Internet Measurement and Data Analysis (1)

Kenjiro Cho

2011-09-28

introduction

how does the entire Internet look like?

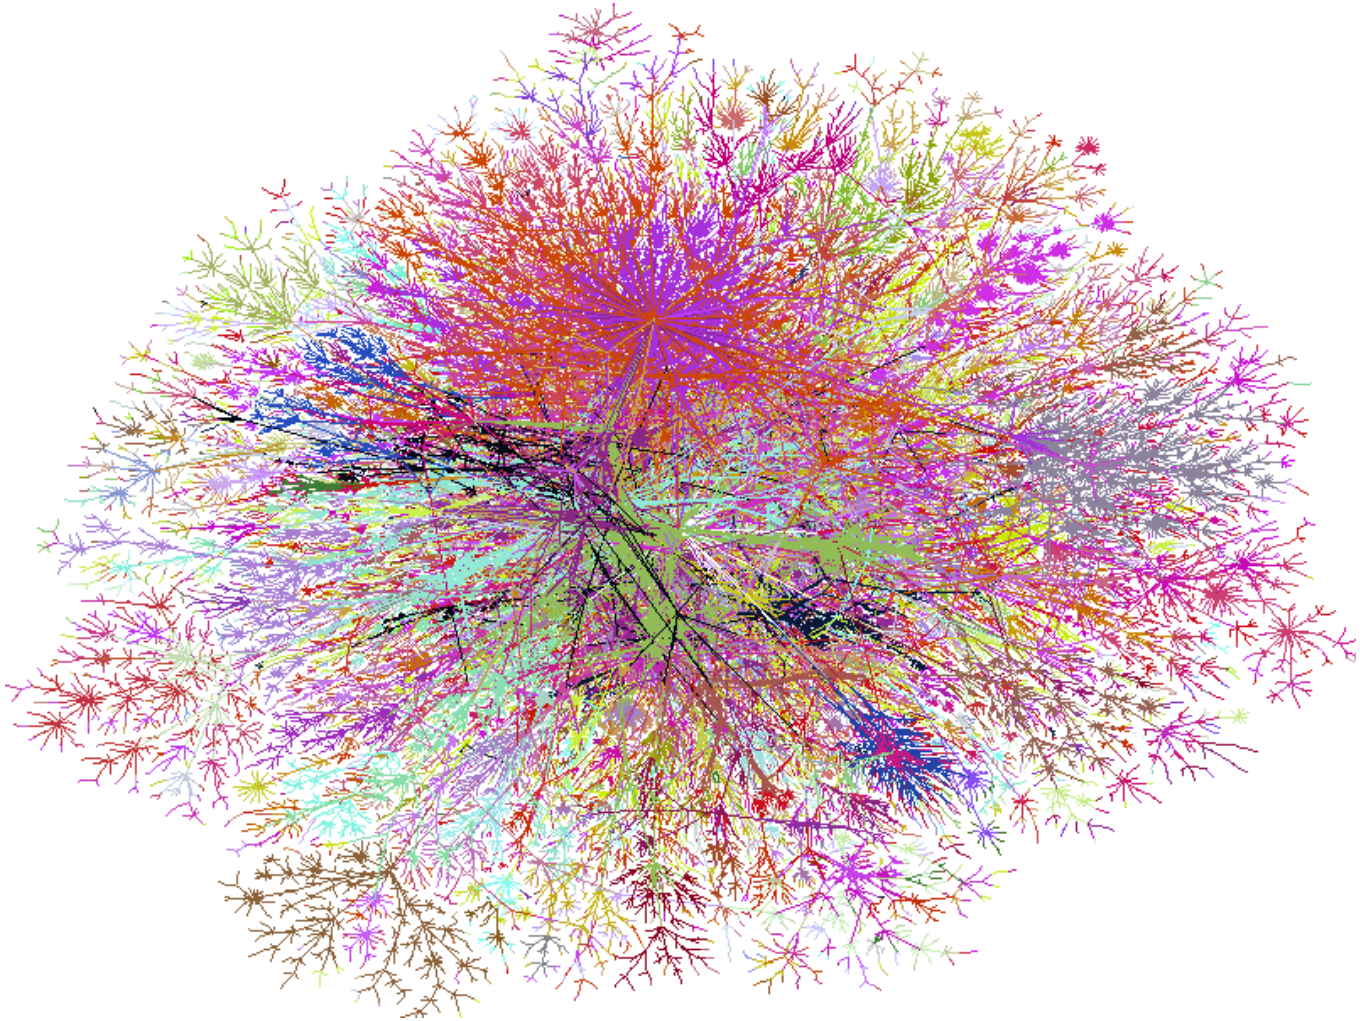

lumeta internet mapping http://www.lumeta.com

http://www.cheswick.com/ches/map/

introduction (cont'd)

how does the entire Internet look like?

- no one knows
- but, everyone is interested

the theme of the class

- looking at the Internet from different views
  - how to measure what is difficult to measure
  - how to extract useful information from huge data sets

this kind of approach will be increasingly important in the future information society.

self-introduction

Kenjiro Cho

- positions
  - Research Director, IIJ Research Lab
  - Guest Professor, Keio SFC
  - Adjunct Professor, JAIST
  - Board member, WIDE Project
- bio
  - BE in electronics from Kobe University in 1984.
    - started as a hardware engineer at Canon, Inc, then became interested in operating systems
  - M.Eng in computer science from Cornell University in 1993
    - studied computer science and distributed systems
  - Researcher at Sony Computer Science Labs from 1996
    - research on the Internet
  - Ph.D. (Media and Governance) from Keio University in 2001
  - Researcher at IIJ from 2004
- research topics
  - Internet measurement and management
  - quality of Service communications
  - networking support in operating systems

Internet measurement and data analysis

- Faculty: Kenjiro Cho ⟨kjc@sfc.keio.ac.jp⟩
- TA: Yohei Kuga ⟨sora@sfc.wide.ad.jp⟩
- SA: Midori Kato ⟨katoon@ht.sfc.keio.ac.jp⟩
  Ryo Nakamura ⟨upa@sfc.wide.ad.jp⟩
- URL: http://web.sfc.keio.ac.jp/~kjc/classes/sfc2011f-measurement/
- support email (facaulty, TA, SA): ⟨imda@sfc.wide.ad.jp⟩
- textbooks, references: the lecture slide materials will be provided online.
- programming: data processing exercises by Ruby
- evaluation: 2 assignments and a final report

class overview

Now that the Internet has become a social infrastructure, it becomes increasingly important to understand the current usage and behavior of the Internet and predict the future, not only for technical aspects but also for investment decisions and policy making.

However, it is challenging to grasp the Internet that is gigantic and complex systems; while it is not realistic to perform large-scale measurement covering the entire Internet, it is often the case that traditional sampling methods cannot be applied. Moreover, there are various technical, social, economical, and legal constraints, and we need to solve problems under these constraints.

In this class, you will learn about the overview of Internet measurement and large-scale data analysis, and basic skills for the forthcoming information society to obtain new knowledge from massive information.

class overview (cont'd)

Theme, Goals, Methods

In this class, you will learn about Internet measurement and data analysis methods, to obtain knowledge and understanding of networking technologies and large-scale data analysis. Each class will provide specific topics where you will learn problems, constraints, and solutions. At the same time, you will learn technical and theoretical backgrounds of the topics such as networking technologies, statistics, and algorithms. Each class consists of a lecture, and exercises on data analysis.

Prerequisites

The prerequisites for the class are basic programming skills and basic knowledge about statistics.

In the exercises and assignments, you will need to write programs to process large data sets, using the Ruby scripting language and the Gnuplot plotting tool. To understand the theoretical aspects, you will need basic knowledge about algebra and statistics. However, the focus of the class is to understand how mathematics is used for engineering applications.

class schedule (1/5)

- Class 1 Introduction (9/28)
  - network measurement and Internet measurement
  - network management tools
  - network measurement tools
  - exercise: introduction of Ruby scripting language
- Class 2 Measuring the size of the Internet (10/5)
  - the number of users and hosts
  - the number of web pages
  - precision, errors, significant digit
  - how to make good graphs
  - exercise: graph plotting by Gnuplot
- Class 3 Data recording and log analysis (10/12)
  - data format
  - log analysis methods
  - exercise: log data and regular expression

class schedule (2/5)

- Class 4 Measuring the speed of the Internet (10/19)
  - bandwidth measurement
  - inferring available bandwidth
  - mean, standard deviation
  - linear regression
  - exercise: mean, standard deviation, linear regression
  - **assignment 1**
- Class 5 Measuring the structure of the Internet (10/26)
  - Internet architecture
  - network layers
  - topologies
  - graph theory
  - exercise: topology analysis
- Class 6 Measuring the characteristics of the Internet (11/2)
  - delay, packet loss, jitter
  - correlation and multivariate analysis
  - principal component analysis
  - exercise: correlation analysis

class Schedule (3/5)

- Class 7 Measuring the diversity and complexity of the Internet (11/9)
  - sampling
  - statistical analysis
  - histogram
  - exercise: histogram, CDF
- Class 8 Distributions (11/16)
  - normal distribution and other distributions
  - confidence intervals
  - statistical tests
  - exercise: generating distributions, confidence intervals
  - **assignment 2**
- Class 9 Measuring time series of the Internet (11/30)
  - Internet and time
  - network time protocol
  - time series analysis
  - exercise: time series analysis

class schedule (4/5)

- Class 10 Measuring traffic of the Internet (11/30?)
  - traffic measurement
  - exercise: traffic measurement
- Class 11 Measuring paths of the Internet (12/7)
  - routing protocols
  - EGP and IGP
  - exercise: routing behavior analysis
- Class 12 Measuring anomalies of the Internet (12/14)
  - anomaly detection
  - spam filters
  - Bayes' theorem
  - exercise: anomaly detection
- Class 13 Data mining (12/21)
  - pattern extraction
  - classification
  - clustering
  - exercise: clustering

class schedule (5/5)

- Class 14 Scalable measurement and analysis (1/11)
  - distributed parallel processing
  - cloud technology
  - exercise: large-scale data processing
- Class 15 Summary (1/18)
  - summary of the class
  - Internet measurement and privacy issues

network measurement and Internet measurement

- network measurement
  - measurement in limited environment
  - snapshot at a time
- Internet measurement
  - measurement of the Internet as a large-scale open system
    - large-scale distributed system
    - open system (continuously changing)

Internet measurement – measuring unmeasurable Internet

- need for generic measurement data for the Internet
  - example: typical packet size distribution
- the Internet is an open system continuously changing, evolving, and expanding
  - no central point, representative locations, different behaviors are observed depending on observing location and time
  - seeking for generality of the Internet: measuring unmeasurables
- for operation of the Internet, for development of protocols and equipment
  - seeking for the best estimates, predicting the future, and revisiting the existing knowledge
- need to consider not only from technical aspects but also from social, political and economical aspects

importance of measurement

measurement is a basis of all technologies

- for networking, it is an attempt to observe invisible networks
- needed for operation, design, implementation, and research

- however, it has become difficult by commercialization of the Internet and widespread use
  - traffic data is confidential for providers and will not be disclosed
  - risks of leaking private information

goals of measurement and data analysis

- operational goals
  - trouble-shooting
  - tuning for performance and reliability
  - monitoring the usage, usage reports
  - long-term planning, cost evaluation of network capacity and equipment
- engineering goals (software, hardware, protocol design and implementations)
  - design trade-offs (e.g., buffer size and its cost)
  - testing and evaluation
  - observing unexpected behaviors (in complex systems)
- research goals (theory, modeling, new findings)
  - characteristics of network behaviors
  - modeling (e.g., behavior of web services)
  - behaviors of complex systems
    - abundant data and tools
- inputs for policy or investment plans

characteristics of network data and behavior

- skewed distributions with large variance
  - inherent mechanism to make burst transfer
  - skewed utilization: e.g., a handful users generate most traffic
- anomalies everywhere
  - bugs, mis-configurations, spec mismatches, accidents, maintenance's
- interferences among various mechanisms
  - e.g., congestion control: Ethernet's collision avoidance, packet queueing, TCP's congestion control, capacity provisioning
- traffic aggregation
  - complex behavior as a whole (more than the sum of the individual components)
- limitations of network measurement
  - many practical issues and limitations exist
  - measurement affects the observed behavior

measurement needs combined skills

- goals could be operational, engineering, scientific
  - all inseparable, all skills required
    - knowledge of operational environment
    - engineering of measurement tools
- output can be facts, findings, new ideas
  - new ideas are not always necessary
  - facts, especially long-term measurement, are valuable
- but you should have clear goals
  - better to start with real problems to solve
    - there are many issues and problems but some are more important than others

why traffic measurement of Internet is so hard?

- massive, diverse and changing traffic
- mechanisms at different layers in different time scale
  - interact with each other
- dynamics
  - Internet mechanisms are adaptive and resilient
  - traditional measurement techniques are often not applicable
- pathological traffic is not unusual
  - by bugs, misconfigurations, errors, mismatches, accidents
- we still don't have good understanding

massive volume of traffic

- unprecedented scale with unprecedented growth
  - e.g., traffic volume: 1Gbps traffic
    - 120MB/sec 7GB/minute 420GB/hour 9.8TB/day
- far more data than we can analyze
  - techniques needed to reduce data size
    - filtering: e.g., record only TCP SYN packets
    - aggregation: e.g., flow-based accounting
    - sampling: e.g., record 1 in n packets
  - also, techniques needed to reduce dimensionality
- still, details matter
  - a big impact often comes
    - from small fraction
    - from minor differences

diverse traffic

- large variation in traffic mix between sites
- backbone vs. access links
  - access line types: fiber, ADSL, modem, wireless, satellite
    - differences in bandwidth, delay, loss

typical traffic doesn't exist!

constant change of traffic pattern

- daily, weekly traffic pattern
- trend changes over time
  - web in 90s and p2p in 2000s completely changed traffic pattern
- hard to predict future!

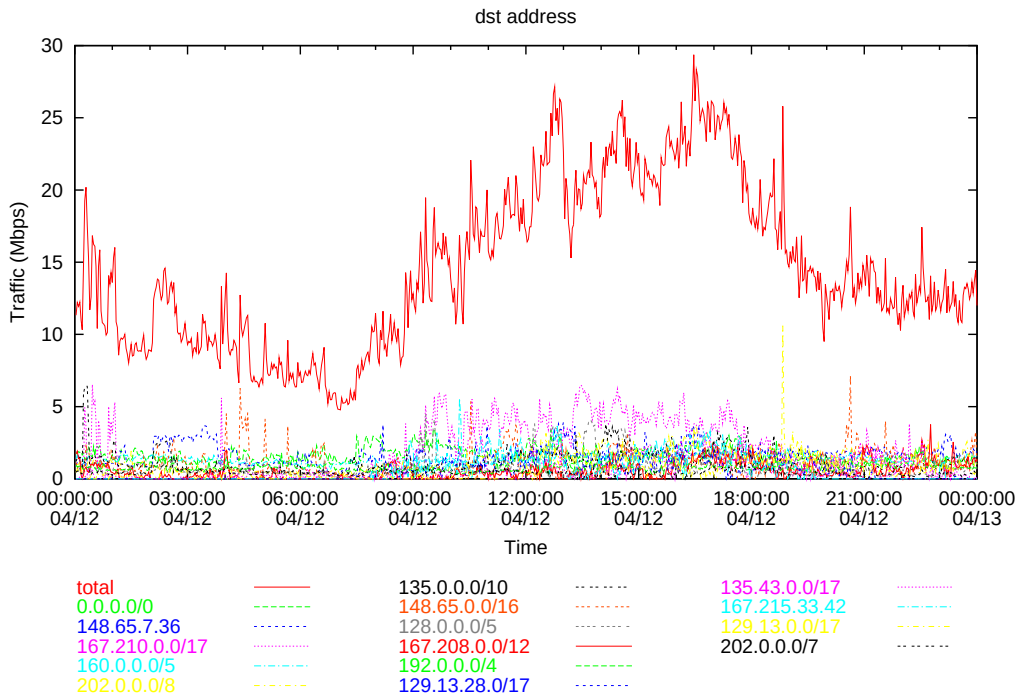

limitations of Internet measurement

- problems often occur at boundaries of different networks
  - cooperation needed but not easy
- measurement affects the behavior of the observed network
- need understanding and help from operators
  - need to understand operational requirements and find suitable methods for measurement
- cost: measurement doesn't come free
  - limitations to measure high-end routers with a PC
- privacy and confidential information in data
  - barriers for researchers to access commercial data

summary

Internet measurement and data analysis

- measurement is basis for all technologies
- for networking, it is an attempt to observe invisible networks
- need to consider not only from technical aspects but also from social, political and economical aspects

theme of the class

- Internet measurement and data analysis as case studies
- learn how to measure what is difficult to measure
- learn how to extract useful information from huge data sets

Network Management Tools

commonly-used management tools

network management tools (originally not designed for measurement)

- ping
  - reachability, round-trip time
- traceroute
  - path detection
- tcpdump
  - packet capturing
- SNMP
  - usage monitoring, network equipment status monitoring

ping

- a popular and widely-available tool to check connectivity
- ICMP-echo request/reply
- limitations
  - ping responses do not mean network is working correctly
  - ICMP is not representative of host/network performance

router architecture

- fast path: hardware assisted processing
- slow path: software processing
  - ICMP packets are processed via slow path

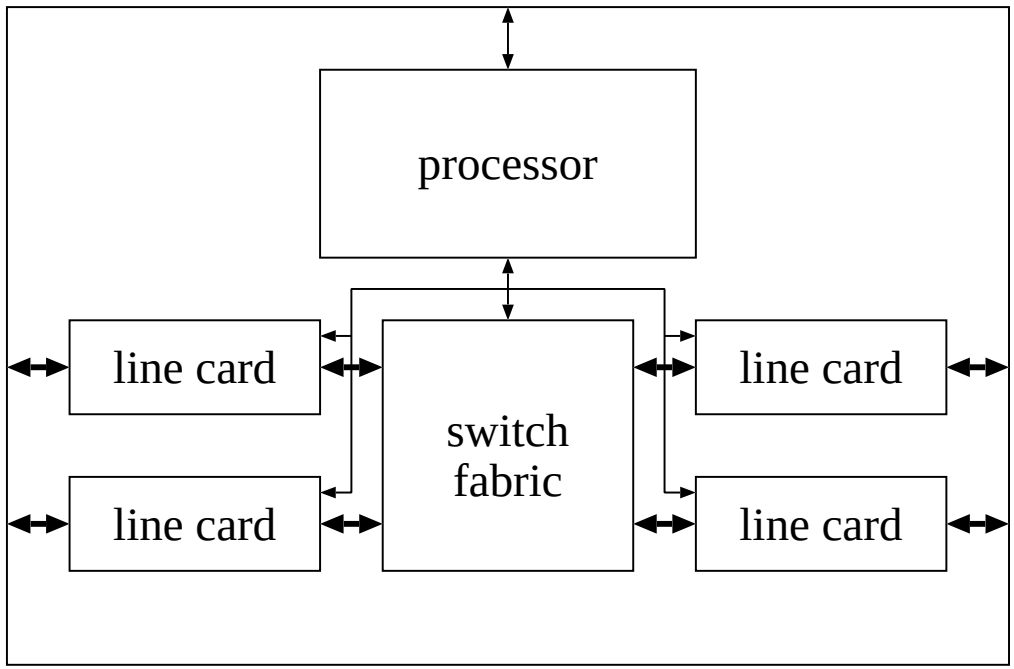

ping sample output

```
% ping -c 10 www.ait.ac.th
PING www.ait.ac.th (202.183.214.46): 56 data bytes
64 bytes from 202.183.214.46: icmp_seq=0 ttl=114 time=112.601 ms
64 bytes from 202.183.214.46: icmp_seq=1 ttl=114 time=106.730 ms
64 bytes from 202.183.214.46: icmp_seq=2 ttl=114 time=106.173 ms
64 bytes from 202.183.214.46: icmp_seq=3 ttl=114 time=111.704 ms
64 bytes from 202.183.214.46: icmp_seq=4 ttl=114 time=112.412 ms
64 bytes from 202.183.214.46: icmp_seq=5 ttl=114 time=114.603 ms
64 bytes from 202.183.214.46: icmp_seq=6 ttl=114 time=111.755 ms
64 bytes from 202.183.214.46: icmp_seq=7 ttl=114 time=115.273 ms
64 bytes from 202.183.214.46: icmp_seq=8 ttl=114 time=106.525 ms
64 bytes from 202.183.214.46: icmp_seq=9 ttl=114 time=111.562 ms

--- www.ait.ac.th ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 106.173/110.934/115.273/3.142 ms
```

traceroute

- exploit TTL (time-to-live) of IP designed for loop prevention
  - TTL is decremented by each intermediate router
  - router returns ICMP TIME EXCEEDED to the sender when TTL becomes 0
- limitations
  - path may change over time
  - path may be asymmetric
    - can observe only out-going paths
  - report from one of the interfaces of the router
    - hard to identify interfaces belonging to same router

traceroute sample output

```
% traceroute www.ait.ac.th
traceroute to www.ait.ac.th (202.183.214.46), 64 hops max, 40 byte packets
 1  202.214.86.129 (202.214.86.129)  0.687 ms  0.668 ms  0.730 ms
 2  jc-gw0.IIJ.Net (202.232.0.237)  0.482 ms  0.390 ms  0.348 ms
 3  tky001ix07.IIJ.Net (210.130.143.233)  0.861 ms  0.872 ms  0.729 ms
 4  tky001bb00.IIJ.Net (210.130.130.76)  10.107 ms  1.026 ms  0.855 ms
 5  tky001ix04.IIJ.Net (210.130.143.53)  1.111 ms  1.012 ms  0.980 ms
 6  202.232.8.142 (202.232.8.142)  1.237 ms  1.214 ms  1.120 ms
 7  ge-1-1-0.toknf-cr2.ix.singtel.com (203.208.172.209)  1.338 ms  1.501 ms
 1.480 ms
 8  p6-13.sngtp-cr2.ix.singtel.com (203.208.173.93)  93.195 ms  203.208.172.
229 (203.208.172.229)  88.617 ms  87.929 ms
 9  203.208.182.238 (203.208.182.238)  90.294 ms  88.232 ms  203.208.182.234
(203.208.182.234)  91.660 ms
10  203.208.147.134 (203.208.147.134)  103.933 ms  104.249 ms  103.986 ms
11  210.1.45.241 (210.1.45.241)  103.847 ms  110.924 ms  110.163 ms
12  st1-6-bkk.csloxinfo.net (203.146.14.54)  131.134 ms  129.452 ms  111.408
 ms
13  st1-6-bkk.csloxinfo.net (203.146.14.54)  106.039 ms  105.078 ms  105.196
 ms
14  202.183.160.121 (202.183.160.121)  111.240 ms  123.606 ms  112.153 ms
15  * * *
16  * * *
17  * * *
```

tcpdump

- packet capturing tool
  - capture the first N bytes of packets
- flexible filtering
  - e.g., capture only TCP SYN from host X
- enables detailed analysis
- limitations
  - huge volume
  - difficult to capture on high-speed links

tcpdump sample output

```
18:45:29.767497 IP 202.214.86.132.50052 > 202.210.220.18.80: \
  S 3304970307:3304970307(0) win 65535 <mss 1460,nop,nop,sackOK,nop, \
  wscale 1,nop,nop,timestamp 710778973 0>
18:45:29.770038 IP 202.210.220.18.80 > 202.214.86.132.50052: \
  S 3129218301:3129218301(0) ack 3304970308 win 65535 <mss 1460,nop, \
  ywscale 1,nop,nop,timestamp 2523776361 710778973,nop,nop,sackOK>
18:45:29.770090 IP 202.214.86.132.50052 > 202.210.220.18.80: \
  . ack 1 win 33304 <nop,nop,timestamp 710778973 2523776361>
18:45:29.787084 IP 202.214.86.132.50052 > 202.210.220.18.80: \
  P 1:521(520) ack 1 win 33304 <nop,nop,timestamp 710778975 2523776361>
18:45:29.791392 IP 202.210.220.18.80 > 202.214.86.132.50052: \
  P 1:222(221) ack 521 win 33304 <nop,nop,timestamp 2523776363 710778975>
18:45:29.887024 IP 202.214.86.132.50052 > 202.210.220.18.80: \
  . ack 222 win 33304 <nop,nop,timestamp 710778985 2523776363>
18:45:34.792726 IP 202.210.220.18.80 > 202.214.86.132.50052: \
  F 222:222(0) ack 521 win 33304 <nop,nop,timestamp 2523776864 710778985>
18:45:34.792763 IP 202.214.86.132.50052 > 202.210.220.18.80: \
  . ack 223 win 33304 <nop,nop,timestamp 710779475 2523776864>
18:45:42.528539 IP 202.214.86.132.50052 > 202.210.220.18.80: \
  F 521:521(0) ack 223 win 33304 <nop,nop,timestamp 710780249 2523776864>
18:45:42.531088 IP 202.210.220.18.80 > 202.214.86.132.50052: \
  . ack 522 win 33303 <nop,nop,timestamp 2523777637 710780249>
```

SNMP (Simple Network Management Protocol)

- SNMP allows a remote user to
  - query information, store information, set traps
  - by UDP (unreliable)
- standardized set of traffic statistics
  - supported by most of routers, switches, host OS
  - many management/monitoring products
- MIB (Management Information Base)
  - tree structured database of SNMP objects
    - e.g., interfaces.ifTable.ifEntry.ifOutOctets
    - standard MIBs and private MIBs
  - get, set, get-next to access MIB
- limitations
  - supported statistics are limited
    - most counter statistics are hard-coded, e.g.,interface counters
  - accessing to MIB objects is expensive

MRTG

- popular tool to show SNMP data
- time series data aggregated over time
  - daily, weekly, monthly to bound the storage size
- inbound/outbound traffic
  - can be used for other types of time series data

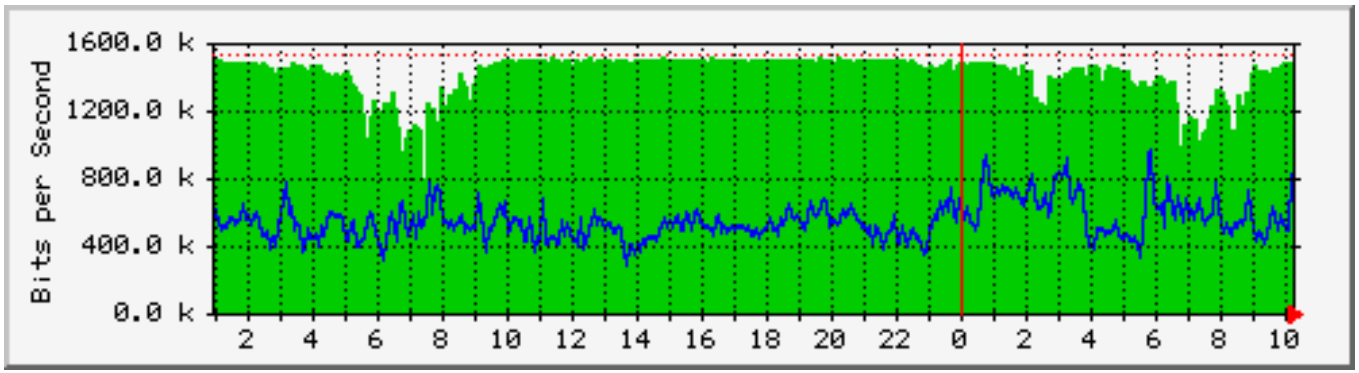

RRDtool

- RRDtool: successor of MRTG
  - flexible configuration, graphing
  - can be used for any time-series data
- flowscan: visualizes netflow data by rrdtool

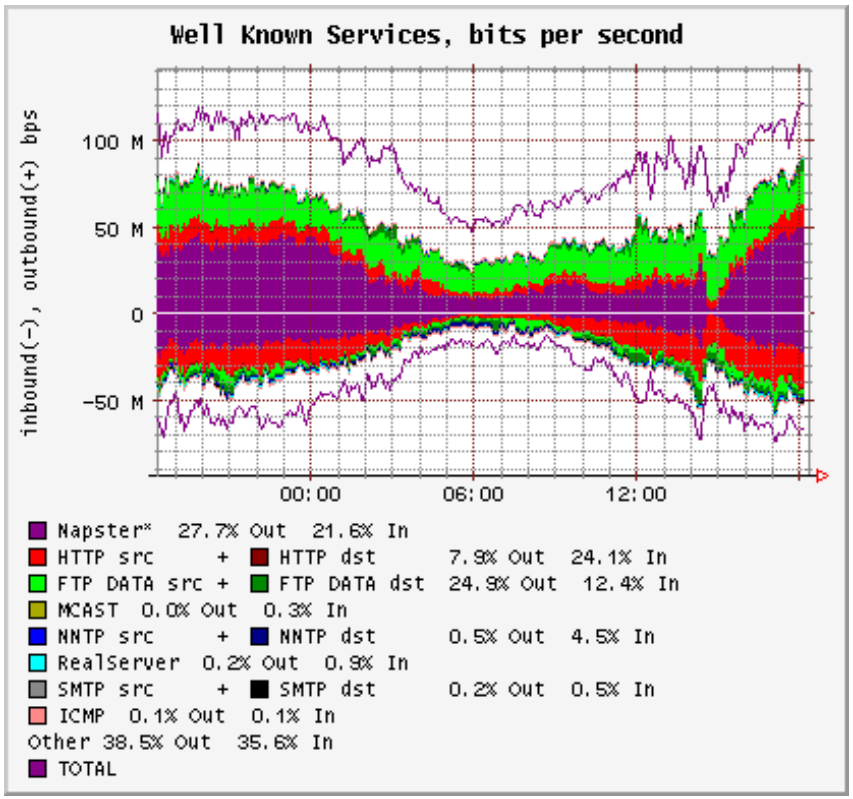

from caida web site

summary

network management tools

- not originally designed for measurement
- still often used for measurement
- when using for measurement, need to understand the mechanisms and limitations

Introduction to Ruby

Ruby

- a scripting language for object-oriented programming
- supports wide range of functions for text processing and system management
- free software started in 1993
- original author: Yukihiro Matsumoto
- became popular for Ruby on Rails (a web application framework)

Ruby information
Ruby official site: `http://www.ruby-lang.org/`
Ruby reference manual: `http://www.ruby-lang.org/en/documentation/`
Ruby の歩き方: `http://jp.rubyist.net/magazine/?FirstStepRuby`

Ruby characteristics

- interpreter language: no need to compile for execution
- highly portable: runs on most platforms
- simple syntax
  - no predefined data type for variables, variables can store any data and are dynamically typed
  - no need to declare variables, variable types (local variables, global variables, instance variables) can be inferred from variable names
- garbage collection: users do not need to manage memory
- object-oriented
  - everything is an object
  - class, inheritance, methods
  - iterator and closure
    - control structures and procedures can be written in object-oriented manner
- unique object-oriented features (instance specific methods, mixin, etc)
- powerful string operations/regular expressions
- built-in support for large integers
- Ruby's shortcomings: a bit slower than its competitors

Ruby commands

- irb: Ruby's interactive interface

```
$ irb --simple-prompt
>> puts "Hello"
Hello
```

- ruby: Ruby main program

```
$ ruby test.rb
```

or,

```
$ ruby -e 'puts "Hello".reverse'
olleH
```

- rdoc: extract text in RDoc format in a source file. RDoc is a format to be built into a Ruby script.
- ri: a command line tool to read Ruby documents in the RDoc format

```
$ ri String#upcase
```

exercise: a program to count text lines

count the number of text lines in a file given by the argument

```
filename = ARGV[0]
count = 0
file = open(filename)
while text = file.gets
  count += 1
end
file.close
puts count
```

write to “count.rb” and then run it

```
$ ruby count.rb foo.txt
```

rewrite it in a more rubyish way

```
#!/usr/bin/env ruby
count = 0
ARGF.each_line do |line|
  count += 1
end
puts count
```

next class

Class 2 Measuring the size of the Internet (10/5)

- the number of users and hosts
- the number of web pages
- precision, errors, significant digit
- how to make good graphs
- exercise: graph plotting by Gnuplot

references

[1] Ruby official site. `http://www.ruby-lang.org/`

[2] gnuplot official site. `http://gnuplot.info/`

[3] Mark Crovella and Balachander Krishnamurthy. *Internet measurement: infrastructure, traffic, and applications*. Wiley, 2006.

[4] Antonio Nucci and Konstantina Papagiannaki. *Design, Measurement and Management of Large-Scale IP Networks: Bridging the Gap Between Theory and Practice* . Cambridge University Press, 2008.

[5] Pang-Ning Tan, Michael Steinbach and Vipin Kumar. *Introduction to Data Mining*. Addison Wesley, 2006.

[6] Raj Jain. *The art of computer systems performance analysis*. Wiley, 1991.

[7] あきみち、空閑洋平. **インターネットのカタチ**. オーム社, 2011.

[8] 井上洋, 野澤昌弘. **例題で学ぶ統計的方法**. 創成社, 2010.

[9] 平岡和幸, 掘玄. **プログラミングのための確率統計**. オーム社, 2009.