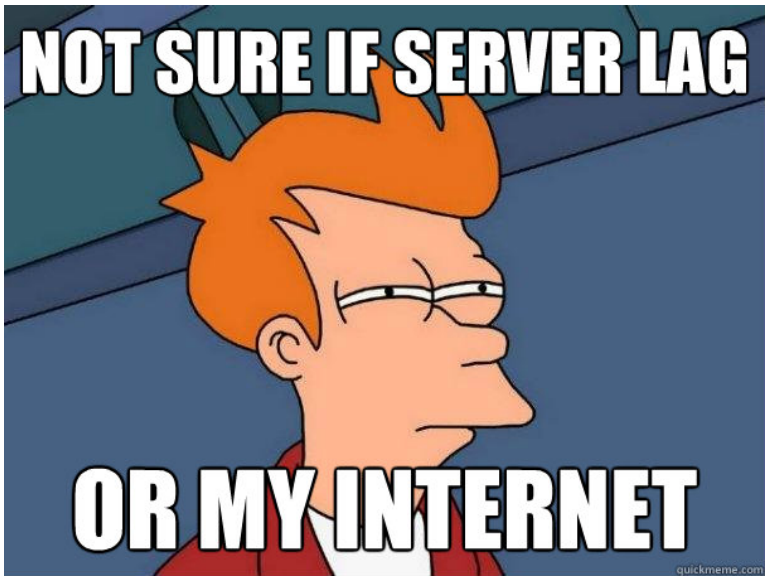

Motivations

Monitor Internet-wide delays over time

- Measure millions of hosts RTTs
- Assess network performance at large-scale
- Report significant RTT fluctuations

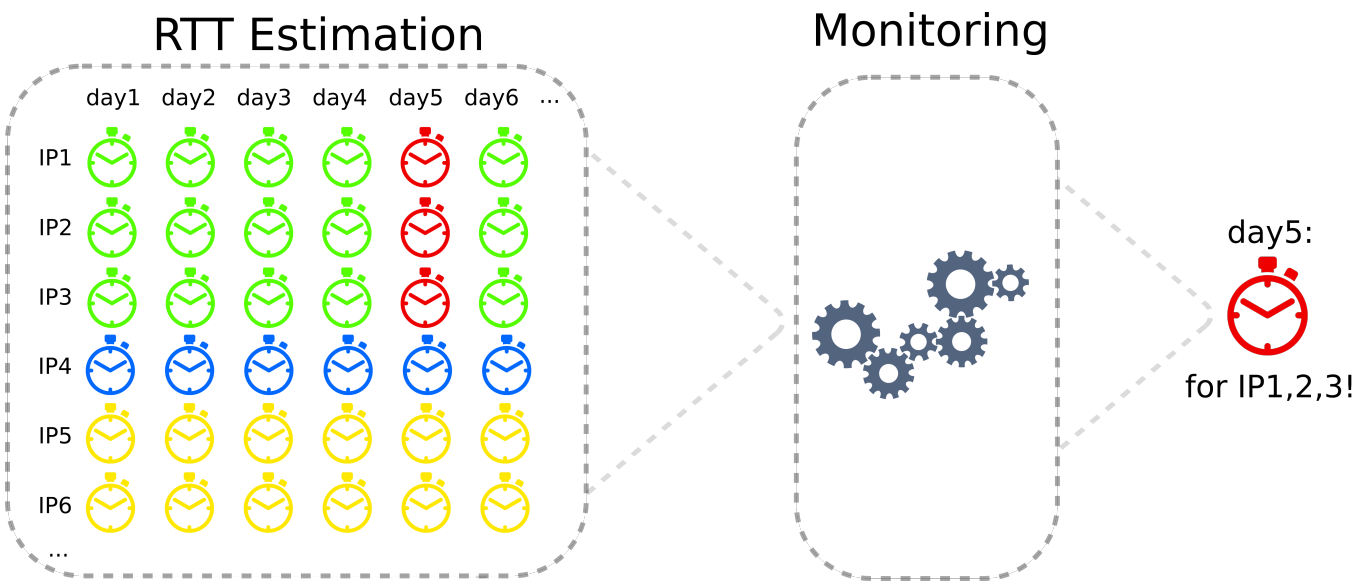

RTT Estimation

RTT from passive measurements

- Measure traffic at backbone network
- RTT estimation from TCP traffic

Advantages

- Non-Intrusive (~~Ping the entire IP space~~)
- Monitor RTT experienced by Internet users

RTT Estimation in the network

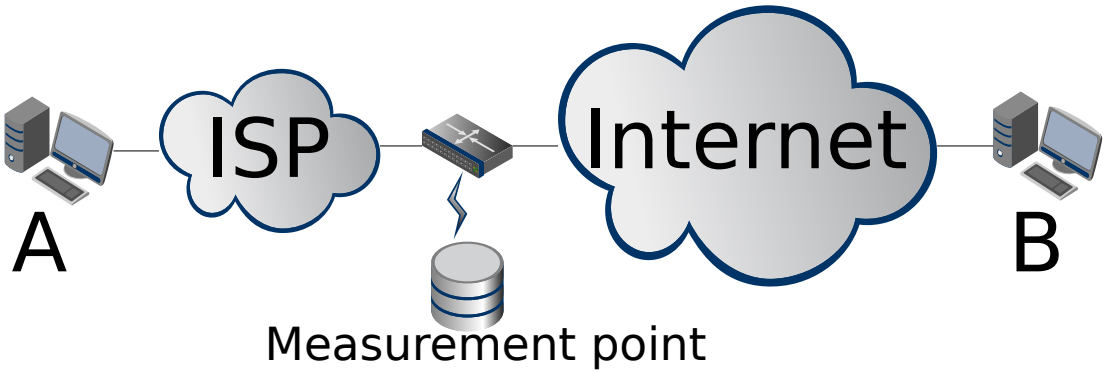

RTT Estimation in the network

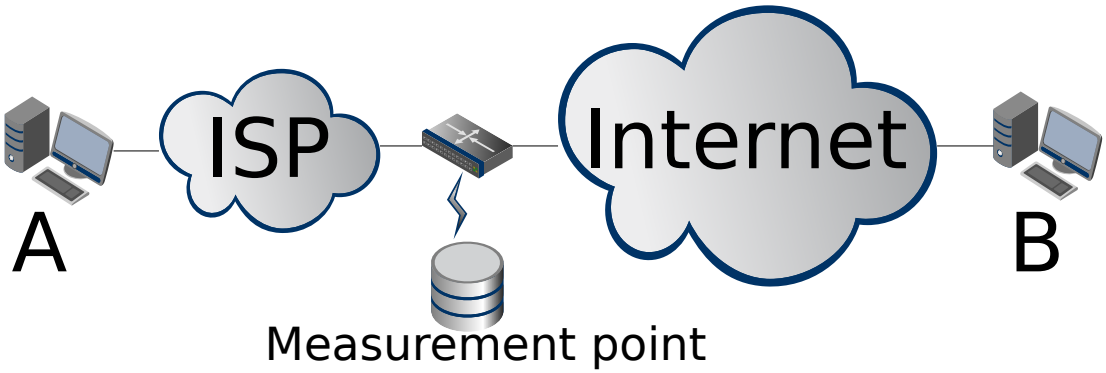

Based on Karn's algorithm

- For a certain host A
- Compute delay samples
  $\delta = \theta_{ACK} - \theta_{SEQ}$
- Ignore retransmitted packets

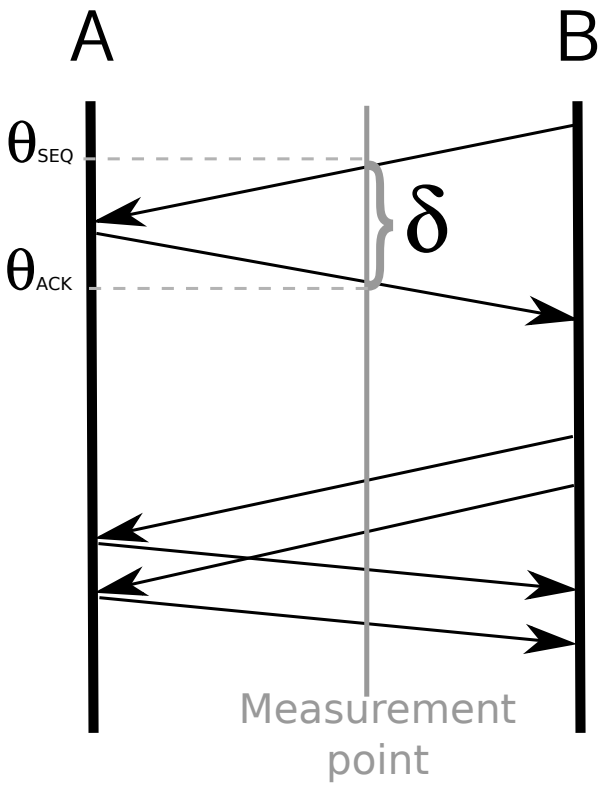

Problem: Understanding RTT from numerous hosts?

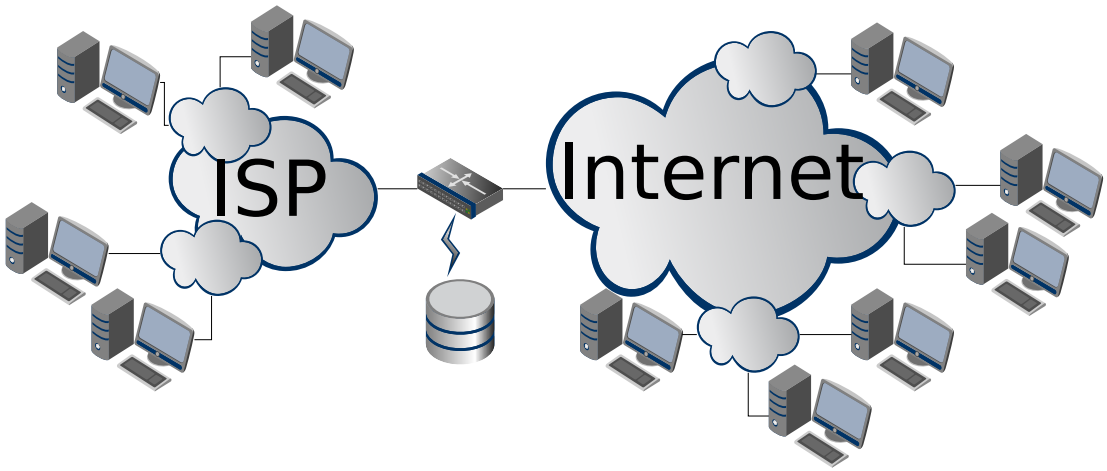

Problem: Understanding RTT from numerous hosts?

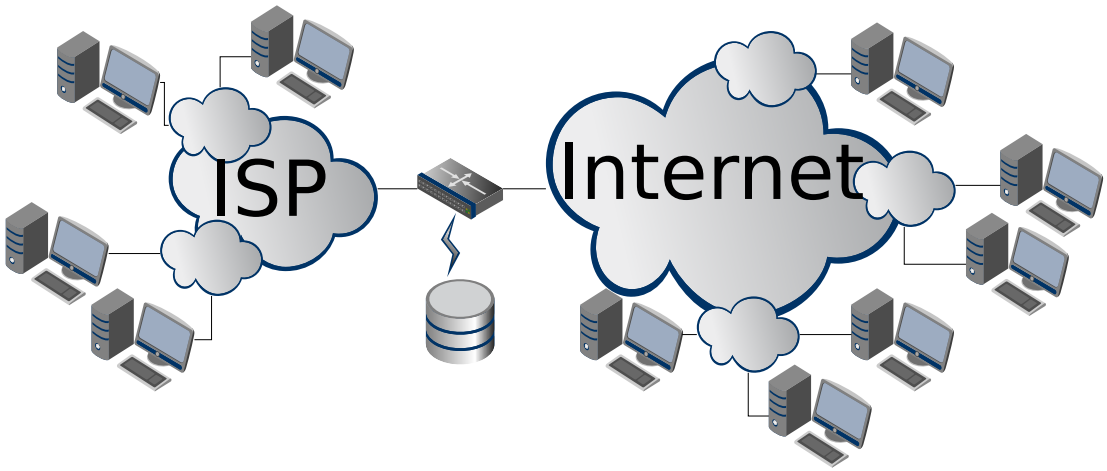

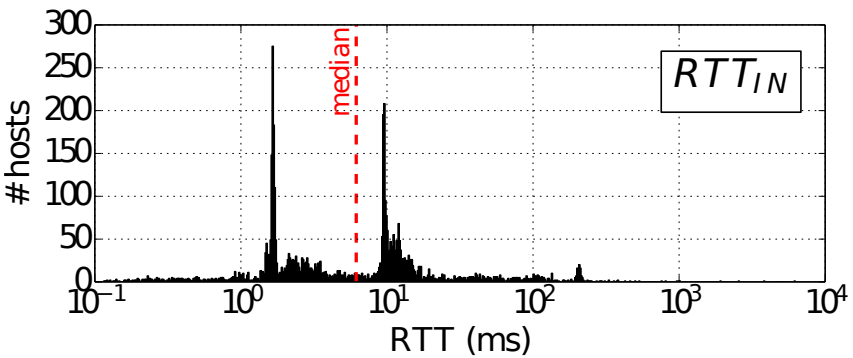

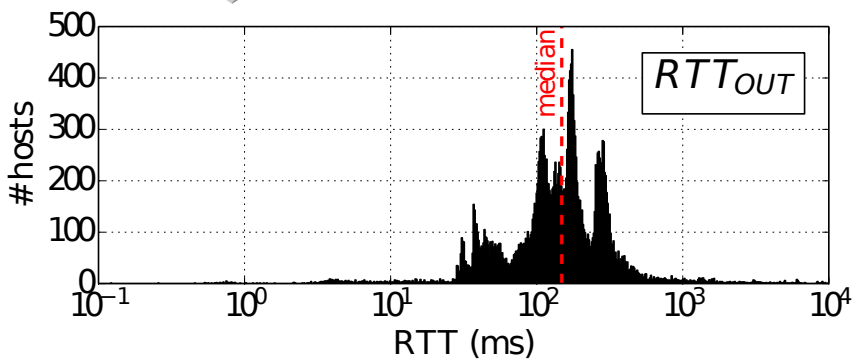

- **Multimodal** distributions
- Median value is misleading! (Don't use it!)

Monitoring RTT

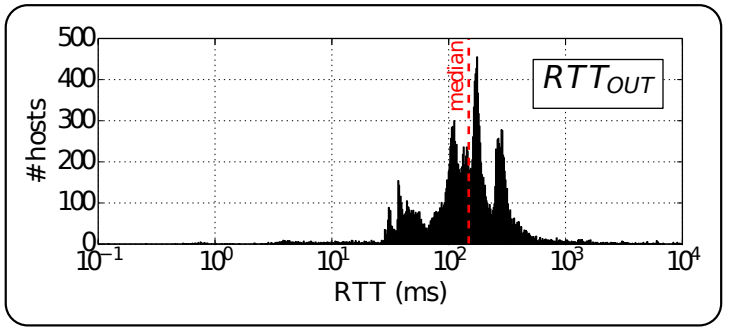

Find and monitor typical RTTs

- Identify usual RTTs experienced by Internet hosts
- Characterize, and monitor spatial and temporal dynamics of RTTs
- Detect abnormal RTTs fluctuations for both host population or specific hosts

Proposed Model

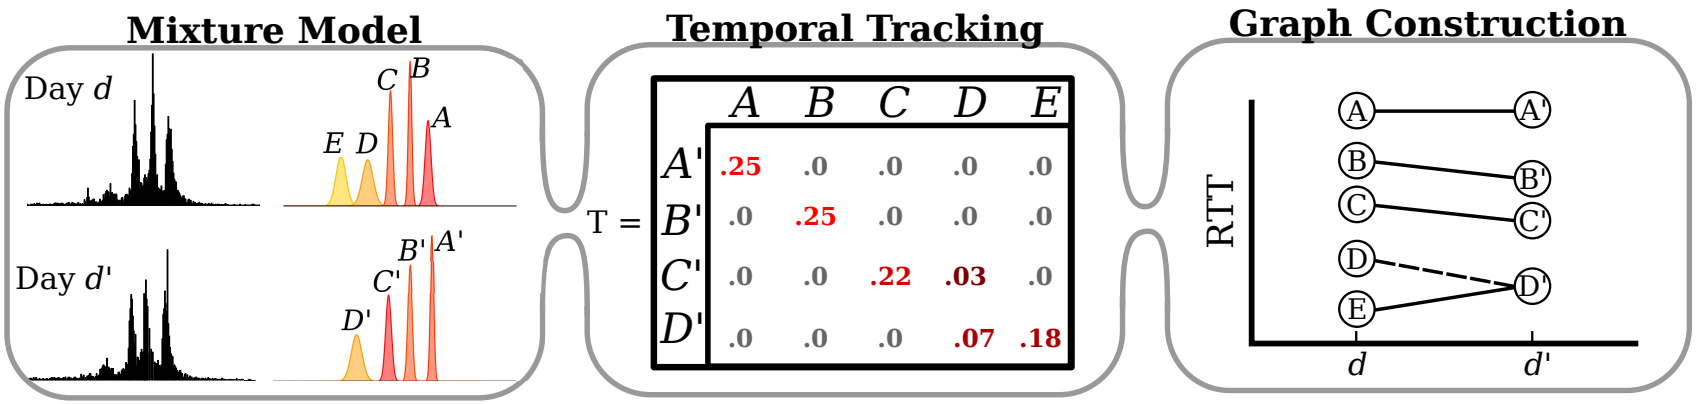

1. Uncover the daily RTT distributions using a mixture model
2. Link RTT distributions from similar sub-population of IPs across time
3. Formalize RTTs time evolution in a graph for further systematical analysis

Mixture Model

Identify RTT sub-populations:

- Unknown number of mixed components
- Dirichlet process mixture model
- log-normal distribution

→ Obtain the mean and std. deviation of typical RTTs (μ, σ)

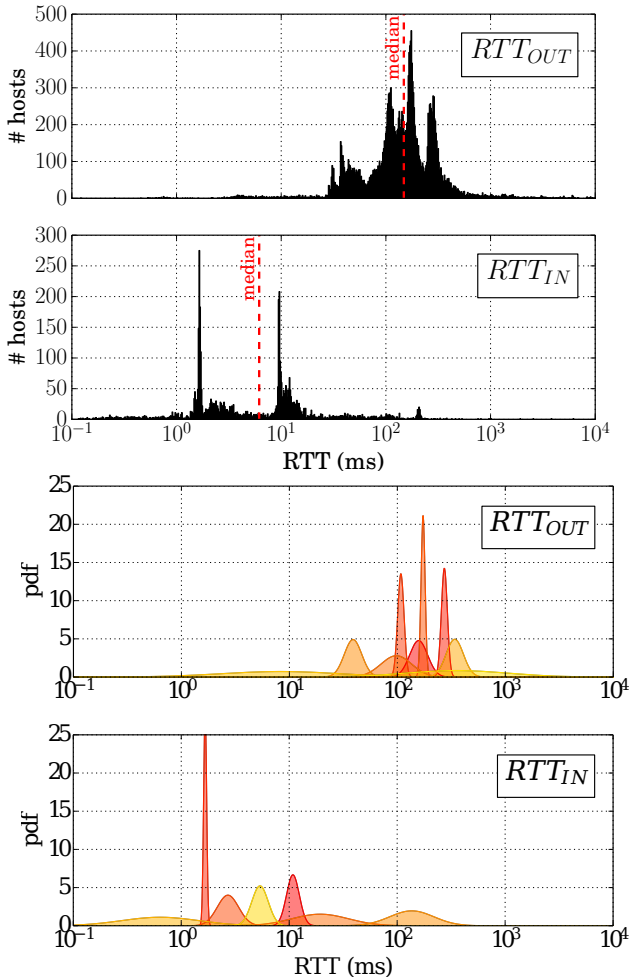

Temporal Tracking

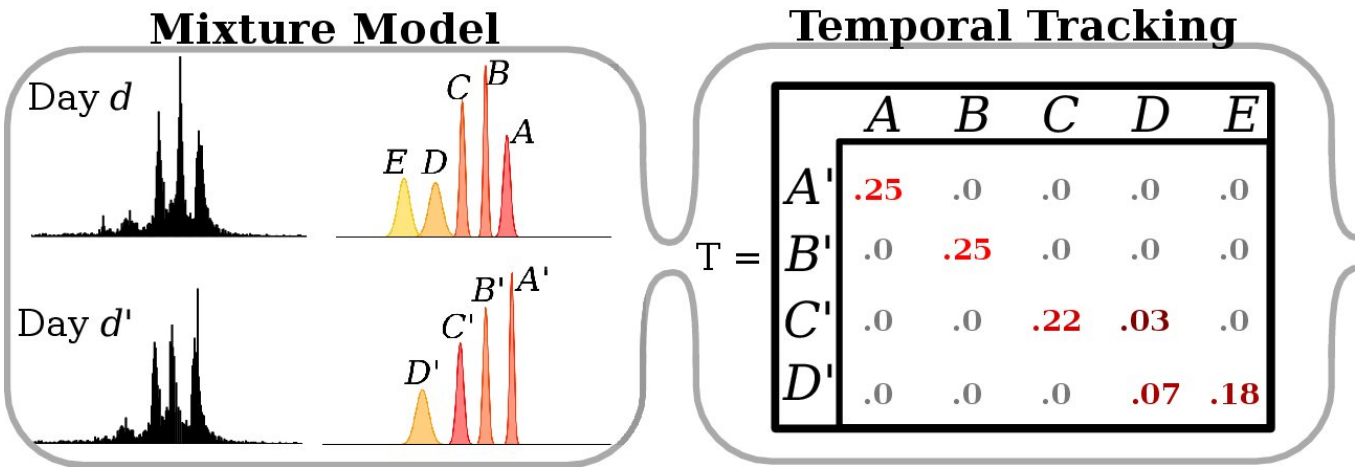

Number of components from d and d' might differ

- IPs from E moved to D'?
- or they are not active in day d'?

Temporal Tracking (cont.)

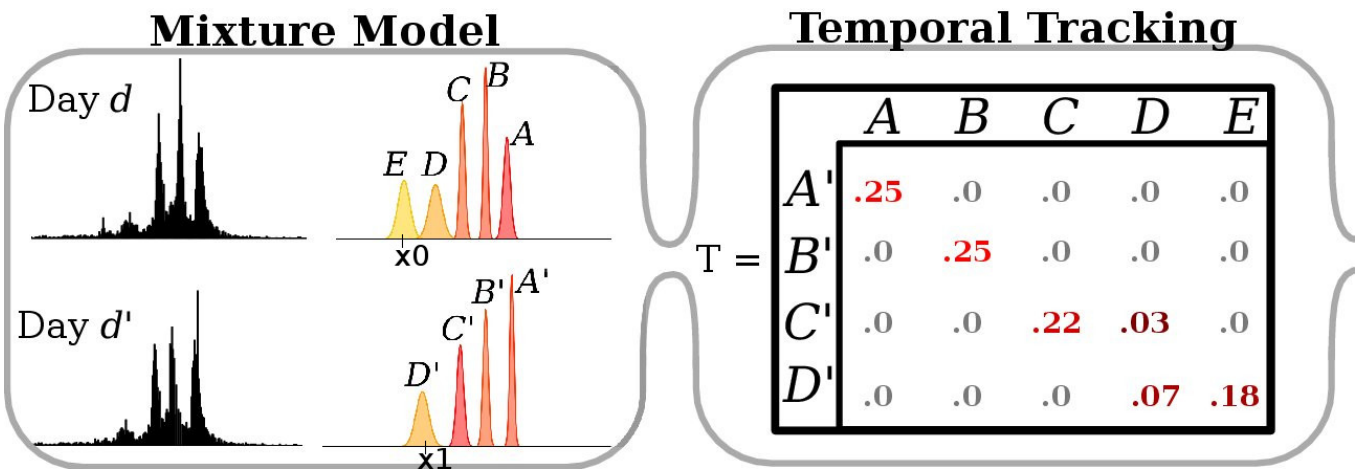

Connect distributions from different days:

- See components as probability density functions
- Compute probability of IPs to fall in A and A', A and B', ...

→ Transition matrix from day d to day d'

Graph Construction

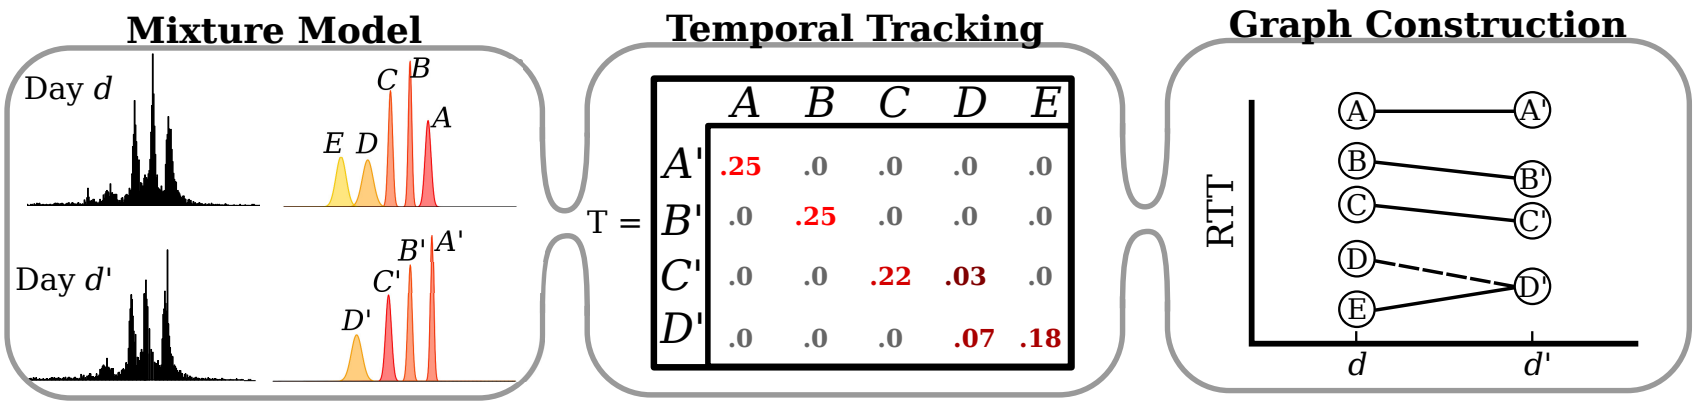

Graph from transition matrix

1. Nodes: identified modes / typical RTTs for one day
2. Edges: relate modes with similar IPs

Evaluation

Dataset: MAWI Archive

- transit link between the WIDE network (ASN2500) and the Internet
- 15 minutes of traffic everyday from Jan. 2001 to Mar. 2014
- 4678 traces (pcap files)
- RTT estimates from 12 millions unique IP addresses
- Separate RTTs to hosts inside the WIDE network (RTT_{IN}) and outside (RTT_{OUT})

Longitudinal Study

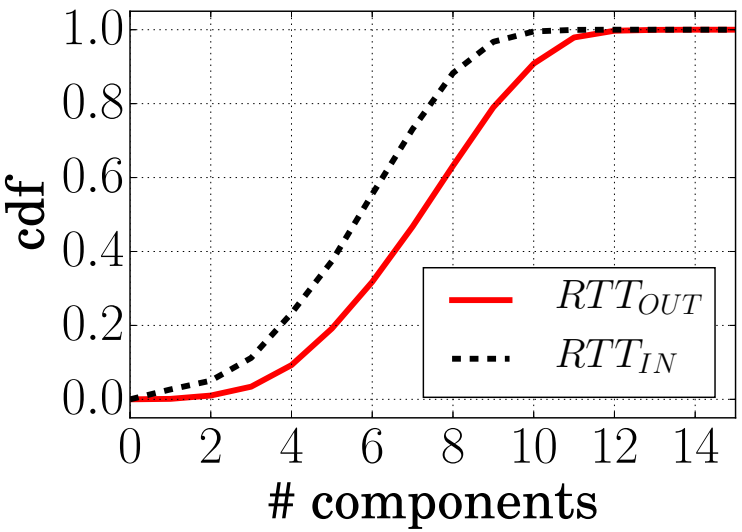

(a) CDF of the number of identified modes per day.

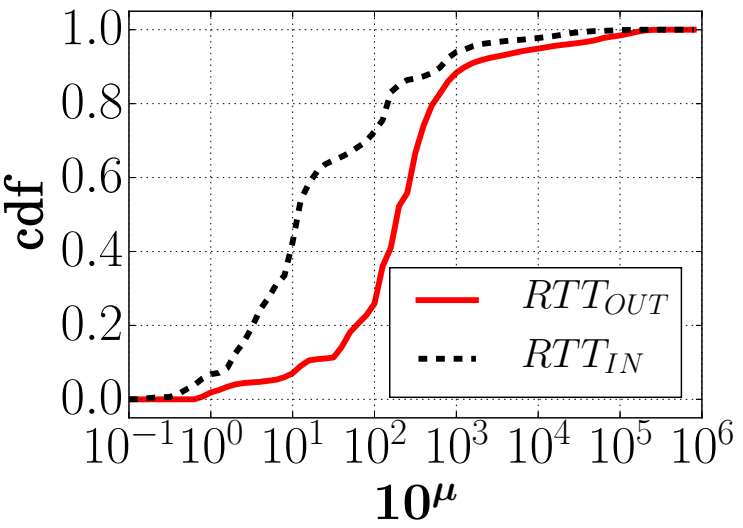

(b) CDF of 10^{μ}, the modeled RTT of the identified modes (milliseconds).

$\rightarrow$ RTT_{OUT} contains more components and higher RTTs

$\rightarrow$ $RTT_{IN} = 500ms$ $(10^{2.7})$ satellite link!

Geolocation

Example: RTT_{OUT} from 2014/03/21 to 2014/03/31

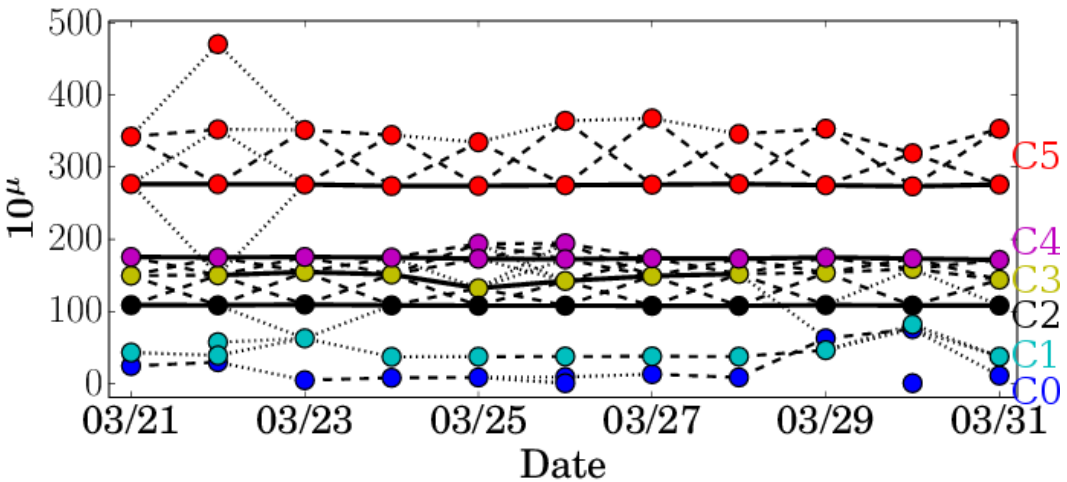

→ Cluster modes using community mining (Louvain algorithm)

Geolocation

Example: RTT_{OUT} from 2014/03/21 to 2014/03/31

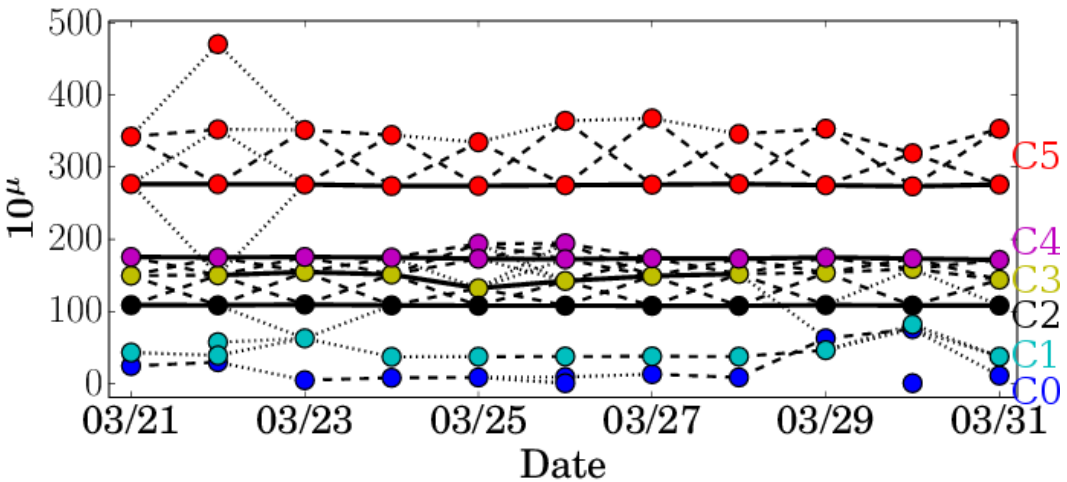

→ Cluster modes using community mining (Louvain algorithm)

	JP	KR	US	CA	EU	CN	$\overline{RTT}$
C5			8%		**73%**	3%	289 ms
C4			**87%**	4%			175 ms
C3			**73%**			11%	149 ms
C2			**91%**				108 ms
C1		**97%**					44 ms
C0	**98%**						19 ms

Table: Hosts geolocation breakdown using Maxmind Geo-IP

Application 1

Look at communities RTT fluctuations:

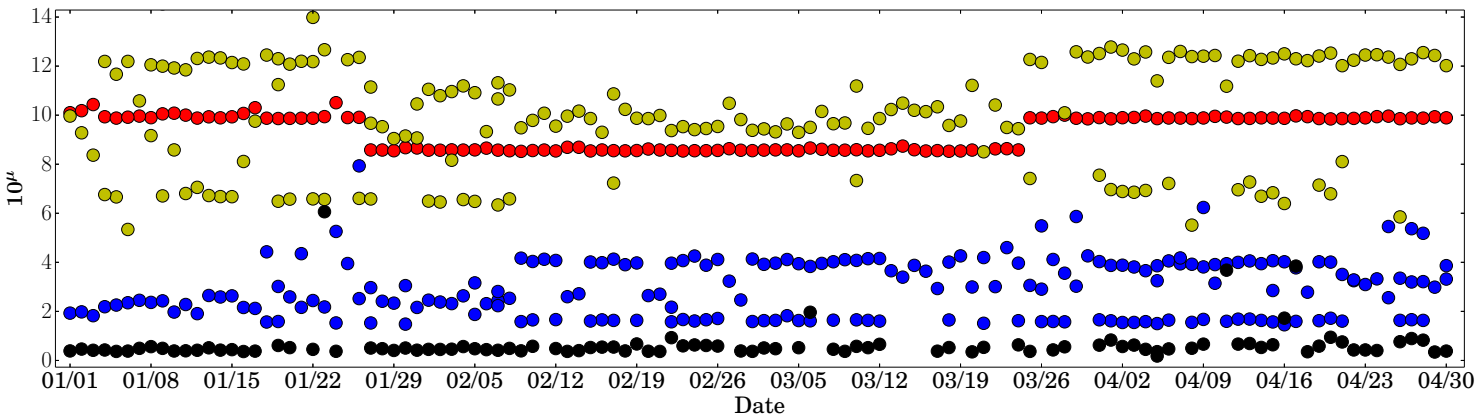

- $\overline{RTT}$ fluctuation = avg. RTT day2 - avg. RTT day1 (normalized)
  - ≈ 0 means the RTTs are stable
  - if deviate from 0 means RTTs of numerous hosts have changed

$\rightarrow$ $\overline{RTT}$ fluctuation depicts important RTT changes

$\overline{RTT}$ fluctuations & BGP updates?

- BGP Route Information Base (RIB) from Route Views Project
- Ratio of IPs affected by a BGP route vs. $\overline{RTT}$ fluctuations:

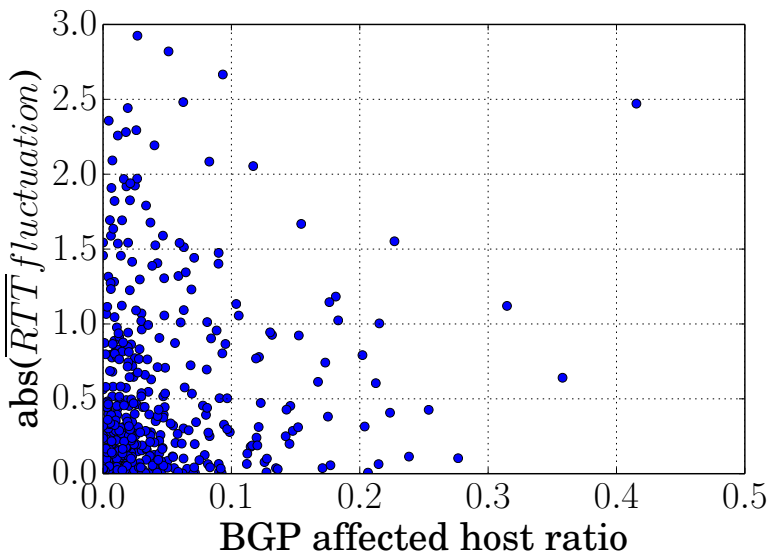

$\rightarrow$ 66% of the BGP updates affecting $>$ 15% clustered IPs exhibit $>$ 0.15 $\overline{RTT}$ fluctuations (similar to Rimondini et al. PAM'14)

$\overline{RTT}$ fluctuations & network congestion?

- Assuming MAWI throughput is proportional to network congestion
- Compare MAWI throughput and $\overline{RTT}$ fluctuations

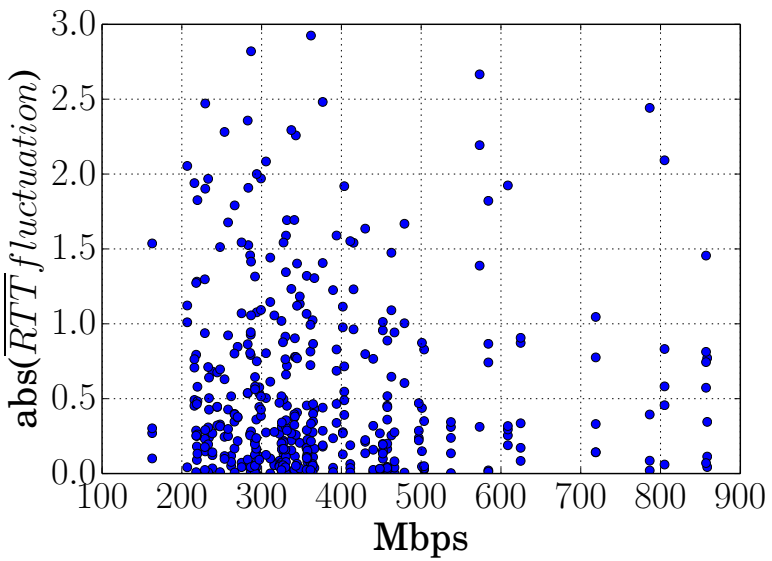

→ higher $\overline{RTT}$ fluctuations when average throughput is higher than 500 Mbps

$\overline{RTT}$ fluctuations: Example

Tohoku earthquake (2011/03/07-2011/03/14)

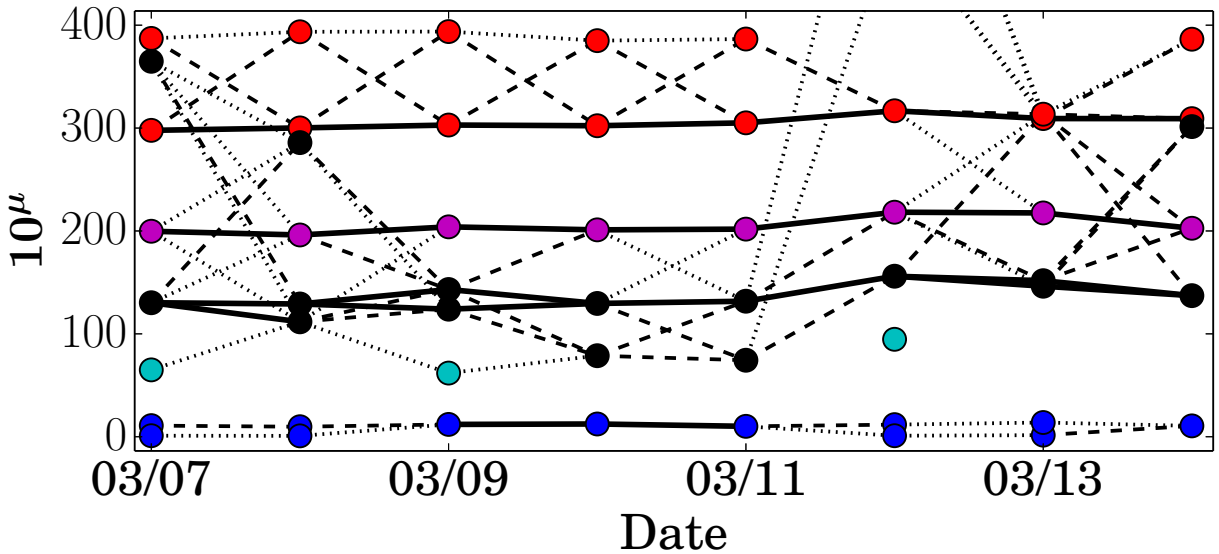

March 12 2011:

- 20ms RTT increase for all hosts outside of Japan
- RTT inside Japan are unchanged

$\rightarrow$ Impact of damaged trans-Pacific links and intra-AS route change

Application 2

Consistency check

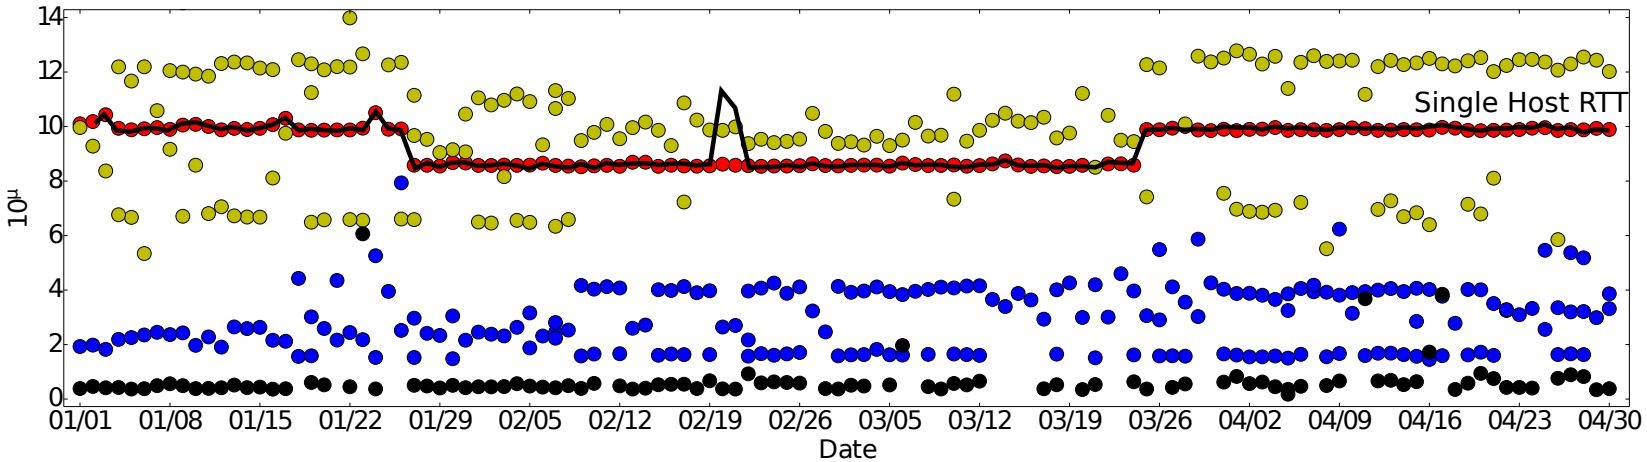

- Verify if a host is consistent with "its" cluster
- Compare the host RTTs with the identified RTT distributions
- Take into account the RTT variance
- Consistency score:
  - ≈ 1 means the host behaves like other hosts
  - ≈ 0 means the host deviates from other hosts

Consistency check: Examples

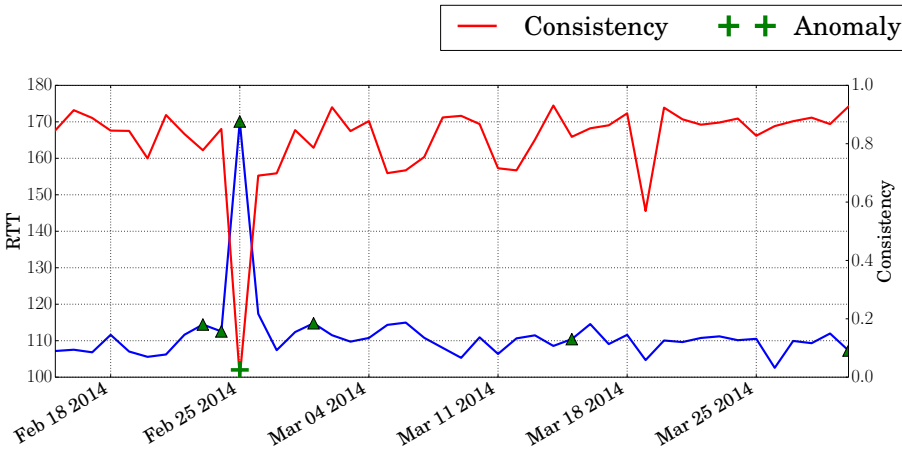

(c) TLD DNS server affected by route change

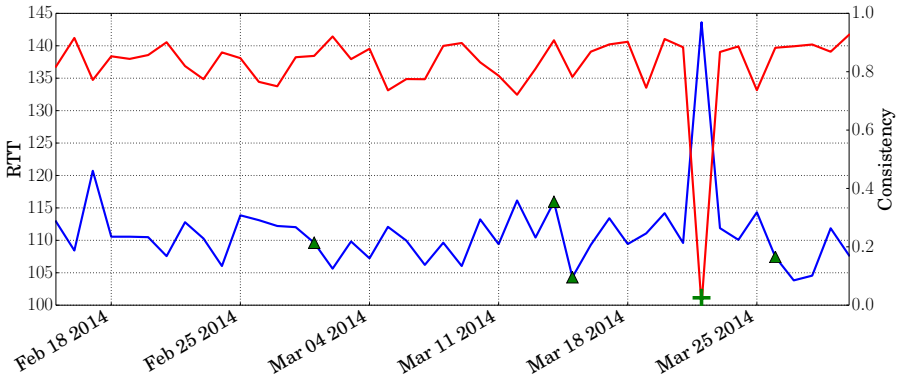

(d) Amazon host experiencing suspicious RTT peak

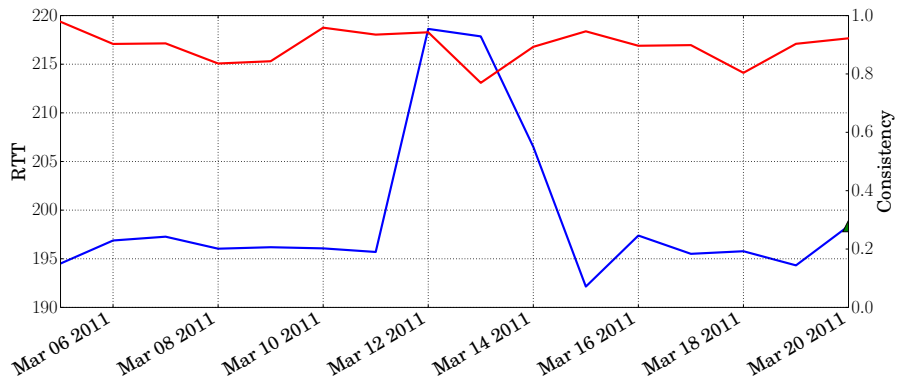

(e) RTT fluctuation during the Tohoku earthquake

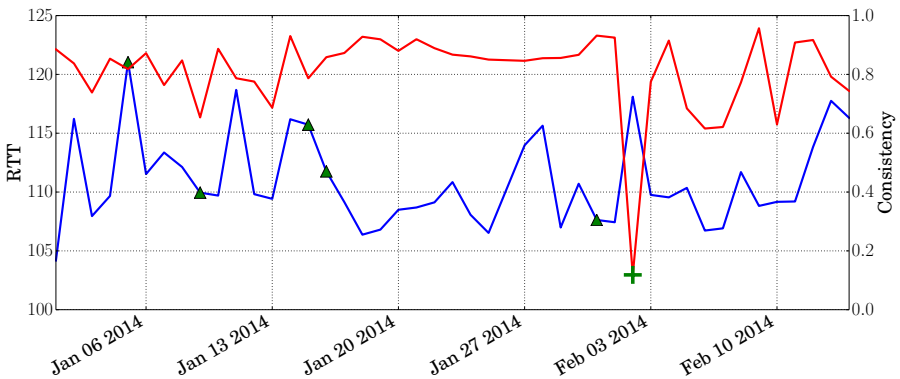

(f) Suspicious RTT fluctuations identified with the consistency check

→ Proposed model gives more insights than simple RTT analysis

Discussions

Implementation

- Low memory usage and computational complexity
- Suitable to sampled traffic
- Accuracy decreases with distance ($Tokyo \neq Osaka$, $FR = UK$)

Empirical Approach

- Monitor RTT experienced by Internet users
- Cluster IP based on RTT values (not AS)

Possible applications

- DDoS detection?
- BGP hijack?
- $\rightarrow$ Difficult to evaluate

Conclusions

Proposed mixture model

- Coarse view of numerous hosts RTTs
- Track typical RTTs time evolution
- Formalize RTT dynamics in a graph

Applications

- Provides insights into the delays experienced by a large population of IP hosts
- Reference to find hosts deviating from their population

R. Fontugne, J. Mazel, K. Fukuda. "An Empirical Mixture Model for Large-Scale RTT Measurements", INFOCOM 2015