Measurements

SIGCOMM Topic Preview Session 5

Walter Willinger
wwillinger@niksun.com

August 23, 2018

Internet Measurements: **How? Why? What?**

How to do Internet Measurements?

- Ethical considerations
  - Many Internet measurements concern end users (humans)
  - Beyond IRBs: Aim for transparency, obtain user consent
  - SIGCOMM, IMC: Ethical considerations included in CFP
  - New ACM Code of Ethics and Professional Conduct (2018)

- Reproducibility
  - New ACM policy on result and artifact review and badging (2018)
  - "An experimental result is not fully established unless it can be independently reproduced."
  - ACM SIGCOMM 2017 Reproducibility Workshop

Why Internet Measurements?

- They aid scientific discovery
  - "To measure is to know." *(Lord Kelvin)*
- They are the "bread and butter" for network management
  - "You can't manage what you don't measure." (*W. E. Deming)*
- They help us reveal "unintended consequences" of our system/protocol designs
  - "I didn't think; I experimented." *(Wilhelm Roentgen)*

What Internet Measurements?
The Case of Static Data

What Internet Measurements?
The Case of Static Data

What Internet Measurements?
The Case of Static Data

- The type of data we are all used to …
  - A fixed, given dataset (from own or third-party experiments)
  - Multiple passes over the data
  - Centralized storage (lots of available memory)
  - Centralized analysis (lots of available CPU)

- The type of data we all know how to analyze ...
  - Statistics
  - Machine learning, deep learning, …
  - Visualization

What Internet Measurements?
The Case of Dynamic Data

What Internet Measurements?
The Case of Dynamic Data

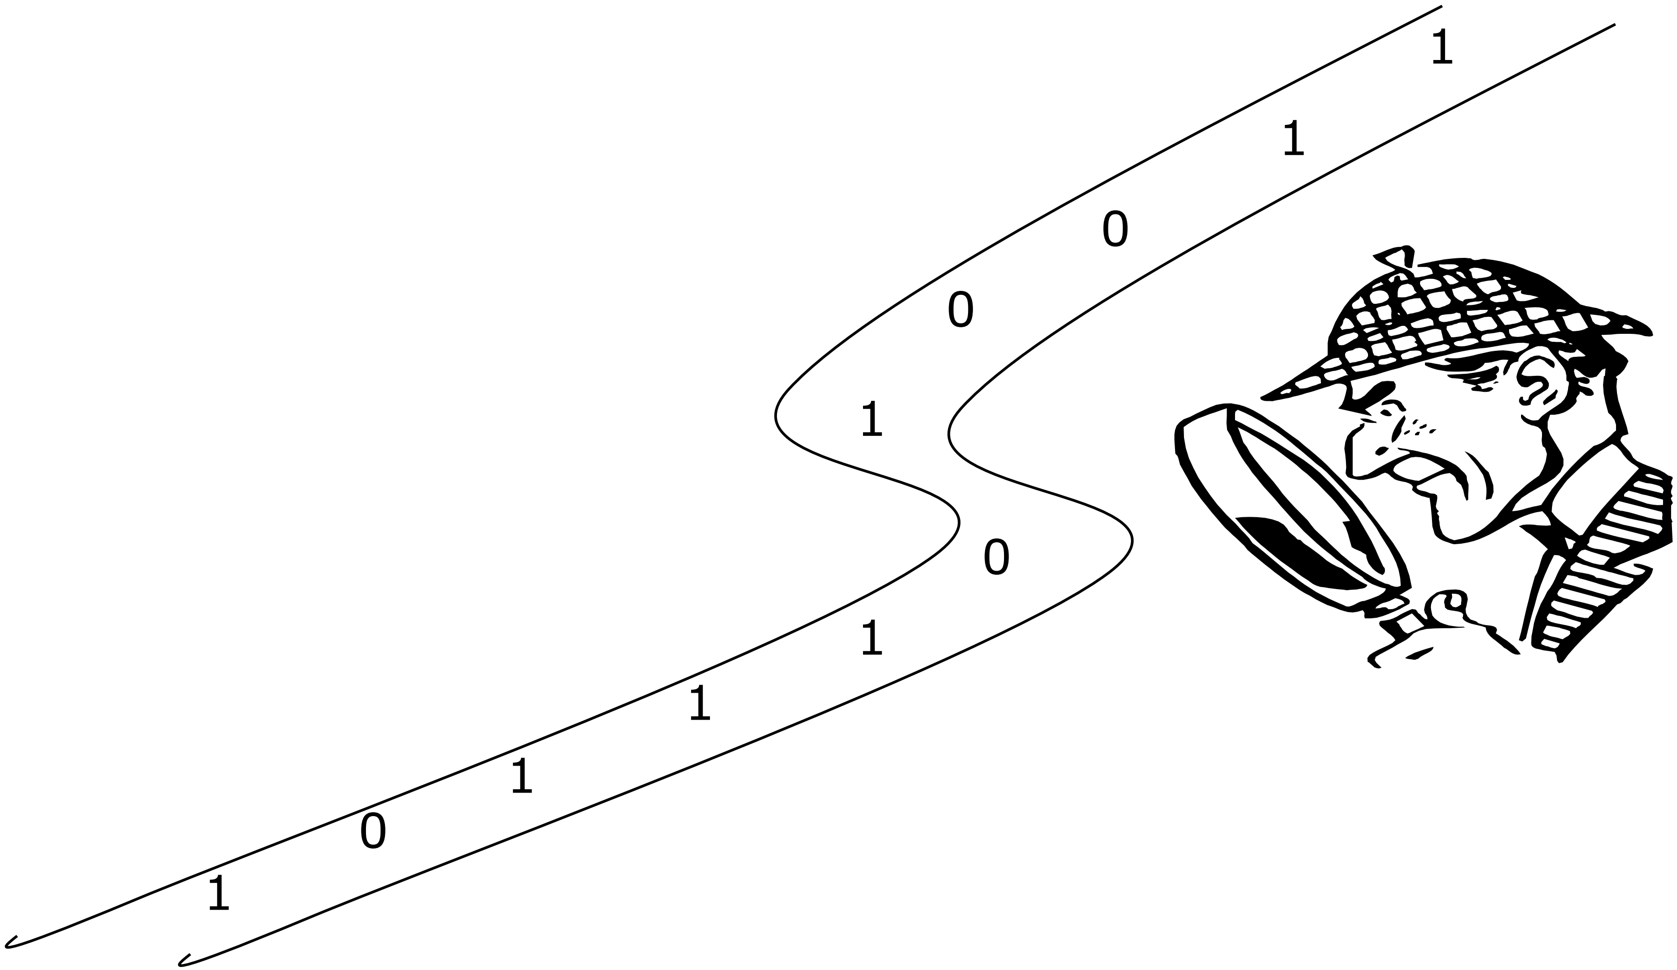

What Internet Measurements?
The Case of Dynamic Data

- The future of Internet measurement - streaming data
  - Append at the "front" of steam only, delete at the "end" only
  - "One-pass" – only one chance to look at/use the current record/packet, at the time the record/packet is captured
  - 'Big (Internet) data" – high volume, velocity, variety, and value

- The type of data most of us don't know how to analyze
  - Small space – log or polylog in data stream size
  - Small time (speed is of essence for real-time decision making)
  - One-pass

What Internet Measurements?
The Case of Distributed Dynamic Data

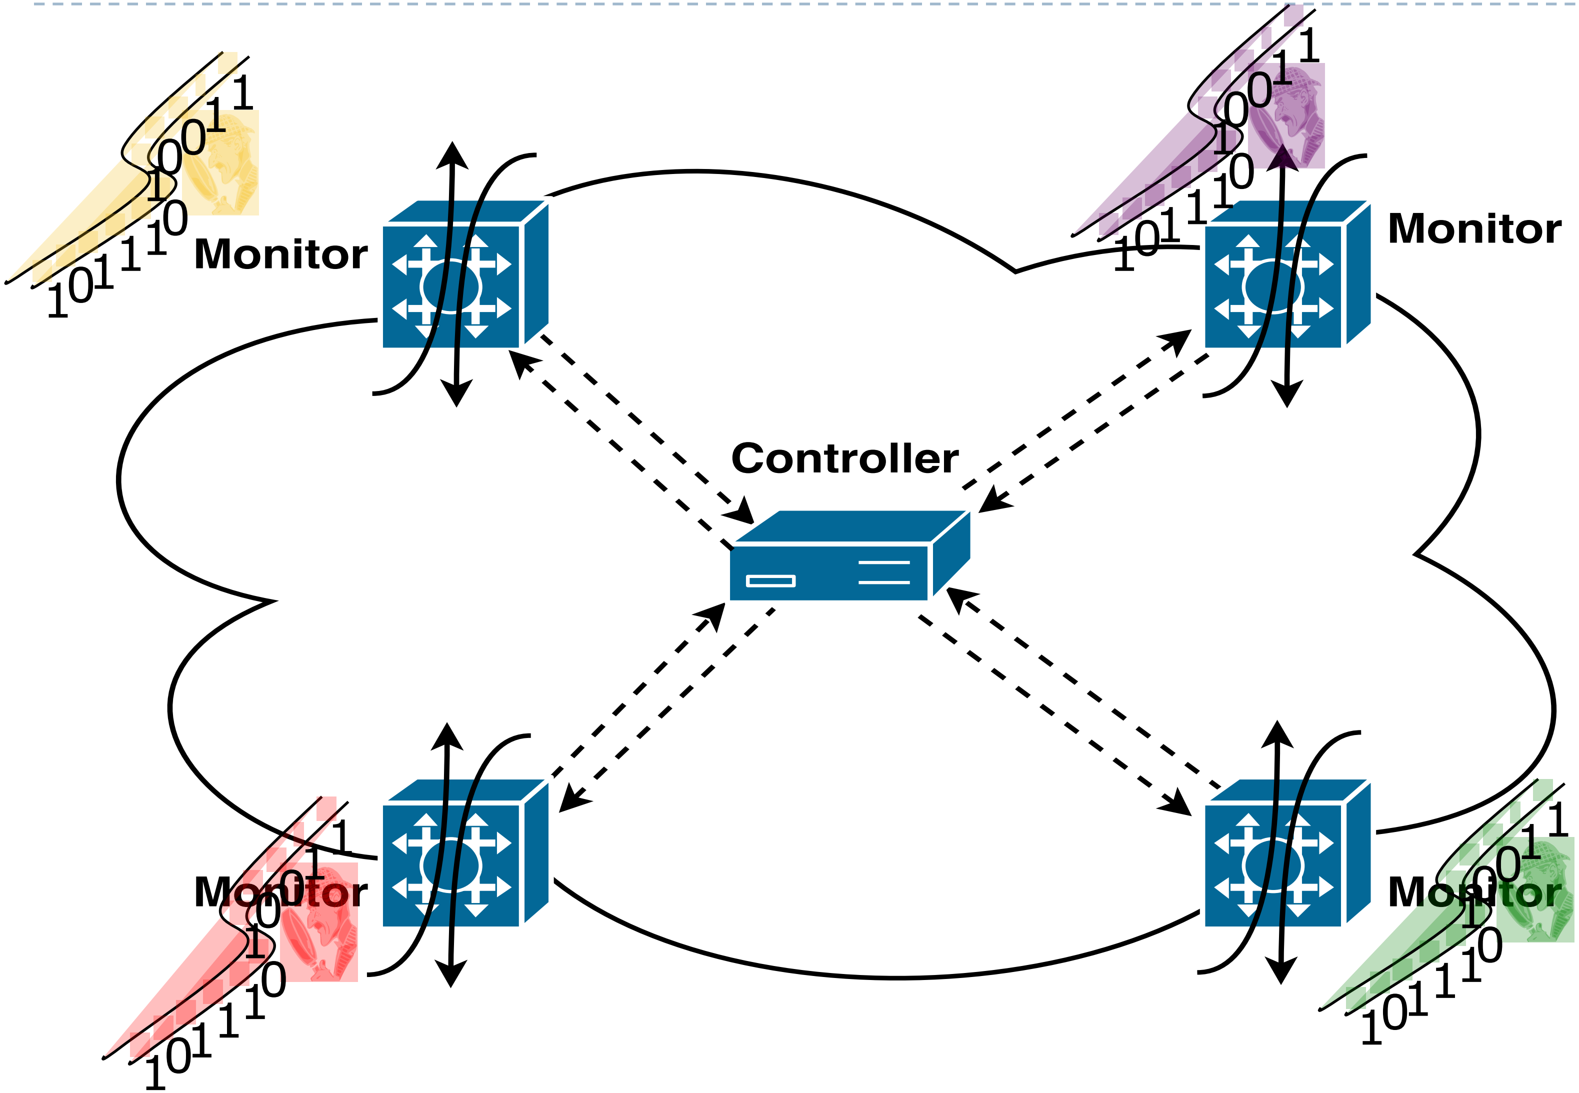

What Internet Measurements?
The Case of Distributed Dynamic Data

- Requirements for analyzing distributed streaming data
  - Each remote location sees its own set of streaming data
  - "One-pass" over streaming data in remote location
  - "Beefy" (i.e., resource-rich) remote nodes
  - "Control & Command"-like communication structure

- Analyzing distributed streaming data: Do's and don'ts
  - Move processing to the data (i.e., remote nodes)
  - Don't move the data to a central location for processing

Analysis of Streaming Data 101

- Complexity arises from massive size of streaming data
- Computing exact answers for streaming data is generally hard
- Randomization to the rescue
  - Approximate answers (based on approximate measurements) are often "good enough" - (ε, δ)-approximations
  - Answer is within $\varepsilon\%$ of correct result with high probability $(1-\delta)$
- Approximate answers from approximate measurements
  - The "old" way: Sampling
  - The "new" way: Sketching

Analysis of Streaming Data 101
Sampling

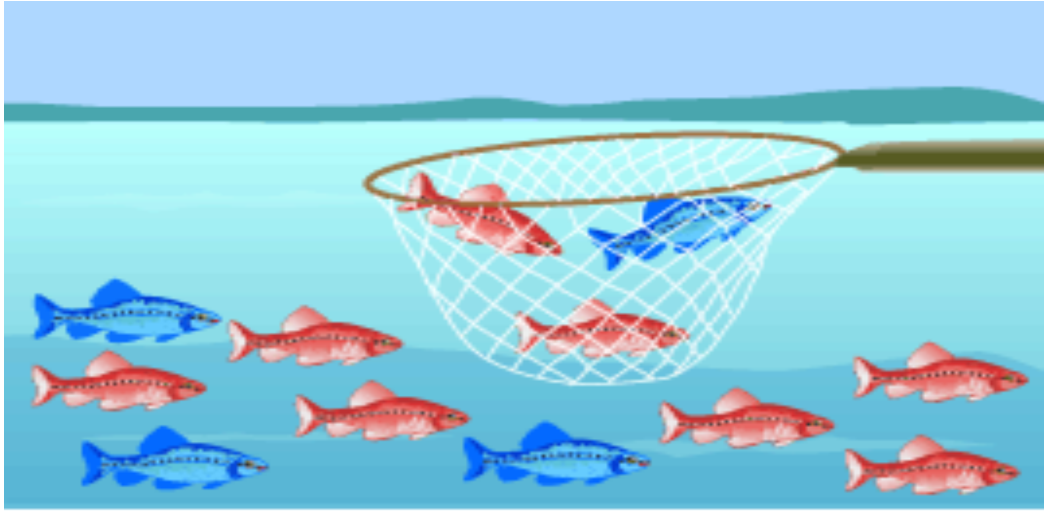

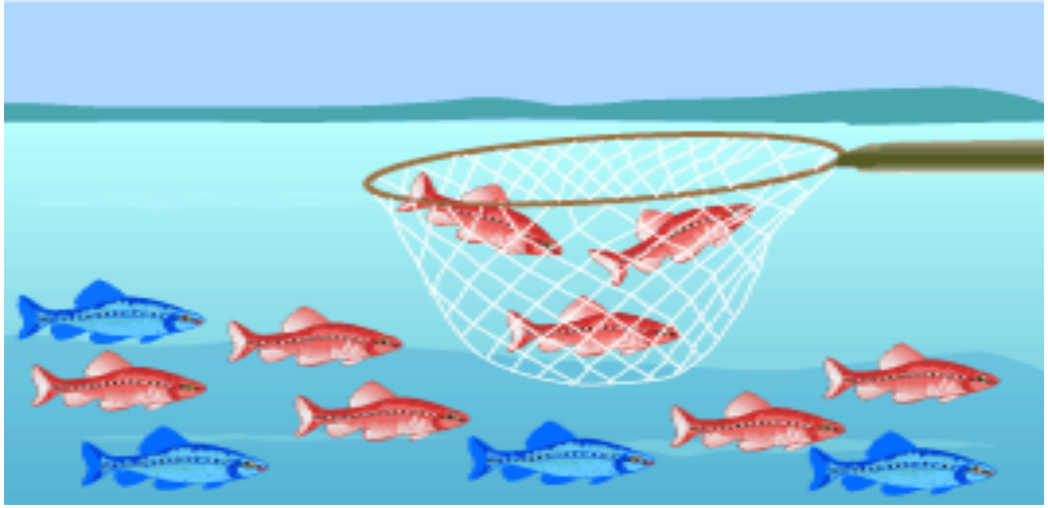

Sampling

- Basic challenges with sampling
  - Sampling techniques don't see all the data!
  - How (i.e., when and how often) to sample k items uniformly from a stream when the length of the stream is not known?
- Not every problem can be solved with sampling
  - Example: counts of distinct items in the stream
  - If a large fraction of items aren't sampled, we don't know if they are all the same or all different …

Analysis of Streaming Data 101
Sketching

Analysis of Streaming Data 101
Sketching

- Sketching techniques have the property that they "see" all the data even if they can't "remember" it all

- Popular example: Count-Min (CM) Sketch
  - A probabilistic data structure that serves as a frequency table of items in a stream of data (random linear projection of data)
  - Randomness enters in the form of hash functions map items to frequencies
  - Number (d) and width (w) of hash functions determine accuracy of answer
- Maintain matrix of size $d \times w$ to obtain (ε,δ)-approximations
  - with $w = 2/\varepsilon$ and $d = \log 1/\delta$ (space-accuracy tradeoff)
  - Single pass, small space, small time (low processing time)

Session 10: Measurements
Thursday, Aug 23, 4:30 pm - 6:10 pm

	Static Data	Dynamic Data
Scientific discovery	RADWAN	
(Real-time) Network Management		Elastic Sketch SketchLearn
Unintended consequences	C-SAW	

Incentivizing Censorship Measurements via Circumvention

Aqib Nisar, Aqsa Kashaf, Ihsan Ayyub Qazi, Zartash Afzal Uzmi*
LUMS, Pakistan

ABSTRACT

We present C-Saw, a system that measures Internet censorship by offering data-driven censorship circumvention to users. The adaptive circumvention capability of C-Saw incentivizes users to opt-in by offering small page load times (PLTs). As users crowdsource, the measurement data gets richer, offering greater insights into censorship mechanisms over a wider region, and in turn leading to even better circumvention capabilities. C-Saw incorporates user consent in its design by measuring only those URLs that a user actually visits. Using a cross-platform implementation of C-Saw, we show that it is effective at collecting and disseminating censorship measurements, selecting circumvention approaches, and optimizing user experience. C-Saw improves the average PLT by up to 48% and 63% over Lantern and Tor, respectively. We demonstrate the feasibility of a large-scale deployment of C-Saw with a pilot study.

CCS CONCEPTS

• **Social and professional topics → Censorship**; *Network access control*; • **Security and privacy** → *Human and societal aspects of security and privacy*; • **Networks** → *Network measurement*;

ACM Reference Format:
Aqib Nisar, Aqsa Kashaf, Ihsan Ayyub Qazi, Zartash Afzal Uzmi. 2018. Incentivizing Censorship Measurements via Circumvention. In *SIGCOMM '18: SIGCOMM 2018, August 20–25, 2018, Budapest, Hungary*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3230543.3230568

*Aqsa Kashaf is currently with CMU and Aqib Nisar with USC.

ACM ISBN 978-1-4503-5567-4/18/08…$15.00
https://doi.org/10.1145/3230543.3230568

1 INTRODUCTION

Internet censorship has become increasingly pervasive with nearly 70 countries restricting Internet access to their citizens [38]. It can have a substantial impact on various stakeholders in the Internet ecosystem (e.g., users, content providers, and ISPs) [43] and thus, has drawn considerable interest from systems researchers towards building *censorship measurement systems* [26, 27, 37, 48–50]—that aim to ascertain *what* is blocked, *where* it is blocked, *how* it is blocked, and *when* it is blocked?—as well as *circumvention systems* that aim to bypass censorship [10, 19, 20, 29, 35].

Unfortunately, existing censorship measurement and circumvention systems are designed *independently*, which results in their individual designs to have limited capabilities. For example, building an effective circumvention system requires understanding the capabilities of censors (e.g., what is blocked and how?), which continually evolve over time [52]. However, existing circumvention systems (e.g., Tor [29], Lantern [10], and uProxy [19]) are not driven by such measurements and thus cannot adapt to the deployed blocking mechanism used by censors. This leads to *one-size-fits-all* solutions that are either ineffective (e.g., against some type of blocking) or inefficient, leading to high page load times (PLTs)[1] and thus degrading user experience [22, 24]. On the other hand, an effective censorship measurement system requires deployment of geographically distributed probe points—possibly volunteers—who find little to no incentive, beyond altruism, to help collect continuous measurements [37, 50].

In this paper, we call for *consolidating* censorship measurements and circumvention in a single platform to address the limitations of individual systems. To this end, we propose C-Saw, a system that gathers continuous censorship measurements through crowdsourcing and uses these measurements to offer data-driven circumvention by being *adaptive* to the type of blocking, leading to improved circumvention performance. This, in turn, incentivizes users to opt-in and contribute measurements. As more users crowdsource, the measurement data gets richer, enabling C-Saw to build a comprehensive database of blocked URLs along with the blocking mechanisms used by censoring ISPs. This allows C-Saw to offer even better circumvention capabilities for

[1] Akhoondi et al. [22] reported that Tor can increase latency by more than 5× compared to the direct path.

Thursday, 4:30pm Session

Incentivizing Censorship Measurements via Circumvention

- An example of an unintended consequence of systems/protocol design: Censorship
- Based on a small user study in Pakistan (static data)
- C-Saw is a system that combines censorship measurements with circumvention capabilities
- Ethical considerations are of critical importance

Thursday, 4:30pm Session

RADWAN: Rate Adaptive Wide Area Network

- Aids scientific discovery: boosting throughput on optical links is feasible in practice
- Uses SNR data from 8K links in Microsoft's optical (static data)
- RADWAN is a traffic engineering system that dynamically adapts the rates of optical links to enhance network throughput and availability
- Networking meets fiber-optical communication
- Reproducibility vs proprietary network measurement

RADWAN: Rate Adaptive Wide Area Network

Rachee Singh
rachee@cs.umass.edu
University of Massachusetts, Amherst

Manya Ghobadi
mgh@microsoft.com
Microsoft Research

Klaus-Tycho Foerster
klaus-tycho.foerster@univie.ac.at
University of Vienna

Mark Filer
mafiler@microsoft.com
Microsoft

Phillipa Gill
phillipa@cs.umass.edu
University of Massachusetts, Amherst

ABSTRACT

Fiber optic cables connecting data centers are an expensive but important resource for large organizations. Their importance has driven a conservative deployment approach, with redundancy and reliability baked in at multiple layers. In this work, we take a more aggressive approach and argue for adapting the capacity of fiber optic links based on their signal-to-noise ratio (SNR). We investigate this idea by analyzing the SNR of over 8,000 links in an optical backbone for a period of three years. We show that the capacity of 64% of 100 Gbps IP links can be augmented by at least 75 Gbps, leading to an overall capacity gain of over 134 Tbps. Moreover, adapting link capacity to a lower rate can prevent up to 25% of link failures. Our analysis shows that using the same links, we get higher capacity, better availability, and 32% lower cost per gigabit per second. To accomplish this, we propose RADWAN, a traffic engineering system that allows optical links to adapt their rate based on the observed SNR to achieve higher throughput and availability while minimizing the churn during capacity reconfigurations. We evaluate RADWAN using a testbed consisting of 1,540 km fiber with 16 amplifiers and attenuators. We then simulate the throughput gains of RADWAN at scale and compare them to the gains of state-of-the-art traffic engineering systems. Our data-driven simulations show that RADWAN improves the overall network throughput by 40% while also improving the average link availability.

SIGCOMM '18, August 20–25, 2018, Budapest, Hungary

ACM ISBN 978-1-4503-5567-4/18/08…$15.00
https://doi.org/10.1145/3230543.3230570

CCS CONCEPTS

• **Networks** → *Physical links*; *Traffic engineering algorithms*; *Network economics*; *Network performance analysis*; *Network reliability*; *Wired access networks*;

KEYWORDS

Traffic Engineering, Wide Area Networks, Optical Backbone

ACM Reference Format:
Rachee Singh, Manya Ghobadi, Klaus-Tycho Foerster, Mark Filer, and Phillipa Gill. 2018. RADWAN: Rate Adaptive Wide Area Network. In *SIGCOMM '18: SIGCOMM 2018, August 20–25, 2018, Budapest, Hungary.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3230543.3230570

1 INTRODUCTION

Optical backbones are million-dollar assets, with fiber comprising their most expensive component. Companies like Google, Microsoft, and Facebook purchase or lease fiber to support wide-area connectivity between distant data center locations but have not been able to fully leverage this investment because of the conservative provisioning of the optical network. We show that wide area fiber links exhibit significantly better signal quality (measured by the signal-to-noise-ratio or SNR) than the minimum required to support transmission at 100 Gbps, leaving money on the table in terms of link capacities.

In other words, there is potential to operate fiber links at higher capacity, thereby increasing the throughput of existing optical networks. We analyze historical SNR from 8,000 optical channels in a backbone network and find that the capacity of 64% of the links can be augmented by 75 Gbps or more, leading to a capacity gain of over 134 Tbps in the network. However, we argue that simply raising link capacities to a higher value (*e.g.,* 150 Gbps or 200 Gbps) increases the rate of link failures because the signal quality fluctuates and operating near the SNR threshold makes links susceptible to failure.

Moreover, enforcing a static link capacity forces operators to treat link failures as *binary* events: when the SNR of a link

Elastic Sketch: Adaptive and Fast Network-wide Measurements

Tong Yang
Peking University
yangtongemail@gmail.com

Jie Jiang
Peking University
jie.jiang@pku.edu.cn

Peng Liu
Peking University
liu.peng@pku.edu.cn

Qun Huang
Institute Of Computing Technology, CAS
huangqun@ict.ac.cn

Junzhi Gong
Peking University
gongjunzhi@pku.edu.cn

Yang Zhou
Peking University
zhou.yang@pku.edu.cn

Rui Miao
Alibaba Group
miao.rui@alibaba-inc.com

Xiaoming Li
Peking University
lxm@pku.edu.cn

Steve Uhlig
Queen Mary University of London
steve.uhlig@quml.ac.uk

ABSTRACT

When network is undergoing problems such as congestion, scan attack, DDoS attack, *etc.*, measurements are much more important than usual. In this case, traffic characteristics including available bandwidth, packet rate, and flow size distribution vary drastically, significantly degrading the performance of measurements. To address this issue, we propose the Elastic sketch. It is adaptive to currently traffic characteristics. Besides, it is generic to measurement tasks and platforms. We implement the Elastic sketch on six platforms: P4, FPGA, GPU, CPU, multi-core CPU, and OVS, to process six typical measurement tasks. Experimental results and theoretical analysis show that the Elastic sketch can adapt well to traffic characteristics. Compared to the state-of-the-art, the Elastic sketch achieves 44.6 ~ 45.2 times faster speed and 2.0 ~ 273.7 smaller error rate.

CCS CONCEPTS

• **Networks** → **Network monitoring**; *Network measurement*;

SIGCOMM '18, August 20–25, 2018, Budapest, Hungary

ACM ISBN 978-1-4503-5567-4/18/08.
https://doi.org/10.1145/3230543.3230544

KEYWORDS

Sketches; Network measurements; Elastic; Compression; Generic

ACM Reference Format:

Tong Yang, Jie Jiang, Peng Liu, Qun Huang, Junzhi Gong, Yang Zhou, Rui Miao, Xiaoming Li, and Steve Uhlig. 2018. Elastic Sketch: Adaptive and Fast Network-wide Measurements. In *SIGCOMM '18: SIGCOMM 2018, August 20–25, 2018, Budapest, Hungary.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3230543.3230544

1 INTRODUCTION

1.1 Background and Motivation

Network measurements provide indispensable information for network operations, quality of service, capacity planning, network accounting and billing, congestion control, anomaly detection in data centers and backbone networks [1–9]. Recently, sketch-based solutions[1] [8, 10] have been widely accepted in network measurements [2, 3, 11, 12], thanks to their higher accuracy compared to sampling methods [2, 4, 12] and their speed.

Existing measurement solutions [4, 8, 10, 12–17] mainly focus on a good trade-off among accuracy, speed and memory usage. The state-of-the-art UnivMon [2] pays attention to an additional aspect, generality, namely using one sketch to process many tasks, and makes a good trade-off among these four dimensions. Although existing work has made great contributions, they do not focus on one fundamental need: achieving accurate network measurements no matter how traffic characteristics (including available bandwidth, flow size distribution, and packet rate) vary. Measurements are especially important when network is undergoing problems,

[1] In this paper, sketches refers to data streaming algorithms that can be used for network measurements.

Thursday, 4:30pm Session

Elastic Sketch: Adaptive and Fast Network-wide Measurements

- Analyzes streaming data in the form of packet-level traffic traces for different network management tasks
- Elastic Sketch is a new sketching technique designed to be adaptive to changing traffic conditions.
- Leverages sketches in a distributed environment (i.e., network-wide setting)
- Shows an implementation of Elastic Sketch in both hardware and software

SketchLearn: Relieving User Burdens in Approximate Measurement with Automated Statistical Inference

Qun Huang[†], Patrick P. C. Lee[‡], and Yungang Bao[†]

[†]State Key Lab of Computer Architecture, Institute of Computing Technology, Chinese Academy of Sciences
[‡]Department of Computer Science and Engineering, The Chinese University of Hong Kong

ABSTRACT

Network measurement is challenged to fulfill stringent resource requirements in the face of massive network traffic. While approximate measurement can trade accuracy for resource savings, it demands intensive manual efforts to configure the right resource-accuracy trade-offs in real deployment. Such user burdens are caused by how existing approximate measurement approaches inherently deal with resource conflicts when tracking massive network traffic with limited resources. In particular, they tightly couple resource configurations with accuracy parameters, so as to provision sufficient resources to bound the measurement errors. We design SketchLearn, a novel sketch-based measurement framework that resolves resource conflicts by learning their statistical properties to eliminate conflicting traffic components. We prototype SketchLearn on OpenVSwitch and P4, and our testbed experiments and stress-test simulation show that SketchLearn accurately and automatically monitors various traffic statistics and effectively supports network-wide measurement with limited resources.

CCS CONCEPTS

• **Networks → Network measurement**;

KEYWORDS

Sketch; Network measurement

ACM Reference Format:
Qun Huang, Patrick P. C. Lee, and Yungang Bao. 2018. SketchLearn: Relieving User Burdens in Approximate Measurement with Automated Statistical Inference. . In *Proceedings of SIGCOMM '18*. ACM, New York, NY, USA, 15 pages. https://doi.org/http://dx.doi.org/10.1145/3098822.3098831

SIGCOMM '18, August 20–25, 2018, Budapeset, Hungary
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-4653-5/17/08…$15.00
https://doi.org/http://dx.doi.org/10.1145/3098822.3098831

1 INTRODUCTION

Network measurement is indispensable to modern network management in clouds and data centers. Administrators measure a variety of traffic statistics, such as per-flow frequency, to infer the key behaviors or any unexpected patterns in operational networks. They use the measured traffic statistics to form the basis of management operations such as traffic engineering, performance diagnosis, and intrusion prevention. Unfortunately, measuring traffic statistics is non-trivial in the face of massive network traffic and large-scale network deployment. Error-free measurement requires per-flow tracking [15], yet today's data center networks can have thousands of concurrent flows in a very small period from 50ms [2] down to even 5ms [56]. This would require tremendous resources for performing per-flow tracking.

In view of the resource constraints, many approaches in the literature leverage approximation techniques to trade between resource usage and measurement accuracy. Examples include sampling [9, 37, 64], top-*k* counting [5, 43, 44, 46], and sketch-based approaches [18, 33, 40, 42, 58], which we collectively refer to as *approximate measurement* approaches. Their idea is to construct compact sub-linear data structures to record traffic statistics, backed by theoretical guarantees on how to achieve accurate measurement with limited resources. Approximate measurement has formed building blocks in many state-of-the-art network-wide measurement systems (e.g., [32, 48, 55, 60, 62, 67]), and is also adopted in production data centers [31, 68].

Although theoretically sound, existing approximate measurement approaches are inconvenient for use. In such approaches, massive network traffic competes for the limited resources, thereby introducing measurement errors due to *resource conflicts* (e.g., multiple flows are mapped to the same counter in sketch-based measurement). To mitigate errors, sufficient resources must be provisioned in approximate measurement based on its theoretical guarantees. Thus, *there exists a tight binding between resource configurations and accuracy parameters.* Such tight binding leads to several practical limitations (see §2.2 for details): (i) administrators need

Thursday, 4:30pm Session

SketchLearn: Relieving User Burdens in Approximate Measurement with Automated Statistical Inference

- Analyzes streaming data in the form of packet-level traffic traces for different network management tasks
- Introduces SketchLearn, a new multi-level data structure that simplifies the tricky task of configuring sketches such that the obtained (approximate) answers have guaranteed error bounds
- Shows statistical properties of SketchLearn that have practical implications

For further reading

- ACM Artifact Review and Badging(2018) https://www.acm.org/publications/policies/artifact-review-badging
- ACM Code of Ethics and Professional Conduct (2018) https://www.acm.org/code-of-ethics
- V. Paxson, Strategies for sound Internet measurement, Proc. IMC'04
- M. Crovella and B. Krishnamurthy, Internet Measurement: Infrastructure, Traffic and Applications, Wiley (2006)
- C. Partridge and M. Allman, Ethical considerations in network measurement papers, CACM (2016)
- G. Cormode and S. Muthukrishnan. An improved data stream summary: The count-min sketch and its applications. Journal of Algorithms (2005)
- G. Cormode, Count-Min Sketch, http://dimacs.rutgers.edu/~graham/pubs/papers/cmencyc.pdf

Thanks!

QUESTIONS?