

Measurement Activities at WIDE

Kenjiro Cho
IIJ/WIDE Project


August 24 2009



WIDE Project

WIDE: Widely Integrated Distributed Environment

- ▶ a research consortium in Japan since 1998
 - ▶ about 100 sponsor companies
 - ▶ 40 universities
 - ▶ 200 active members



GLIF world map on GeoCosmos
courtesy of Global Lambda Integrated Facility

Research at WIDE

motto: research on our left hand, operation on our right hand.
supporting social infrastructure with both hands.

- ▶ protocols and middleware
 - ▶ KAME/USAGI, NEMO, MANET, DVTS
- ▶ testbeds
 - ▶ AI3, StarBED, JGN2plus, GLIF, WiMAX
- ▶ real space Internet
 - ▶ AutoID, Live E!, Locky, InternetCAR
- ▶ social interaction
 - ▶ SOI/SOI-Asia, Lens



characteristics of WIDE traffic

- ▶ WIDE
 - ▶ internet research through live network
- ▶ WIDE has its own backbone operated by members
 - ▶ backbone includes
 - ▶ international links
 - ▶ IXes
 - ▶ root name servers
 - ▶ various link types up to 10GbE
 - ▶ carrying both commodity traffic and experiments
 - ▶ commodity: university traffic, WIDE members
 - ▶ experiments: new products, our technologies under development
 - ▶ IPv6 everywhere
 - ▶ events (including firedrills)
 - ▶ not a typical internet but a showcase

traffic measurement and analysis in WIDE

- ▶ measurement activities across research groups
- ▶ broad perspectives
 - ▶ tracking long-term trends
 - ▶ analysis (with wide range of granularity)
 - ▶ operational tools (trouble-detection/shooting)
 - ▶ evaluation of new technologies
- ▶ emphasis on
 - ▶ wide-area
 - ▶ multi-point
 - ▶ measurement on backbone
 - ▶ long-term
 - ▶ continuation by group effort

traffic measurement activities within WIDE

1. MAWI traffic repository
2. Residential Broadband Traffic Study
3. IX Traffic Study
4. Gulliver Project
5. Regional AS Topology Structures
6. Anomaly Detection by Sketch and Non Gaussian Multiresolution Statistical Detection Procedures
7. Longitudinal Statistical Analysis based on Robust Estimation
8. Host Clustering by Communication Patterns

international collaboration

- ▶ CAIDA (the Cooperative Association for Internet Data Analysis)
 - ▶ collaboration since 2003 on DNS, topology, routing
- ▶ CASFI
 - ▶ joint-workshops, data sharing
- ▶ CNRS
 - ▶ measurement and modeling of emerging applications and security threats
- ▶ RSSAC (ICANN Root Server Systems Advisory Committee)
 - ▶ root name server measurement
 - ▶ WIDE, CAIDA, ISC OARC, USC/ISI
- ▶ other collaboration
 - ▶ routeviews, RIPE, INRIA, AIT
- ▶ A day in the life of the Internet
 - ▶ simultaneous measurement worldwide to promote research based on re-examinable data
 - ▶ CAIDA, CASFI, WIDE, DNS root ops, etc.

MAWI Traffic Repository

- ▶ pcap packet traces from WIDE backbones
- ▶ anonymized traces publicly available
- ▶ many papers used MAWI traces

<http://mawi.wide.ad.jp/mawi/>

Kenjiro Cho, Koushirou Mitsuya and Akira Kato.
Traffic Data Repository at the WIDE Project.
USENIX FREENIX Track, San Diego, CA, June 2000.

Residential Broadband Traffic Study (1/2)

key question: what is the macro level impact of video and other rich media content on traffic growth at the moment?

- ▶ traffic growth is one of the key factors driving research, development and investment in technologies and infrastructures
- ▶ crucial is the balance between demand and supply
- ▶ measurements: 2 data sets
 - ▶ aggregated SNMP data from 6 ISPs covering 42% of Japanese traffic
 - ▶ Sampled NetFlow data from 1 ISP

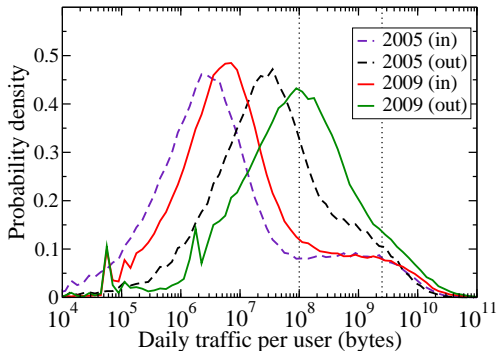
K. Cho, K. Fukuda, H. Esaki, and A. Kato.
Observing Slow Crustal Movement in Residential User Traffic.
ACM CoNEXT2008, Madrid, Spain, Dec. 2008.

K. Cho, K. Fukuda, H. Esaki, and A. Kato.
The impact and implications of the growth in residential user-to-user traffic.
ACM SIGCOMM2006, Pisa, Italy, Aug. 2006.

Residential Broadband Traffic Study (2/2)

daily traffic volume per user

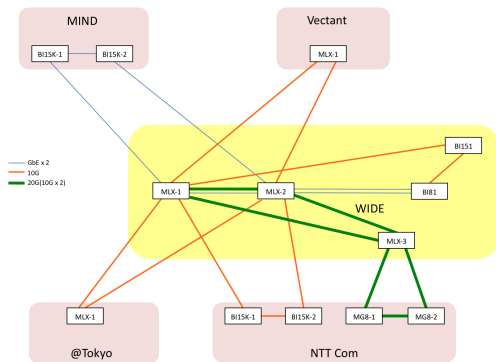
- ▶ increase in download volume of client-type users
 - ▶ out mode: from 32MB/day to 114MB/day
 - ▶ in mode: from 3.5MB/day to 6MB/day
- ▶ while peer-type dist. isn't growing much (mode:2GB/day)



changes in daily traffic per user (2005 vs. 2009)

IX Traffic Study


- ▶ the 1st IX in Japan started by WIDE as research experiment
- ▶ one of the 3 major IXes in Japan



DIX-IE topology (Aug 2009)

sFlow Measurement at IXes

- ▶ planning to deploy sflow measurement at DIX-IE
- ▶ currently, being tested in WIDE backbone (not at IX)
- ▶ tools: nfsen (<http://nfsen.sourceforge.net/>)



IPv4 vs IPv6 traffic from a WIDE backbone link

Gulliver Project (1/4)


- ▶ distributed active measurement project
 - ▶ DNS reachability
 - ▶ traceroute
- ▶ small box as measurement platform
 - ▶ NetBSD-based router product
 - ▶ remote management framework by IIJ
 - ▶ long MTBF (no HDD)
 - ▶ low management cost



IIJ seil plus

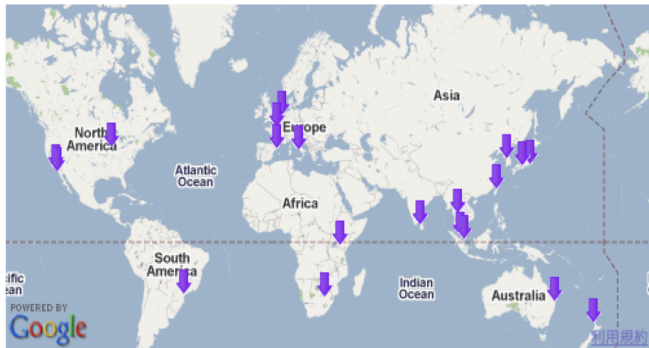
Gulliver Project (2/4)

- ▶ started in August 2007
- ▶ <http://gulliver.wide.ad.jp/>




Gulliver Project (3/4)

- ▶ 30 probes as of August 2009
- ▶ around the world, focusing on developing countries



gulliver probe locations

Gulliver Project (4/4)




gulliver: measured RTT to DNS root servers (2008)

Regional AS Topology Structures (1/2)

- ▶ study on topologies from regional views
 - ▶ Internet development affected by geographical constraints
 - ▶ existing global topologies do not capture regionality
- ▶ goals
 - ▶ understand similarities and differences in regional Internet structures
 - ▶ e.g., identifying regional hub ASes, hub cities
- ▶ methodology
 - ▶ extract nodes in the region in traceroute data sets
 - ▶ by reverse DNS names, geo-IP mapping, etc
- ▶ visualization by CAIDA's AS core map

Yohei Kuga, Kenjiro Cho, Osamu Nakamura.
On inferring regional AS topologies.
AINTEC2008. Bangkok, Thailand. Nov. 2008.

Regional AS Topology Structures (2/2)



regional topology visualization by AS Core Map

Europe (top left), Asia (top right), US-west (bottom left), US-east (bottom right)

Anomaly Detection by Sketch and Non Gaussian Multiresolution Statistical Detection Procedures (1/2)

collaboration with Abry's team at ENS-Lyon


- ▶ features

- ▶ generates self-reference from the target traffic, no training required
- ▶ can detect small hidden anomalies
- ▶ works with uni-directional data (applicable to backbone)

Guillaume Dewaele, Kensuke Fukuda, Pierre Borgnat, Patrice Abry, Kenjiro Cho.
Extracting Hidden Anomalies using Sketch and Non Gaussian Multiresolution
Statistical Detection Procedures.
SIGCOMM2007 LSAD Workshop. Kyoto Japan. August 2007.

Anomaly Detection by Sketch and Non Gaussian Multiresolution Statistical Detection Procedures (2/2)

- ▶ sketch: divides packets into N groups by hashing
- ▶ for each group, extract statistical features from packet arrival distribution in multiple time resolutions
- ▶ compare normalized features among the groups, detect deviations as anomalies



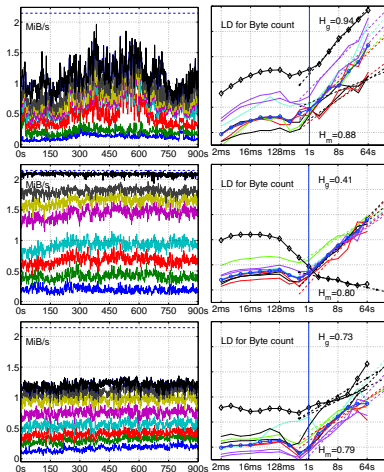
Longitudinal Statistical Analysis based on Robust Estimation (1/3)

the same sketch technique is used to extract typical behaviors for long-term traffic analysis

- ▶ traffic statistics show huge variability
- ▶ median sketch as robust estimation to observe traffic evolution
- ▶ strong and persistent long range dependence found

Pierre Borgnat, Guillaume Dewaele, Kensuke Fukuda, Patrice Abry, Kenjiro Cho.
Seven Years and One Day: Sketching the Evolution of Internet Traffic.
INFOCOM2009. Rio de Janeiro, Brazil. April 2009.

Longitudinal Statistical Analysis based on Robust Estimation (2/3)




top: no congestion, middle: congestion, bottom: severe anomalies

left: sketch sub-trace time series, right: LDs for global, sketches, median-sketch

Longitudinal Statistical Analysis based on Robust Estimation (3/3)

- ▶ LRD over 7 years: global (gray) and median-sketch (black) estimates of H during 2001-2008



Host Clustering by Communication Patterns (1/3)

- ▶ profiling traffic at host level
- ▶ unsupervised statistical classification
- ▶ 9D features for clustering
- ▶ cross validation with existing tools
- ▶ visualization by graphlets

G. Dewaele, Y. Himura, P. Borgnat, K. Fukuda, P. Abry, O. Michel, R. Fontugne, K. Cho, H. Esaki.

Unsupervised host behavior classification from connection patterns.
submitted for publication.

Host Clustering by Communication Patterns (2/3)


9 features for clustering (S: entropy)

- ▶ network connectivity
 - ▶ # of dst addrs
 - ▶ # of src ports / # of dst addrs
 - ▶ # of dst ports / # of dst addrs
- ▶ connection dispersion in address space
 - ▶ $S(IP_2)/S(IP_4)$
 - ▶ $S(IP_3)/S(IP_4)$
- ▶ packet size distribution
 - ▶ mean # of packets/flow
 - ▶ % of small packets (≤ 144 bytes)
 - ▶ % of large packets (≥ 1392 bytes)
 - ▶ $S(\text{medium size packets})$

Host Clustering by Communication Patterns (3/3)

Minimum Spanning Tree (MST) clustering


1. plot hosts in (reduced 2D) feature space, pick a random host to start MST
2. remove edges longer than a threshold (dashed line)
3. dense clusters are divided further



MST clustering illustrated in 2D space

Host Connection Pattern Visualization

- ▶ goal: inspections of anomaly detection/host behavior clustering results
- ▶ graphlets (BLINC[Karagiannis05]) to show src-based 5 tuple patterns
- ▶ automation tool under development



sample graphlets: P2P (left) and scanning (right)

Summary

measurement activities at WIDE

- ▶ operational support
- ▶ tool development
- ▶ data sharing
- ▶ modeling and analysis
- ▶ visualization
- ▶ many more measurement related activities

WIDE is interested in open research collaboration

<http://www.wide.ad.jp/>