Measurement of QoE in Tsinghua university

CERNET Center
Qianli Zhang
zhang@cernet.edu.cn

Background

- Internet has evolved as a complex set of independently developed and deployed protocols, technologies, and core applications
  - From simple content to multimedia content
  - From direct retrieve to complex CDN or P2P
  - From independent applications to everything over HTTP
- For ISPs, it is harder to know the end users' experience

Background(2)
Complicated performance analysis

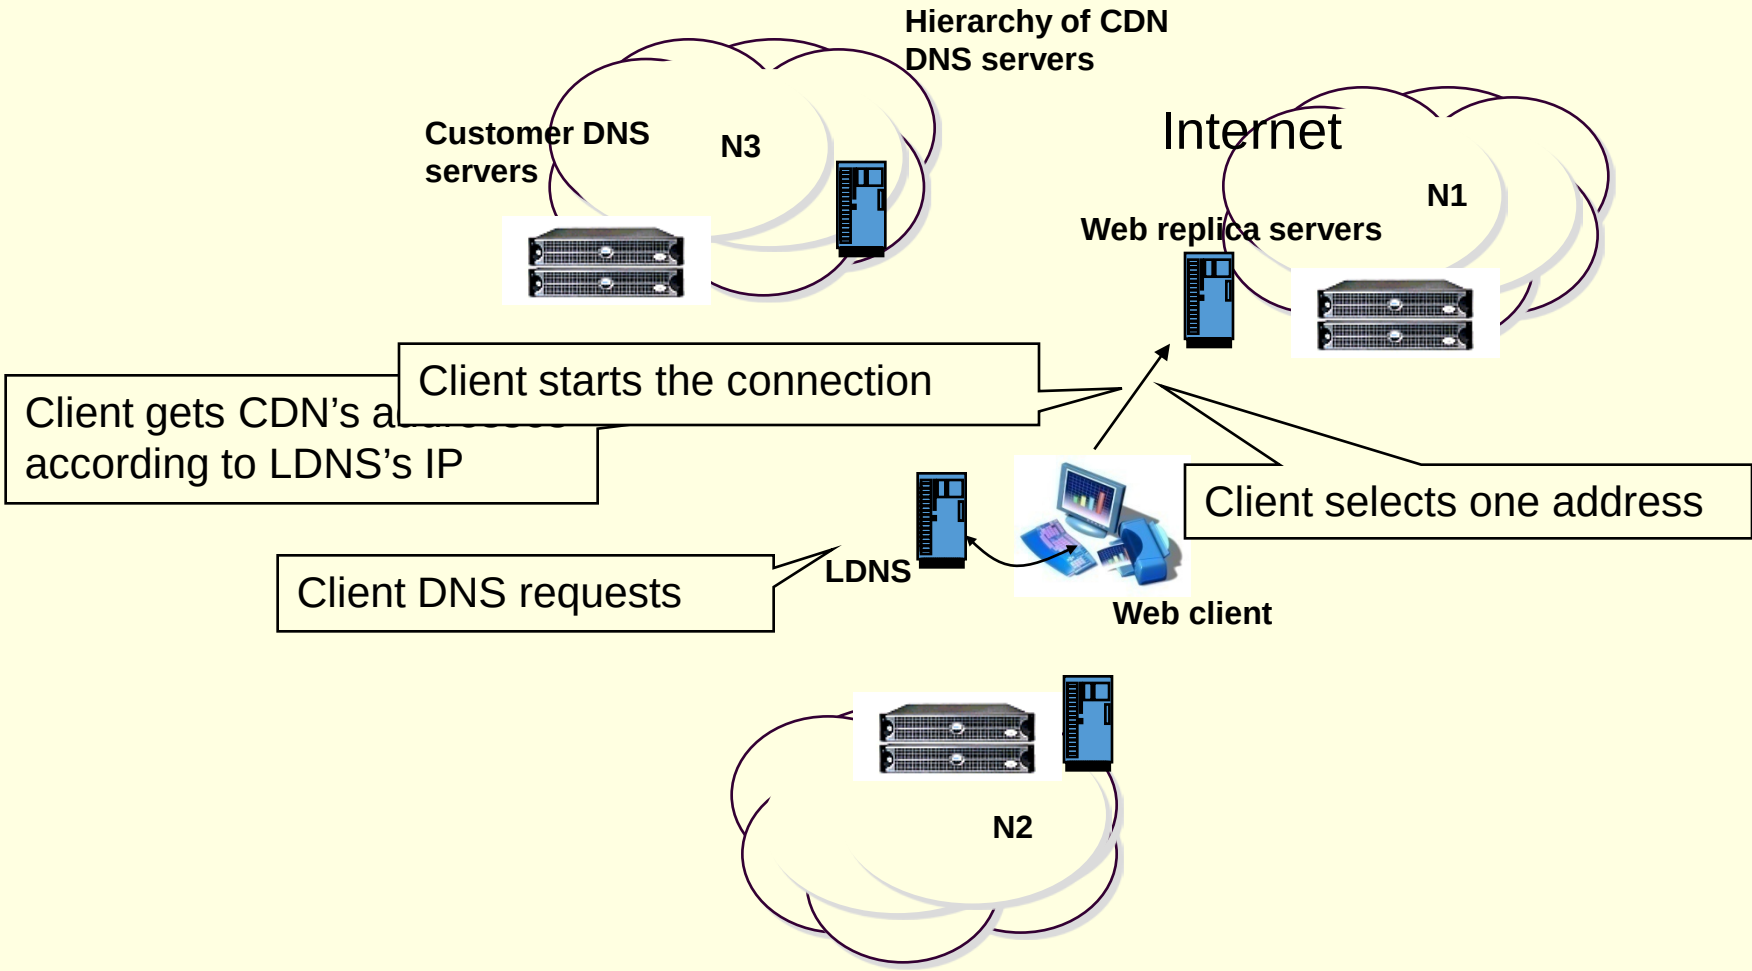

Background(3)
How to evaluate the QoE

- Define QoE
  - There is many factors to evaluate the QoE
    - Latency、Throughput、Response time、Arrival rate、Utilization、Bandwidth、Loss、Stability
  - Measuring Internet performance is easier said than done
  - Decide the one according to your motivation
- Measurement is necessary
  - Motivation is to increase the users' QoE
  - How, when, where, what do we measure?
- Measurement is meaningless without careful analysis
  - Analysis is necessary to derive useful suggestions and findings
  - Analysis of data gathered from networks is often different among networks

Background(4) Measurement

- Active measurement
  - Many tools have been developed to measure/monitor general characteristics of network performance
    - *traceroute* and *ping*
  - Problems:
    - Measurement may not accurately reflect the user's expierence
    - Measurement will send packets and thus add extra load to network
- Passive measurement
  - Will not use extra network resources
    - Packet capture applications (tcpdump) uses packet capture filter (bpf,libpcap)
    - Flow-based measurement tools
    - SNMP tools
  - Problems
    - Privacy issues
    - Very hard to refactor to original user's activity
    - Getting packet scoped in backbone of the network

Background(5)
traceroute and ping

- Traceroute

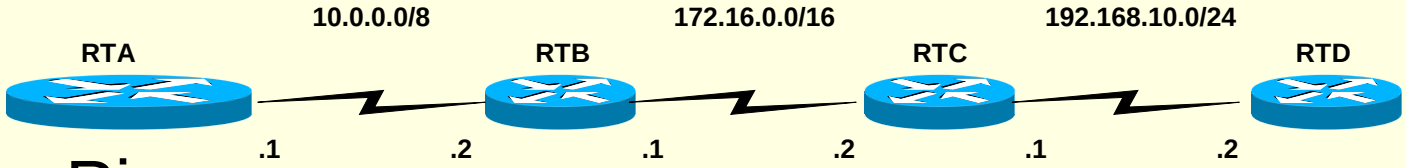

- Ping
  - ICMP ECHO request and ECHO reply

```
root@debian:~# ping 202.112.3.65
PING 202.112.3.65 (202.112.3.65) 56(84) bytes of data.
64 bytes from 202.112.3.65: icmp_req=1 ttl=62 time=0.463 ms
…
^C
--- 202.112.3.65 ping statistics ---
7 packets transmitted, 7 received, 0% packet loss, time 5999ms
rtt min/avg/max/mdev = 0.437/0.470/0.495/0.034 ms
```

Background(6)
Tsinghua University Campus Network

- Tsinghua University campus network is a very large network
  - Six B class addresses
  - Total user number 172,000
  - Active user number 60,000
  - Peak online user number 42,000

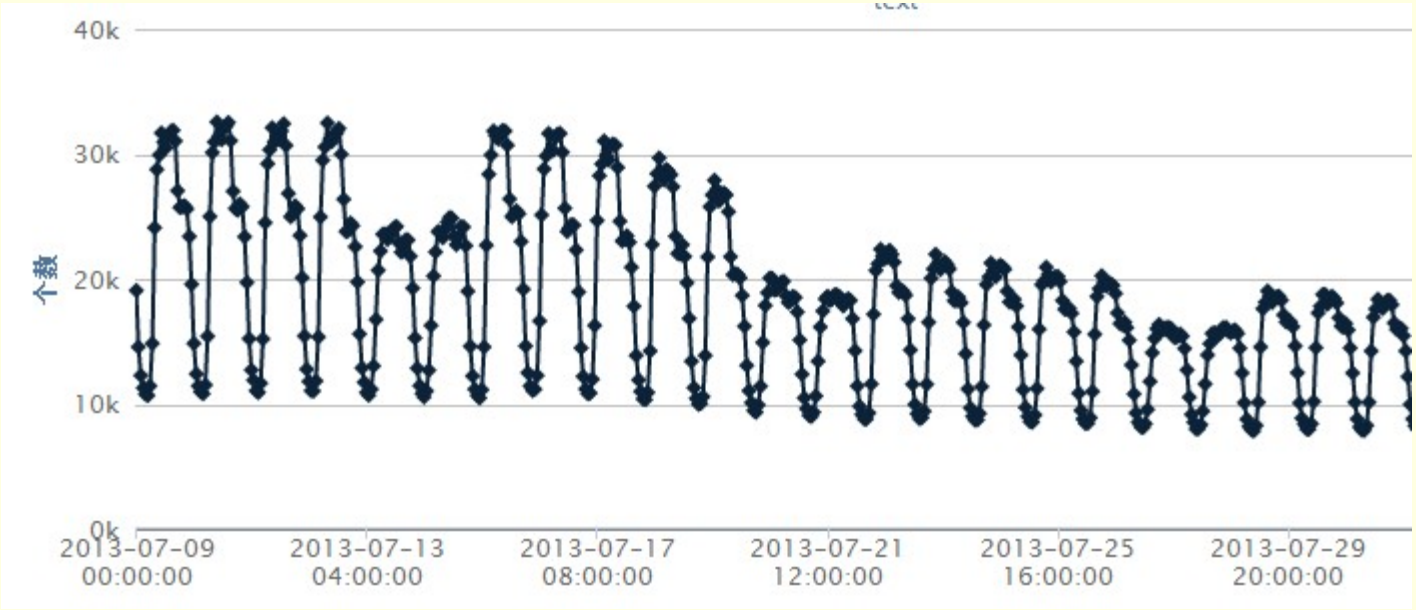

Our practice

- Motivation
  - Evaluate the users' experience in Tsinghua University campus network
  - Find the possible problems in network
- Define the QoE
  - Use the packet loss and RTT of ping
- Analysis
  - Find the reasons of bad experience in some sites
  - Find the solutions to increase users' experience

Active measurement(1): Design

- Define the QoE
  - Packet loss
  - RTT
- Destination selection
  - Alexa top 500 global sites
  - Alexa top 500 sites in China
- Measurement point
  - In tsinghua university campus network, with the common configuration

Active measurement(2): Architecture

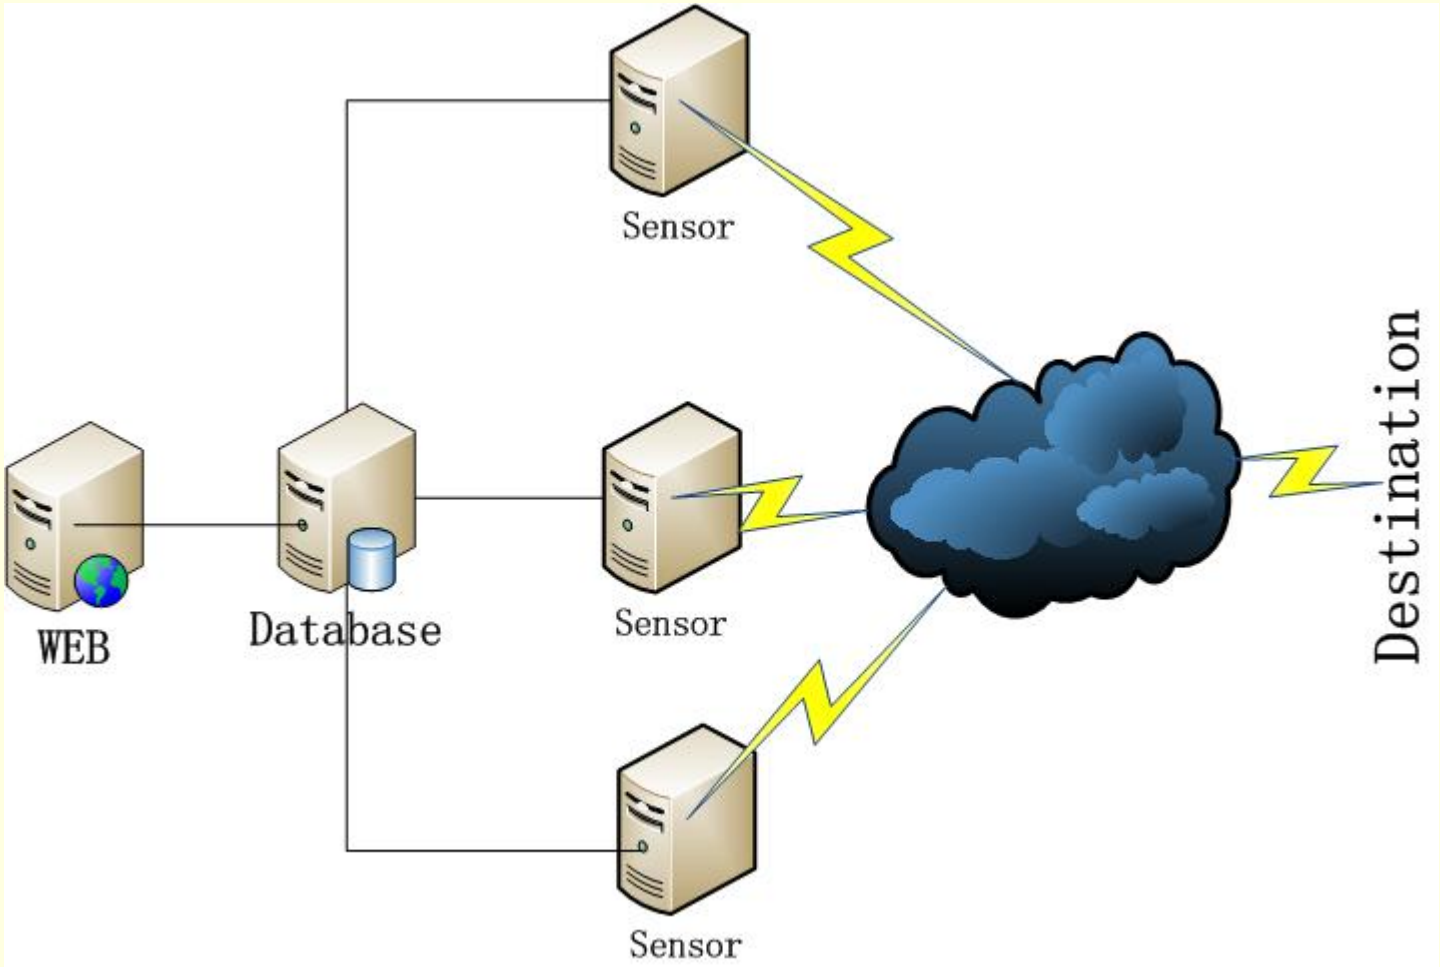

Active measurement(2): Implementation

- Sensor
  - Each 30 minutes, all destinations will be pinged
    - Avoid the congestion caused by too many simultaneous ping packets
  - Each sensor will insert the result into databases
- Database
  - PostgreSQL
- Web
  - Resin based application server
  - Datatable + highchart

Active measurement(3): Implementation(cont)

- In each 30 minutes period, the destination will be grouped to avoid sudden increase of network load
- Incoming

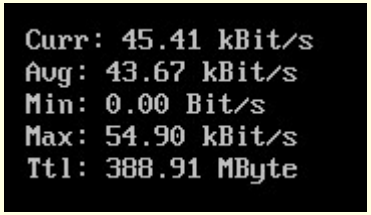

- Outgoing

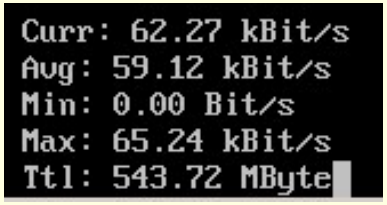

Snapshot(1)
First page

Snapshot(2)

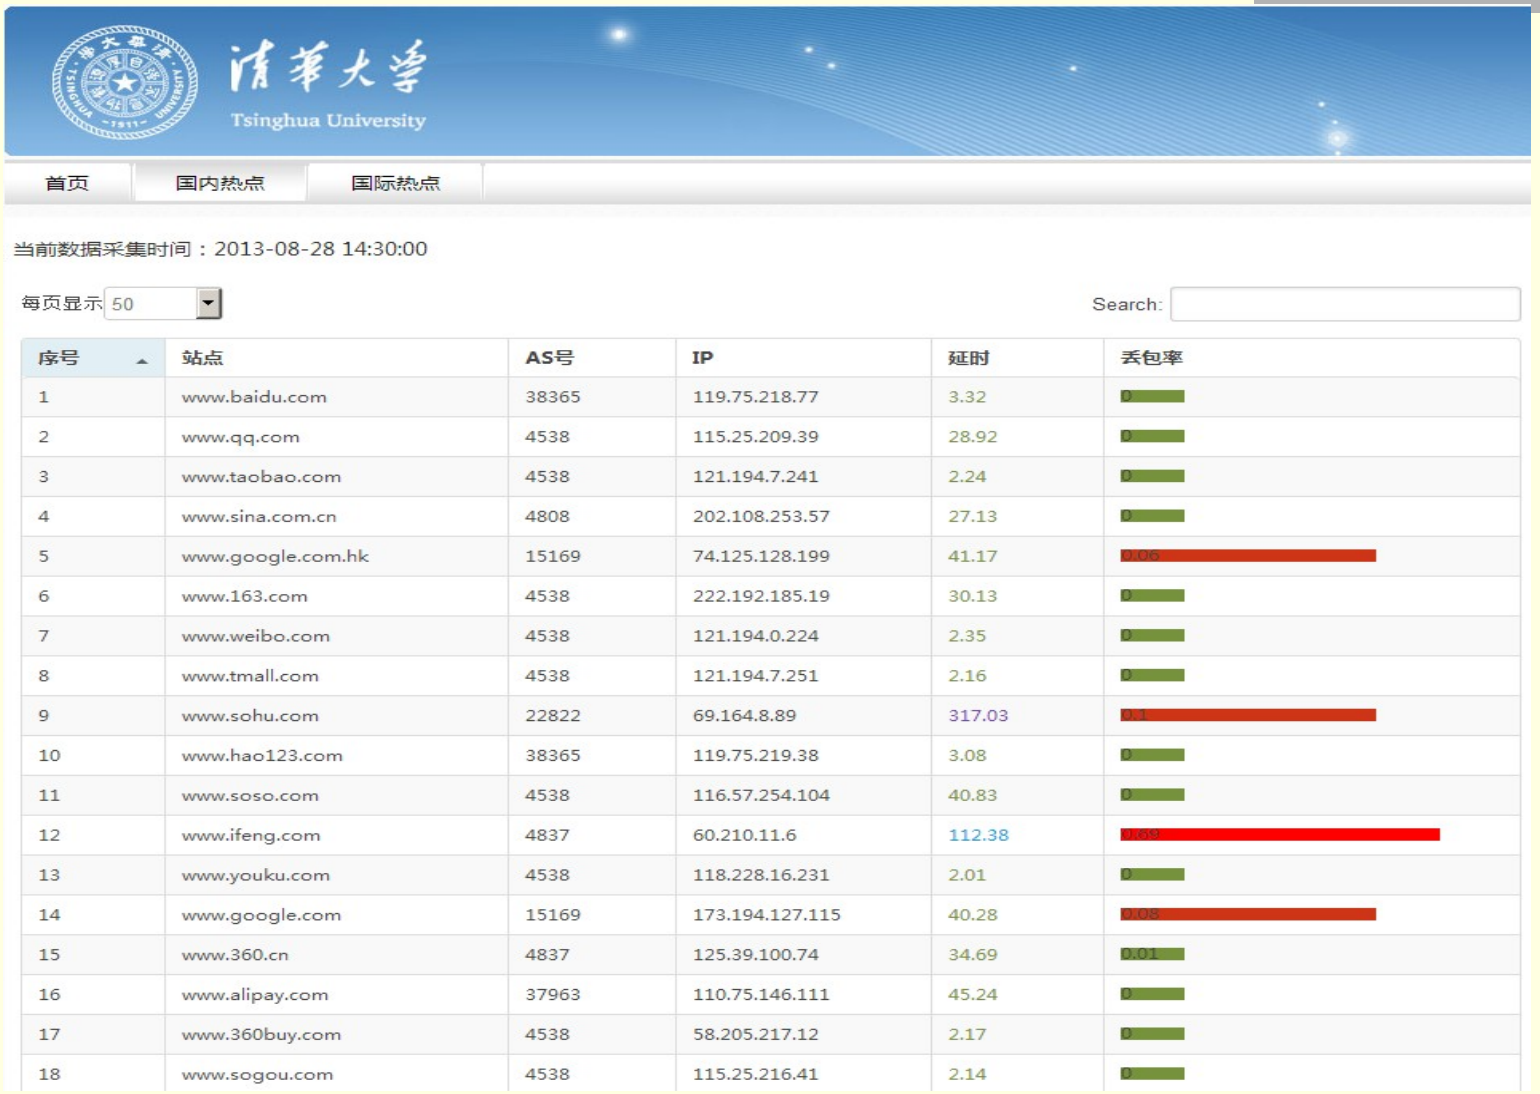

序号	站点	AS号	IP	延时	丢包率
1	www.baidu.com	38365	119.75.218.77	3.32	0
2	www.qq.com	4538	115.25.209.39	28.92	0
3	www.taobao.com	4538	121.194.7.241	2.24	0
4	www.sina.com.cn	4808	202.108.253.57	27.13	0
5	www.google.com.hk	15169	74.125.128.199	41.17	0.06
6	www.163.com	4538	222.192.185.19	30.13	0
7	www.weibo.com	4538	121.194.0.224	2.35	0
8	www.tmall.com	4538	121.194.7.251	2.16	0
9	www.sohu.com	22822	69.164.8.89	317.03	0.1
10	www.hao123.com	38365	119.75.219.38	3.08	0
11	www.soso.com	4538	116.57.254.104	40.83	0
12	www.ifeng.com	4837	60.210.11.6	112.38	0.69
13	www.youku.com	4538	118.228.16.231	2.01	0
14	www.google.com	15169	173.194.127.115	40.28	0.09
15	www.360.cn	4837	125.39.100.74	34.69	0.01
16	www.alipay.com	37963	110.75.146.111	45.24	0
17	www.360buy.com	4538	58.205.217.12	2.17	0
18	www.sogou.com	4538	115.25.216.41	2.14	0

Snapshot(3)


首页
国内热点
国际热点
子网视图
数据采集时间：2013-09-02 数据采集目标：www.google.com.hk
日数据
周数据
月数据
延时曲线
2013-09-02
MS（毫秒）
125
100
75
50
25
延时
2013-08-03
2013-08-04
2013-08-05
2013-08-06
2013-08-07
2013-08-08
2013-08-09
2013-08-10
2013-08-11
2013-08-12
2013-08-13
2013-08-14
2013-08-15
2013-08-16
2013-08-17
2013-08-18
2013-08-19
2013-08-20
2013-08-21
2013-08-22
2013-08-23
2013-08-24
2013-08-25
2013-08-26
2013-08-27
2013-08-28
2013-08-29
2013-08-30
2013-08-31
2013-09-02
丢包曲线
2013-09-02
百分比
0.4
0.3
0.2
0.1
0
-0.1
丢包

Findings#1:
Why these sites are so slow?

83	www.ce.cn	4837	60.210.11.6	112.38	
102	www.yoka.com	4837	60.210.11.6	112.38	
107	www.pcgames.com.cn	4837	60.210.11.6	112.38	
116	www.pclady.com.cn	4837	60.210.11.6	112.38	
129	www.315che.com	4837	60.210.11.6	112.38	
153	www.qianyan001.com	4837	60.210.11.6	112.38	
174	www.weilai.net	4837	60.210.11.6	112.38	
194	www.iche.com	4837	60.210.11.6	112.38	
202	www.discuz.net	17621	112.64.234.169	31.86	
203	www.ynet.com	4837	60.210.11.6	112.38	
215	www.chexun.com	4837	60.210.11.6	112.38	
255	www.ip138.com	4837	60.210.11.6	112.38	
289	www.imanhua.com	4837	60.210.11.6	112.38	
313	www.zjol.com.cn	4837	60.210.11.6	112.38	
346	www.hsw.cn	4837	60.210.11.6	112.38	
348	www.066c.com	4837	60.210.11.6	112.38	
362	www.1616.net	4837	60.210.11.6	112.38	
368	www.bankrate.com.cn	4837	60.210.11.6	112.38	
483	www.discuz.com	17621	112.64.234.169	31.86	
487	www.yokamen.cn	4837	60.210.11.6	112.38	
430	www.doubleclick.com	15169	74.125.235.193	159.09	
148	www.seowhy.com	17672	219.232.254.14	16.8	

Findings#1(cont)
ping different addresses of www.ce.cn

```
root@fping:~#  fping -i 10 -C 100 -q -e
      222.211.64.112 60.210.11.6
222.211.64.112 : 35.70 35.78 50.17 35.87
      35.68 35.66 35.50 35.66 40.08 35.62
      35.72 35.54 35.91 35.92 35.77 35.93
      36.12 37.21 36.22 35.80 35.93 35.48
      36.01 36.01 35.86 35.85 35.80 36.74
      35.76 35.85 36.27 35.79 35.62 35.48
      36.80 36.95 35.60 35.88 35.53 35.79
      35.36 35.52 35.48 35.77 35.79 35.83
      35.38 36.21 35.58 35.32 35.76 35.83
      35.53 35.56 36.77 36.21 36.64 36.04
      35.51 35.54 36.16 35.52 35.76 35.81
      35.96 35.51 35.47 35.81 35.44 35.94
      35.63 36.27 35.78 35.96 35.74 35.50
      35.50 35.78 36.05 35.80 35.54 35.78
      35.44 35.41 35.41 35.61 35.60 35.54
      35.40 35.62 35.56 35.32 35.52 35.56
      35.82 35.57 36.03 35.81 36.79 36.12
60.210.11.6    : 114.02 - - 113.46 - - 115.08 - -
      115.06 - - 114.78 - - - - - - - - - 115.30 - -
      114.52 - - 113.84 - - - - - 115.82 - -
      115.06 - - 114.89 - - 115.05 - - 114.92 - -
      113.67 - - 115.34 - - - - - - - - - - 114.13 -
      - 114.41 - - - - 115.05 - - 113.24 - - - - -
      113.07 113.41 - 114.82 - - 114.21 - -
      115.21 - - - - - - - 114.38
```

www.ce.cn	38283(4)
www.yoka.com	38283
www.pcgames.com.cn	4134
www.pclady.com.cn	4134(2)
www.315che.com	38283(4)
www.qianyan001.com	38283
www.weilai.net	4134(3)
www.iche.com	4134(3)
www.discuz.net	4812
www.ynet.com	4134(3)
www.chexun.com	4134(3)
www.ip138.com	38283(4)
www.imanhua.com	38283(2)
www.zjol.com.cn	4134(3)
www.hsw.cn	4134(3)

Findings #1 Summary
CDN's ignorance

- Most of top 500 sites use commercial CDN
- Usually users in Beijing prefer China unicom
- However, Tsinghua University campus network is connected to ChinaNET(4134, 38283…), while CDNs prefer China Unicom (4837)
- Solution
  - Update CDN's database.
  - Setup a new DNS.
- Summary
  - It is not always right to rely on CDN, since CDN do not have detailed information of all the users

Finding #2
Multiple DNS results from one AS(4837)

```
root@fping:~# dig @166.111.8.28 stockapp.finance.qq.com A

; <<>> DiG 9.7.3 <<>> @166.111.8.28 stockapp.finance.qq.com A
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 2615
;; flags: qr rd ra; QUERY: 1, ANSWER: 3, AUTHORITY: 2, ADDITIONAL: 3

;; QUESTION SECTION:
;stockapp.finance.qq.com.       IN      A

;; ANSWER SECTION:
stockapp.finance.qq.com. 12     IN      A       125.39.127.39
stockapp.finance.qq.com. 12     IN      A       125.39.127.59
stockapp.finance.qq.com. 12     IN      A       111.161.48.41
```

Finding #2(cont)

```
root@fping:/bin# fping -i 10 -C 100  111.161.48.41 125.39.127.39
111.161.48.41 : [0], 96 bytes, 109 ms (109 avg, 0% loss)
111.161.48.41 : [1], 96 bytes, 111 ms (110 avg, 0% loss)
125.39.127.39 : [1], 96 bytes, 127 ms (127 avg, 50% loss)
111.161.48.41 : [2], 96 bytes, 111 ms (110 avg, 0% loss)
111.161.48.41 : [3], 96 bytes, 110 ms (110 avg, 0% loss)
111.161.48.41 : [4], 96 bytes, 111 ms (110 avg, 0% loss)
125.39.127.39 : [4], 96 bytes, 127 ms (127 avg, 60% loss)
111.161.48.41 : [5], 96 bytes, 110 ms (110 avg, 0% loss)
111.161.48.41 : [6], 96 bytes, 111 ms (110 avg, 0% loss)
111.161.48.41 : [7], 96 bytes, 108 ms (110 avg, 0% loss)
125.39.127.39 : [7], 96 bytes, 128 ms (128 avg, 62% loss)
^C
111.161.48.41 : 109.29 111.56 111.52 110.79 111.02 110.33 111.73 108.00
107.98 109.28 110.23 109.79 109.95 110.89
125.39.127.39 : - 127.84 - - 127.90 - - 128.91 - - - - - -
```

Finding #2 Summary: DNS selection among multiple results

- DNS round robin
  - Local DNS server may reorder the sequence of querying results.
  - Usually the end host will use the first item.
    - Chrome
- Problem
  - The first result may be very slow
  - The end users have to wait for the local DNS TTL reached(usually several seconds)
- Solution
  - Local DNS manipulate the DNS records to reflect speed (a little hard)
  - More intelligent clients(browsers) are required.

Finding #3: Performance jitters

- DNS change from one to another
- Sudden network peak

www.ifeng.com(2013-08-30) DNS jitter

Diurnal jitters(www.114so.cn, www.kimiss.com, www.mumayi.com)

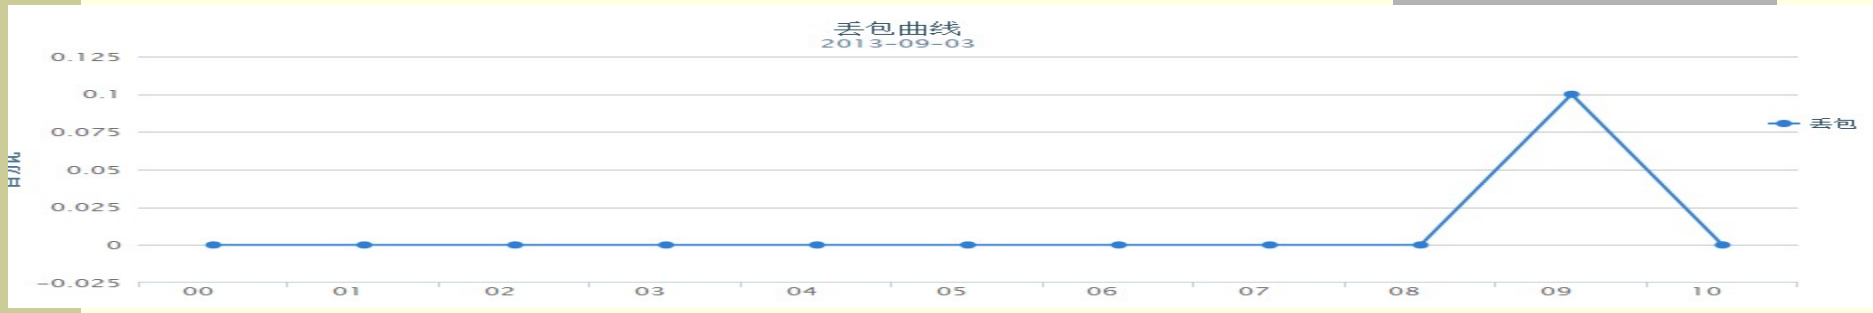

Finding #4
www.adnxs.com

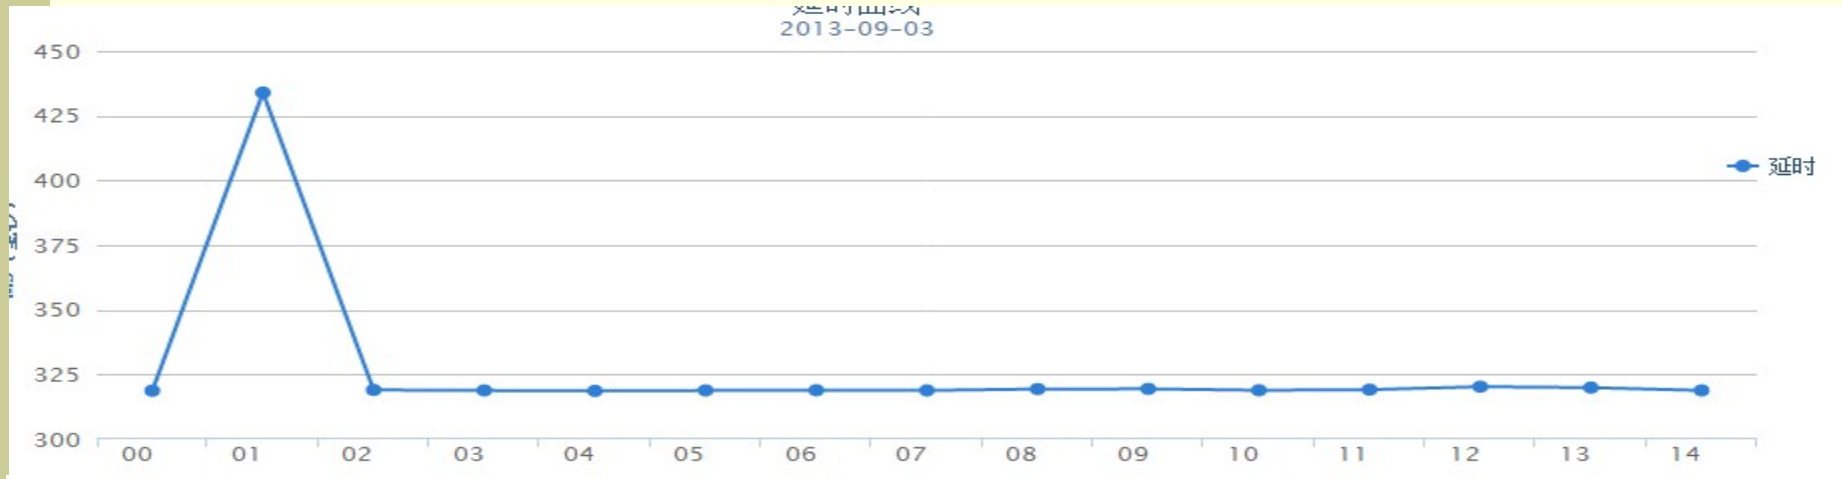

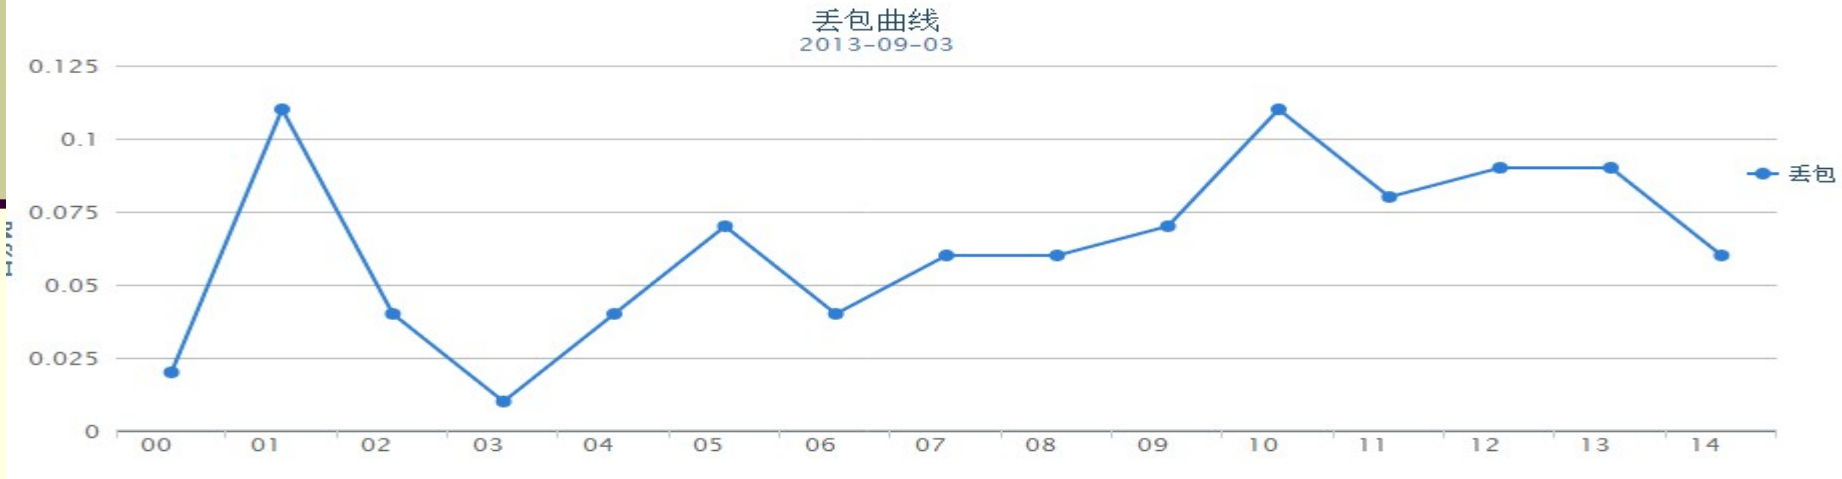

Finding #4
Routing selection(cont)

- While ping from CERNET, the result is much different

151 packets transmitted, 150 received, 0% packet loss, time 150074ms

rtt min/avg/max/mdev = 275.563/275.778/276.395/0.726 ms

- Traceroute results:

traceroute to www.adnxs.com (8.12.226.201), 30 hops max, 60 byte packets

```
 1  203.91.120.129 (203.91.120.129)  0.425 ms  0.438 ms  0.532 ms
 2  202.38.120.189 (202.38.120.189)  0.287 ms  0.274 ms  0.257 ms
 3  202.112.61.253 (202.112.61.253)  0.241 ms  0.223 ms  0.203 ms
 4  202.112.62.49 (202.112.62.49)  1.249 ms  1.251 ms  1.235 ms
 5  202.112.61.158 (202.112.61.158)  0.365 ms  0.429 ms  0.429 ms
 6  202.112.61.122 (202.112.61.122)  1.648 ms  1.007 ms  0.986 ms
 7  202.112.61.18 (202.112.61.18)  54.876 ms  54.882 ms  54.905 ms
…
```

Traceroute from Tsinghua University

traceroute to www.adnxs.com (8.12.226.201), 30 hops max, 60 byte packets
1 106.120.132.1 (106.120.132.1) 0.184 ms 0.168 ms 0.146 ms
2 73.254.202.1.static.bjtelecom.net (1.202.254.73) 2.802 ms 2.920 ms 3.035 ms
3 bj141-158-21.bjtelecom.net (219.141.158.21) 6.717 ms 6.732 ms 6.712 ms
4 bj141-131-53.bjtelecom.net (219.141.131.53) 2.014 ms 118.84.3.25 (118.84.3.25) 3.402 ms 118.84.3.21 (118.84.3.21) 4.622 ms
5 202.97.53.98 (202.97.53.98) 2.103 ms 202.97.53.170 (202.97.53.170) 1.443 ms 202.97.53.98 (202.97.53.98) 2.050 ms
6 202.97.58.118 (202.97.58.118) 2.376 ms 3.520 ms 3.489 ms

Summary

- Performance issues are caused by various reasons
  - DNS
  - CDN
  - Routing
  - …
- It is complicated to optimize a large network
  - Some are caused by local configuration, some are rooted in general design
- It is always helpful to deploy a measurement system
  - Measurement will be more helpful with a purpose

Summary(2)

- Cooperation from multiple sides is very necessary in measurement
  - Measurement from other ASes
  - CDN operators
  - DNS administrators
  - Network operators
  - Web site developers
  - Users
  - ….

Q&A