

# **Lessons learnt from the** *Beijing Olympic Games* **Website Measurement**

Rocky K. C. Chang  
The Hong Kong Polytechnic University  
24 February 2009

# Beyond IPv4 exhaustion, IPv6, translation, ...

---

- Network measurement and monitoring
  - End-to-end **network path quality**
  - Not end-to-end application performance
- Measurement methodology
  - **Active** verses passive measurement
- Active measurement methods
  - **Non-cooperative** verses cooperative of the two end hosts
- Applications: SLA measurement, provider/peer selection, capacity planning, etc

# An unfinished business

---

- The Internet is no longer friendly to measurement probes.
  - Many unfriendly and intelligent middleboxes
  - Measurement Lab from Google, PlanetLab, ...
- Measurement results may not reflect the experience of data packets.
- Continuous monitoring for inter-domain paths is hard without receiving complaints.

# A new measurement method

---

- OneProbe
  - User-level tool
  - TCP data probes (not ICMP, TCP SYN/RST, UDP)
  - Round-trip delay, packet loss, packet reordering, and capacity.
  - Forward-path and reverse-path performance
  - Web service at the remote node
  - For both IPv4 and IPv6
- Tried out OneProbe on a number of websites hosting Beijing Olympic information.

# Outline

---

- Measurement methodology
- Measurement results
  - Diurnal RTT and round-trip loss patterns
  - Discrepancy between Ping RTT and OneProbe RTT
  - Highly symmetric loss rates
  - Impact of configuration changes
  - Some “unresolved” questions
- Conclusions

# Outline

---

- Measurement methodology
- Measurement results
  - Diurnal RTT and round-trip loss patterns
  - Discrepancy between Ping RTT and OneProbe RTT
  - Highly symmetric loss rates
  - Impact of configuration changes
  - Some “unresolved” questions
- Conclusions

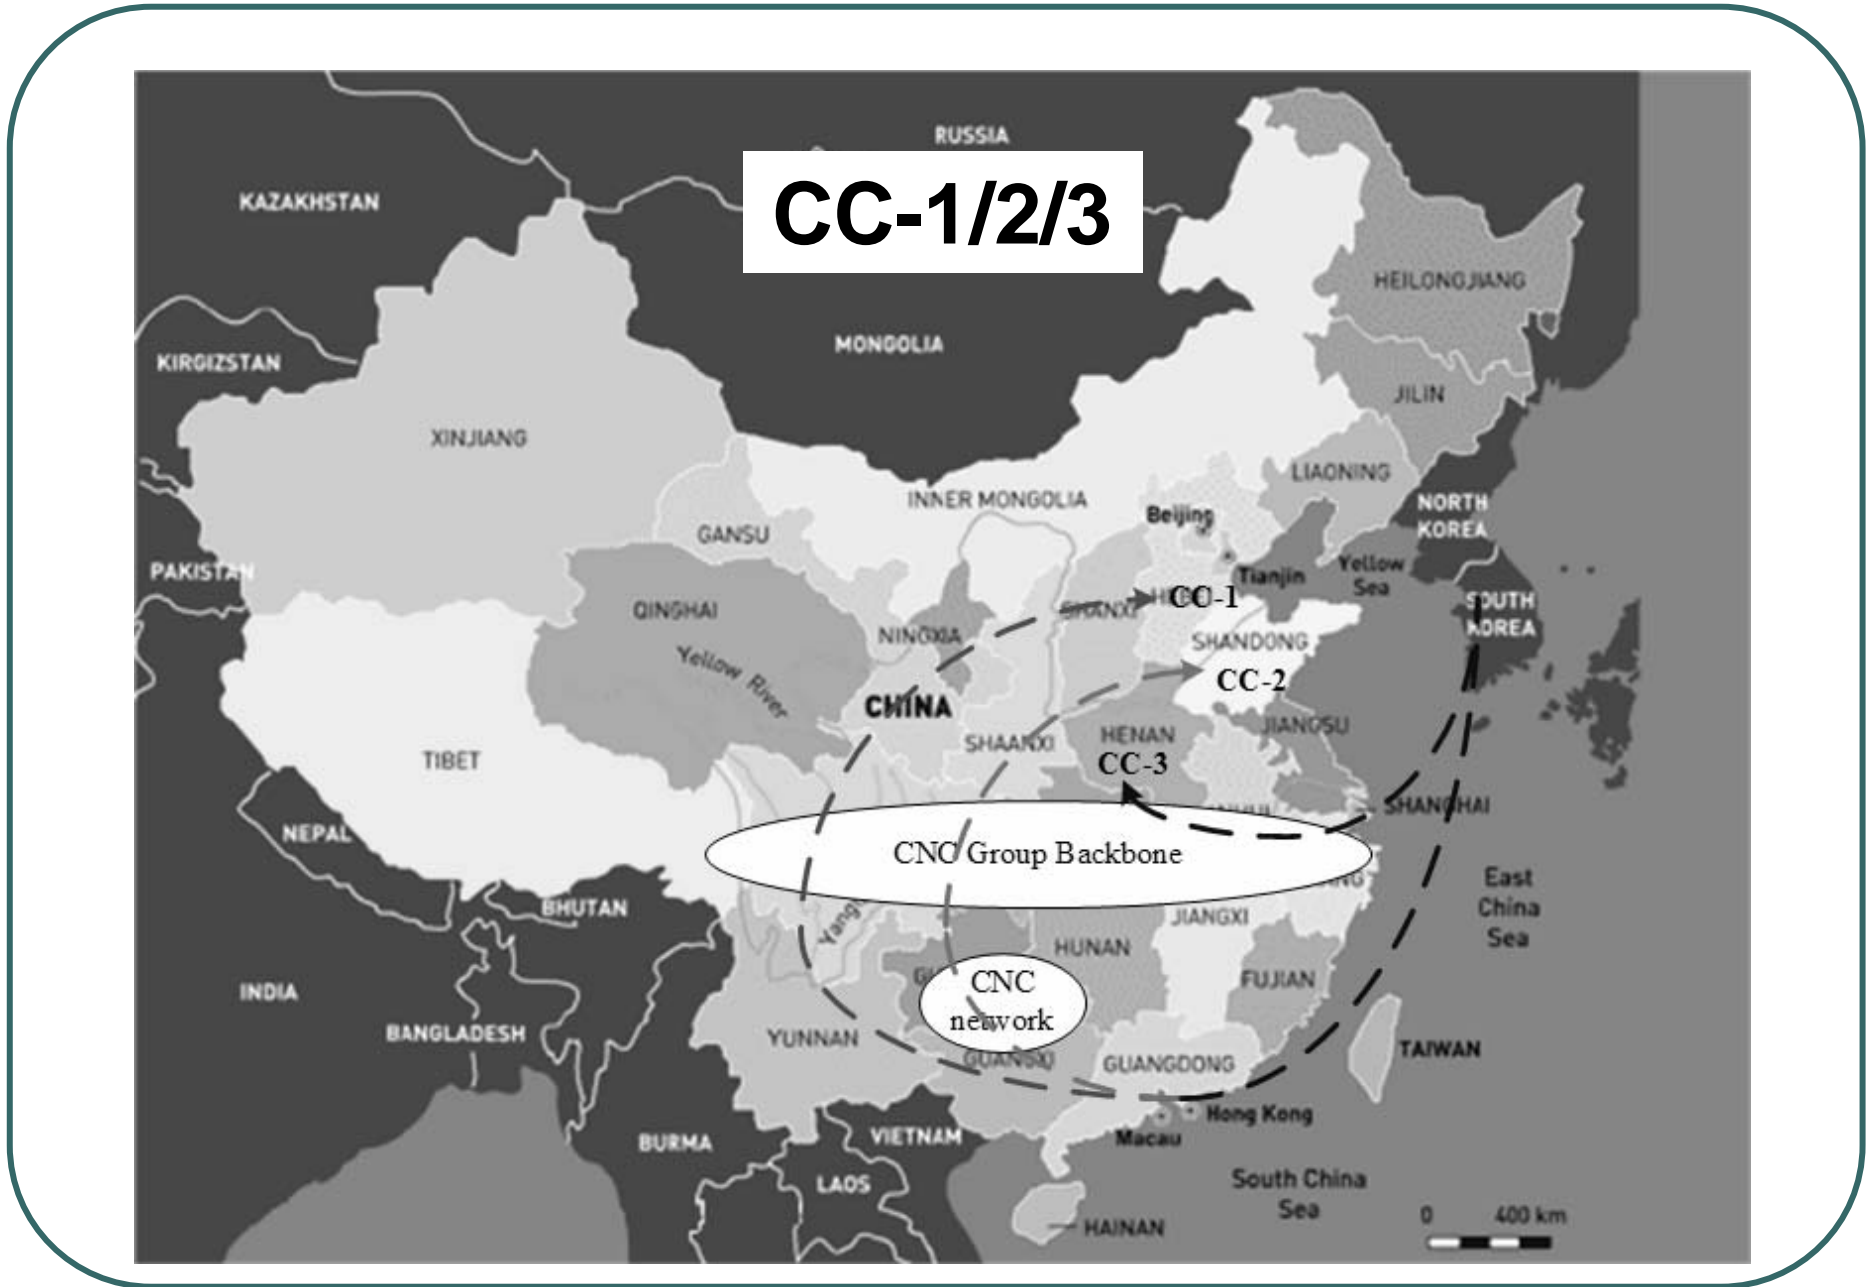
# Measurement methodology

---

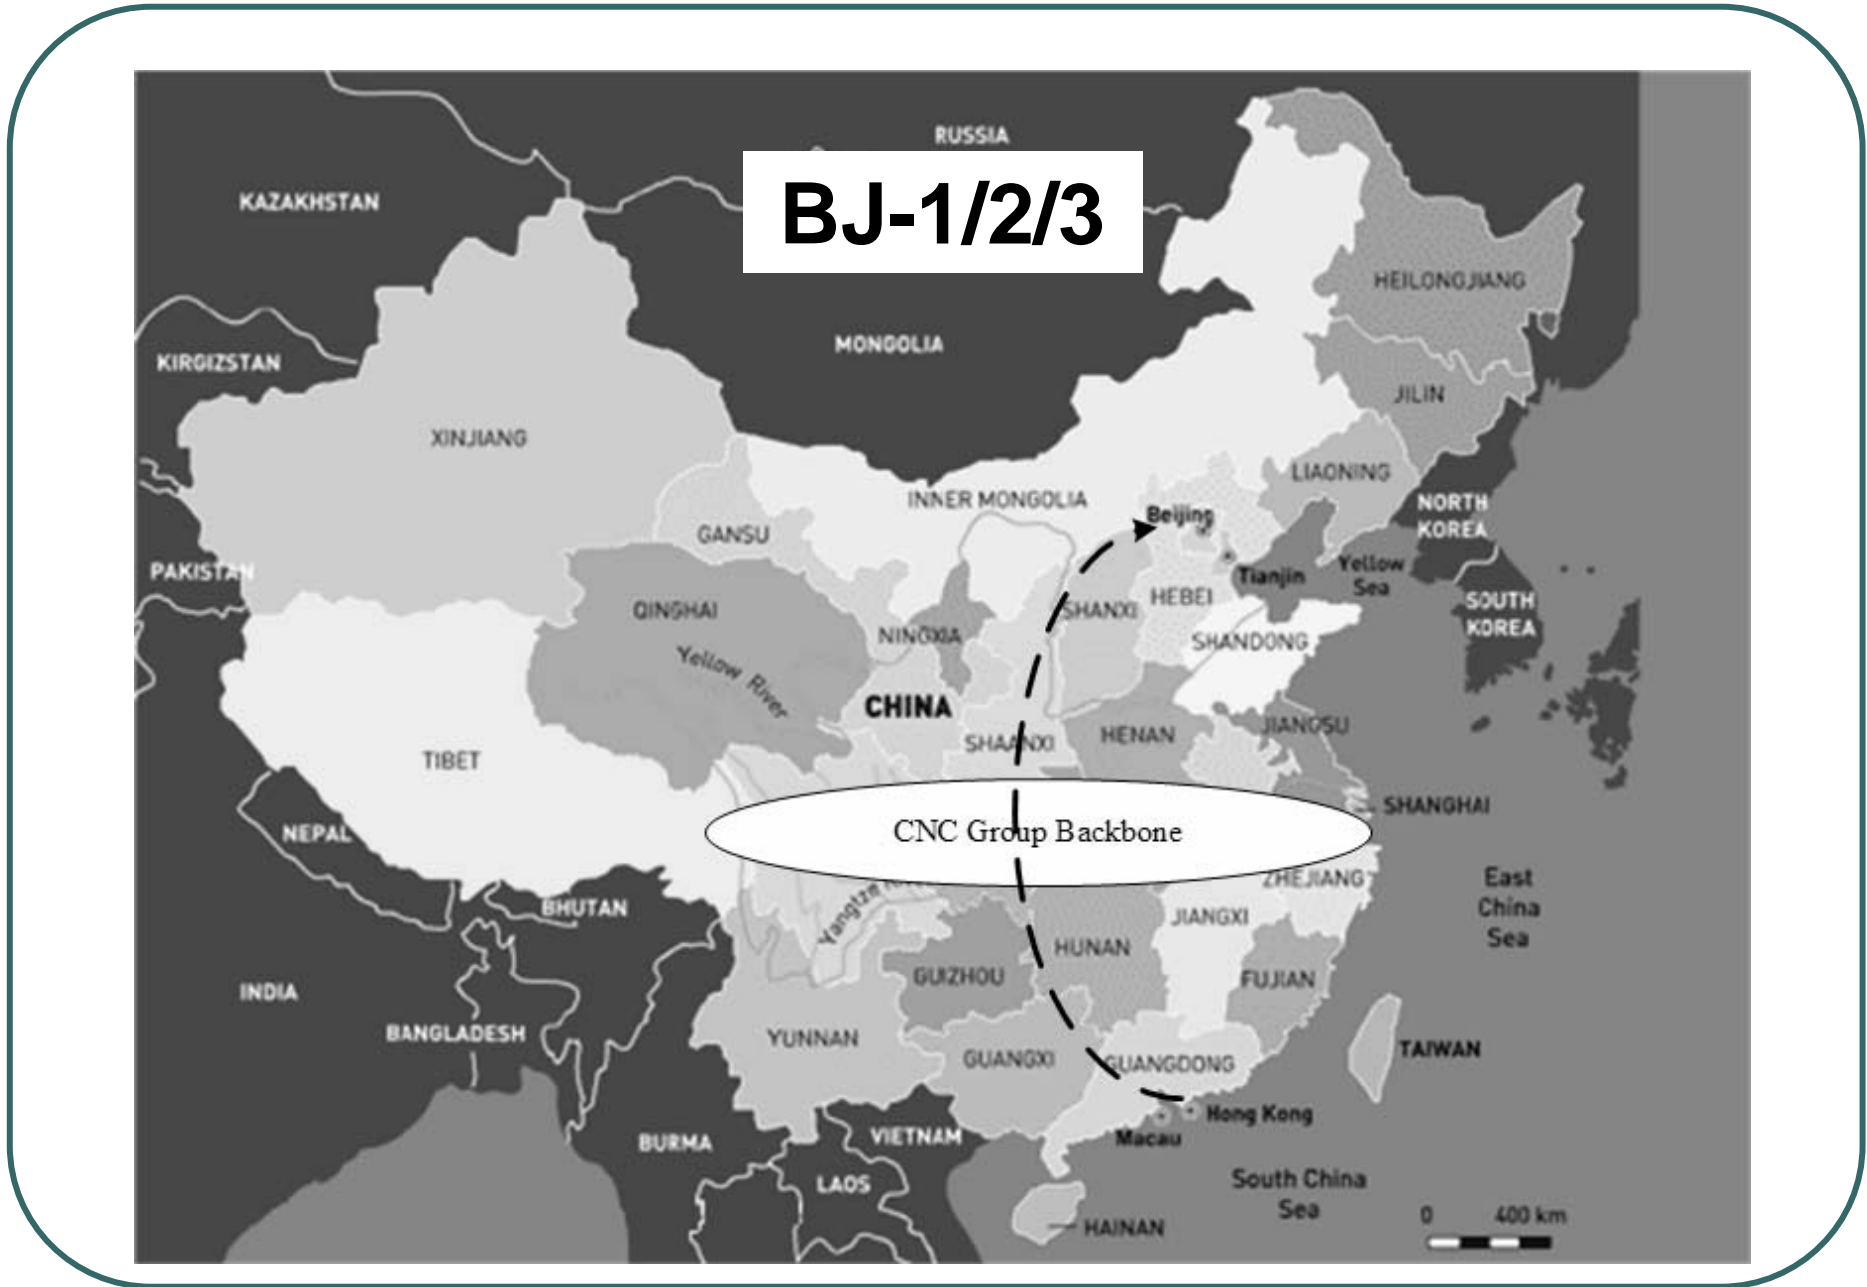
- Solo-active measurement tools
  - OneProbe (TCP data)
  - Ping (ICMP), Pping (TCP SYN-ACK), HTTPing (TCP SYN-ACK and HTTP data)
- Measuring point
  - A data center in Hong Kong
- Destinations
  - Three Beijing origin servers (BJ-1/2/3)
  - Three Chinacache sites (CC-1/2/3)
  - Quest (QT) and Akamai (AK-1/2)

	Probers/ servers	No. hops	Forward path: Region/Network (no. hops)
H-1	CC-1	13	HK (5)→CNCGroup Backbone (4)→Hebei Province Network (4)
	BJ-1	16	HK (5)→AP-TELEGLOBE (2)→CNCGroup Backbone (4)→Beijing Province Network (4)
	QT	5	All in HK
H-2	CC-2	19	HK (5)→CNCGroup Network (6)→CNCGroup Backbone (3)→Shandong Province Network (5)
	BJ-2	16	Almost the same as BJ-1's
	AK-1	5	First four hops same as QT's
H-3	CC-3	16	HK (5)→Korea (2)→CNCGroup Backbone (4)→Henan Province Network (5)
	BJ-3	16	Almost the same as BJ-1's
	AK-2	5	First four hops same as QT's





# CC-1/2/3



# Measurement parameters

---

- Measurement period
  - During the entire Olympic Games period and one week after the Games
- Sampling rate
  - OneProbe: sending a probe every 0.5 secs
  - Others: sending a probe every sec.
- Probe packet size
  - OneProbe: 1500 bytes
  - Ping: 1500 bytes
  - Pping: 40 bytes
  - HTTPping: variable

# Path quality metrics

---

- Ping, Pping, and HTTPing
  - Round-trip time (RTT)
  - round-trip loss rate
- OneProbe
  - RTT
  - One-way loss rates
  - One-way reordering rates

# Outline

---

- Measurement methodology
- **Measurement results**
  - Diurnal RTT and round-trip loss patterns
  - Discrepancy between Ping RTT and OneProbe RTT
  - Highly symmetric loss rates
  - Impact of configuration changes
  - Some “unresolved” questions
- Conclusions

# Overall path quality

---


- Chinacache:
  - RTT: 87ms-260ms
  - Loss rate:  $\leq 10\%$
- Beijing
  - RTT: 78ms-132ms
  - Loss rate:  $\leq 20\%$
- Akamai/Quest
  - RTT: 3.4ms-4.9ms
  - Loss rate:  $\leq 10\%$
- No reordering events observed
- Mostly stable forward-path routes and reverse-path hop counts
  - Detected configuration changes
  - Some persistent load-balancing

# I. Diurnal RTT and round-trip loss patterns

---


- RTT and round-trip loss rate from OneProbe
  - No complaints received!
- Diurnal patterns for both RTT and loss rates
  - Positive correlation
  - More losses on weekends
  - No diurnal patterns observed from the AK/QT sites
- All experienced a sudden increase of 17 ms at 23 August 20:37 UTC.

# The BJ-1 path







# CC-1/2/3



(a) CC-1




(c) CC-2




(e) CC-3


# BJ-1/2/3



(b) BJ-1



(d) BJ-2




(f) BJ-3

## **II. Discrepancy between Ping RTT and OneProbe RTT for CC-3**


---

- Differed by around 100 ms during the peaks for the first 4 days.
- They were similar in the valleys.
- Their RTTs “converged” at 12 Aug. 2008 16:39 UTC (~1.5 hrs into the midnight).
- Discrepancy detected even after the convergence point.

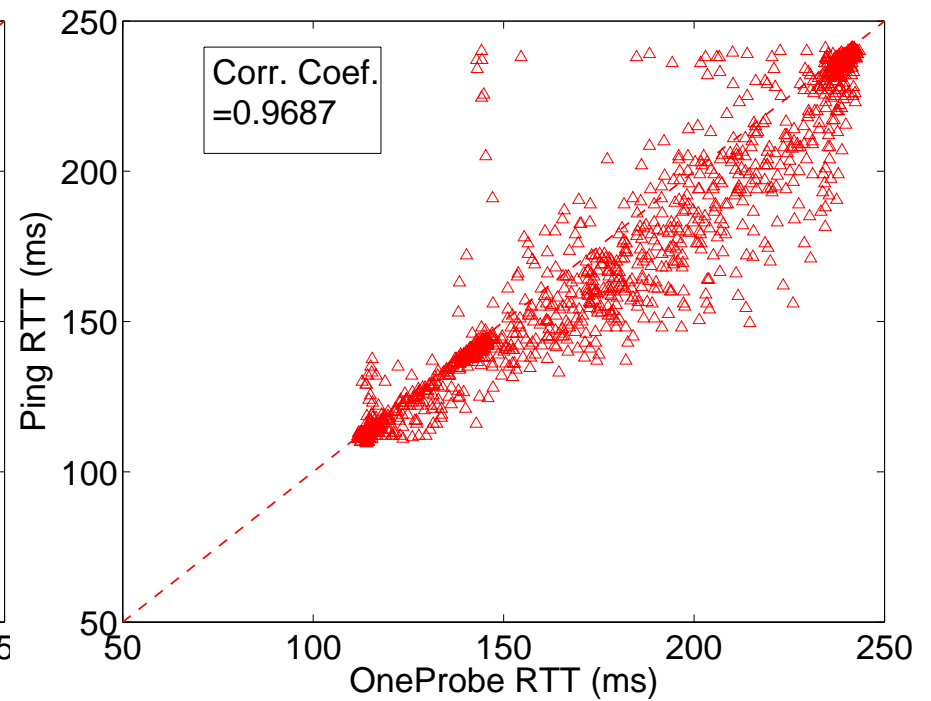
# CC-3



# One-week data for CC-2/3



**CC-2**




**CC-3**

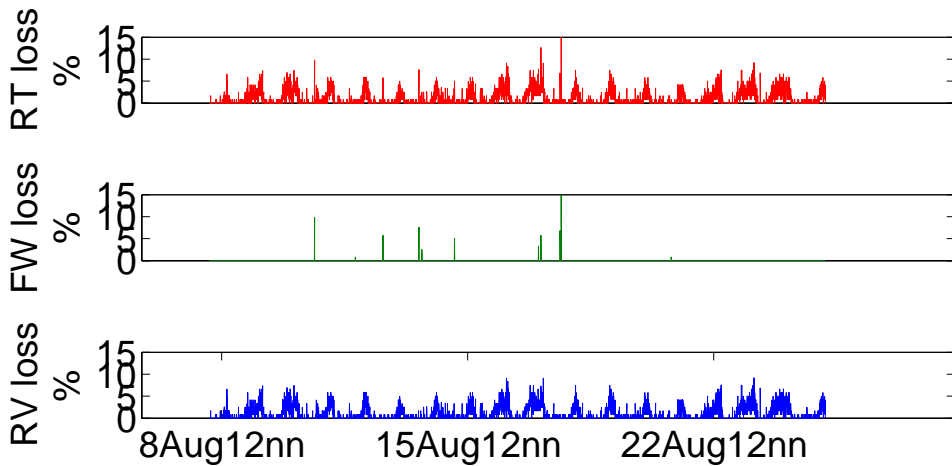
### **III. Highly symmetric loss rates**

---

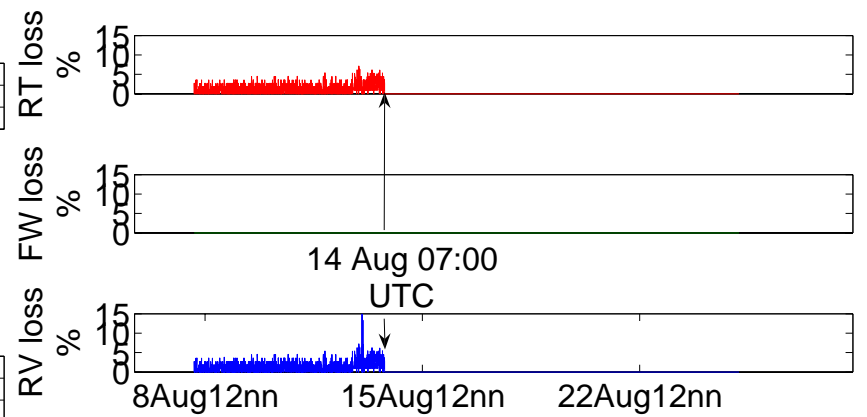
- The round-trip loss rates were mostly dominated by reverse-path losses.
  - Client-server traffic
- CC-3: the forward-path loss rate diminished significantly at 12 Aug. 2008 16:39 UTC.
- QT: the reverse-path losses were persistently high until 14 August 2008 06:50 UTC.



**CC-3**




**BJ-1**



**QT**

# IV. Impact of configuration changes



**AK-2**

## **Some “unresolved” questions**

---



- What (and who) was/were responsible for the diurnal RTT patterns?
- What (and who) was/were responsible for the diurnal loss patterns?
- What caused discrepancy between ICMP measurement and TCP data measurement?



## **What was/were responsible for the diurnal RTT patterns?**


---

- Performed another set of experiments.
  - Same forward path
  - Different reverse path (a different upstream provider)
- This set does not exhibit the diurnal RTT patterns for the CC and BJ paths.
  - What do it mean?



**CC-1**

**QT**



**BJ-1**

# Outline

---

- Measurement methodology
- Measurement results
  - Diurnal RTT and round-trip loss patterns
  - Discrepancy between Ping RTT and OneProbe RTT
  - Highly symmetric loss rates
  - Impact of configuration changes
  - Some “unresolved” questions
- **Conclusions**

# Conclusions

---

- Inaccuracy of the current non-cooperative, active measurement tools.
  - Different forwarding paths
- Round-trip measurement is far from being sufficient.
  - Asymmetric loss rates
- More “sophisticated” methods to reveal the root causes.
  - Network tomography

# Acknowledgments

---

- The experiments and analyses were conducted by the following members in the group:
  - Edmond Chan
  - Waiting Fok
  - Daniel Luo
- This project is partially supported by a grant from the Innovation Technology Fund in Hong Kong.

Questions?