インターネット計測とデータ解析 第13回

長 健二朗

2011年1月19日

前回のおさらい

スケールする計測と解析

- 計算量
- 分散並列処理
- クラウド技術
- MapReduce

今日のテーマ

まとめ

- これまでのまとめと今後の展望
- WIDE プロジェクトの計測研究

科目概要

いまや社会基盤となったインターネットの現状や挙動を把握し、今後を予想することは、技術面のみならず投資判断や政策決定にとっても重要な課題である。

しかし、大規模複雑システムであるインターネットを把握することは難しい。インターネット全体を網羅する大規模な計測は現実的でない一方で、従来のサンプリング手法も適用できない場合が多い。 さらに、技術的、社会的、経済的、法的にも多くの制約があり、その中で問題を解決する必要がある。

本授業は、インターネットの計測技術と大規模データ解析の概要について学び、情報社会で必須となる大量情報から新たな知識獲得をするための基礎能力を身につける。

主題と目的／授業の手法など

インターネット計測とデータ解析手法について学習し、ネットワーク技術と大規模データ処理の総合的な知識と理解を得る。具体的な応用例について、そこでの問題と制約、その工学的な解決手法を学び、同時に、その背後にあるネットワーク技術、数学、統計、アルゴリズムとそれらの関連を理解する。本授業は、システム系科目と解析系科目を関連づけて統合理解する科目である。

授業計画 (1/3)

- 第 1 回 イントロダクション (9/29)
  - ネットワーク計測とインターネット計測
  - ネットワーク管理ツール
  - 計測ツール
- 第 2 回 インターネットのサイズを計る (10/6)
  - ユーザ数、ホスト数
  - ウェブページ数
  - 精度 誤差 有効数字
- 第 3 回 インターネットの構造を計る (10/13)
  - インターネットアーキテクチャ
  - ネットワーク階層
  - 経路制御
  - トポロジー
  - グラフ理論
- 第 4 回 インターネットの速度を計る (10/20)
  - 速度計測
  - 利用可能帯域の推測
  - 平均 標準偏差
  - 線形回帰
  - **課題 1**

授業計画 (2/3)

- 第 5 回 インターネットの特徴量を計る (10/27)
  - 遅延、パケットロス、ジッタ
  - フロー計測
  - 相関と多変量解析
  - グラフによる可視化
- 第 6 回 インターネットの多様性と複雑さを計る (11/10)
  - ロングテールとさまざまな分布
  - サンプリング
  - 統計解析 (ヒストグラム、期待値と大数の法則、検定と信頼区間)
- 第 7 回 インターネットの時間変化を計る (11/17)
  - インターネットと時刻
  - 時系列解析
  - **課題 2**
- 第 8 回 インターネットの挙動を計る (12/8)
  - トラフィック量
  - 経路情報
  - インターネット計測とプライバシー

授業計画 (3/3)

- 第 9 回 インターネットの異常や問題を計る (12/11)
  - 異常検出
  - スパム判定
  - ベイズ理論
- 第 10 回 データの記録とログ解析 (12/15)
  - データフォーマット
  - ログ解析手法
- 第 11 回 データマイニング (12/22)
  - パターン抽出
  - クラス分類
  - クラスタリング
    - 距離と類似度
    - クラスタリング手法
- 第 12 回 スケールする計測と解析 (1/12)
  - 計算量
  - 分散並列処理
  - クラウド技術
  - MapReduce
- 第 13 回 まとめ (1/19)
  - これまでのまとめと今後の展望
  - WIDE プロジェクトの計測研究

反省点

- 予想以上に学生の知識のばらつきが大きかった
- 授業が思ったほどインタラクティブにならなかった
- 計測部分の具体的実例が少なかった
  - 数学や計算機科学の部分の方が教えやすい
- 次期授業に向けての改善点
  - 大量データを解析する技術の習得に重点
  - (データがあれば基本的な解析ができる学生を作る)
  - 演習を増やす (毎回)

Understanding Internet Dynamics from a Global View

the Internet is an evolving open system

- no central point, no typical network or user
- complexity everywhere: topology, usage distribution, etc
- interferences among various components at different layers
- local optimizations often lead to a negative impact to the global system
- importance of global views
  - still, details matter (global impact of local mechanisms)

we still don’t have real science!

WIDE Project

WIDE: Widely Integrated Distributed Environment

- a research consortium in Japan since 1988
  - about 100 sponsor companies
  - 40 universities, 5 national research institutions
  - 200 active members
- 1st decade (1988-1997): building the Internet
  - deployment, internationalization
- 2nd decade (1998-2007): for everyone, anytime, anywhere
  - IPv6, mobility, multimedia
- 3rd decade (2008-2017): beyond Internet
  - real-space networking, sensors, broadcasting

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 (Korean measurement effort)
  - joint-workshops, data sharing
- CNRS
  - measurement and modeling of emerging applications and security threats
- other collaboration
  - ISC OARC, USC/ISI, 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 investiment 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)

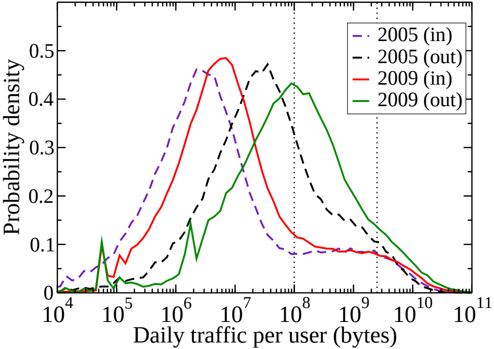

changes in daily traffic per user (2005 vs. 2009)

IX Traffic Study

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

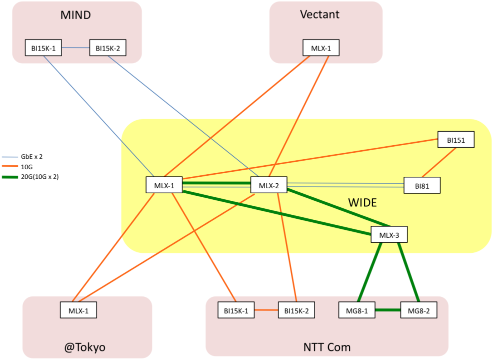

DIX-IE topology (Aug 2009)

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/)

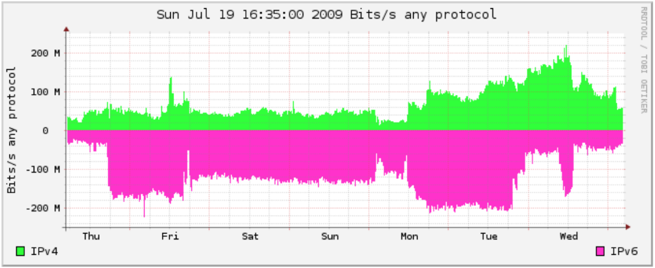

IPv4 vs IPv6 traffic from a WIDE backbone link

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/

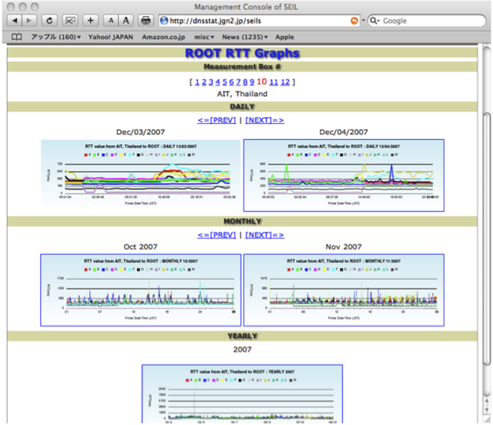

Management Console of SEIL

http://dnsstat.jgn2.jp/seils

ROOT RTT Graphs

Measurement Box #

[1 2 3 4 5 6 7 8 9 10 11 12]

AIT, Thailand

DAILY

<=[PREV] | [NEXT]=>

Dec/03/2007

Dec/04/2007

MONTHLY

<=[PREV] | [NEXT]=>

Oct 2007

Nov 2007

YEARLY

2007

Google
アップル (160)
Yahoo! JAPAN
Amazon.co.jp
misc
News (1235)
Apple
RTT value from AIT, Thailand to ROOT : DAILY 12/03/2007
RTT value from AIT, Thailand to ROOT : DAILY 12/04/2007
RTT value from AIT, Thailand to ROOT : MONTHLY 10/2007
RTT value from AIT, Thailand to ROOT : MONTHLY 11/2007
RTT value from AIT, Thailand to ROOT : YEARLY 2007
A B C D E F G H I J K L M
Probe Date/Time (JST)

Gulliver Project (3/4)

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

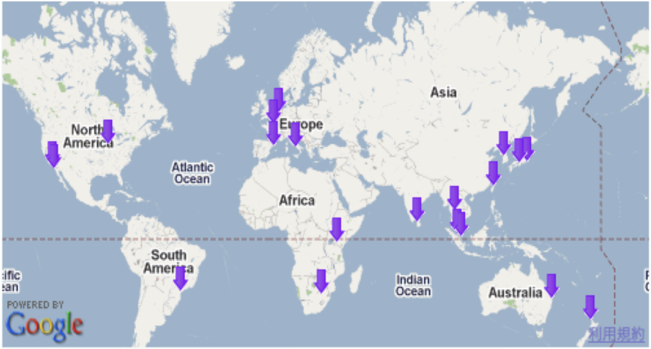

gulliver probe locations

Gulliver Project (4/4)

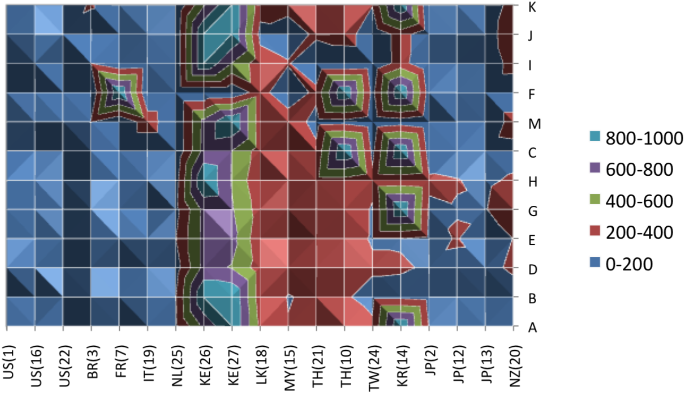

gulliver: measured RTT to DNS root servers (2008)

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)

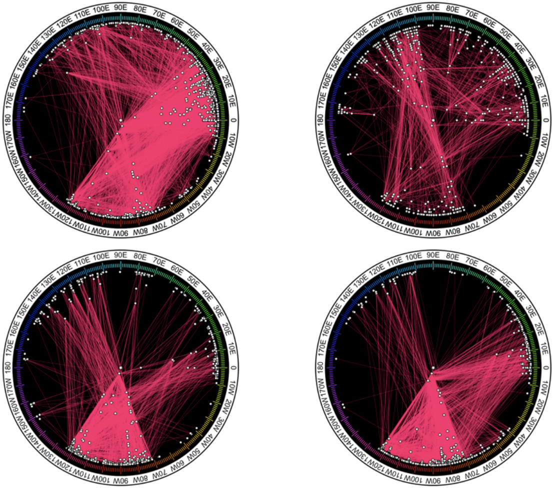

regional topology visualization by AS Core Map

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)

Guilaume 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

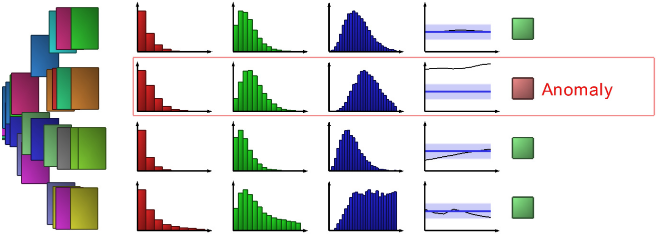

Longitudinal Statistical Analysis based on Robust Estimation (1/3)

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)

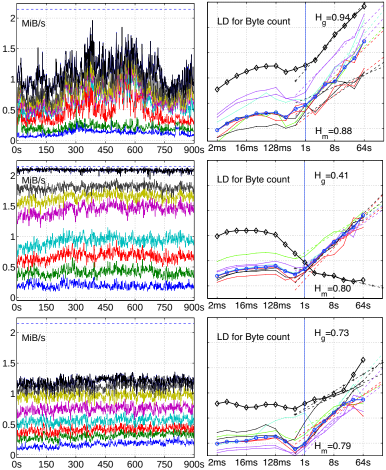

top: no congestion, middle: congestion, bottom: severe anomalies

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

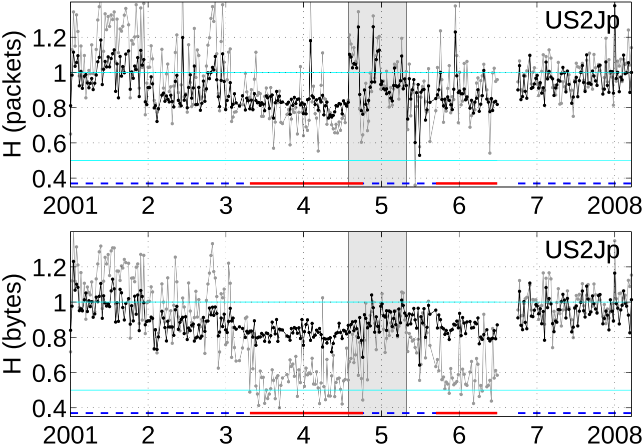

Host Clustering by Communication Patterns (1/3)

- 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 ($\leq$ 144 bytes)
  - % of large packets ($\geq$ 1392 bytes)
  - S(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

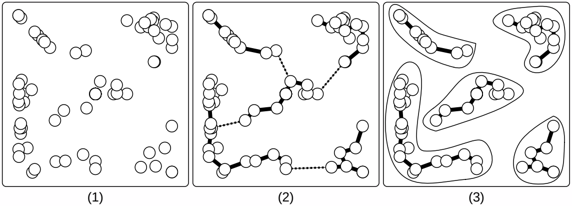

MST clustering illustrated in 2D space

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

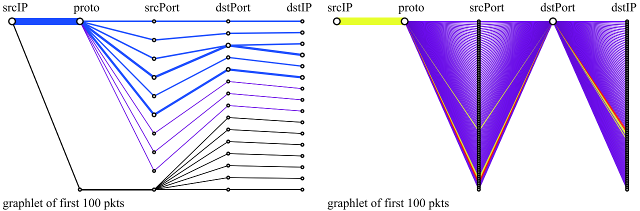

sample graphlets: P2P (left) and scanning (right)

Summary

measurement actitivies 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/

まとめ

まとめ

- これまでのまとめと今後の展望
- WIDE プロジェクトの計測研究