NETWORKING TRENDS SCIENCE DMZ: INTRODUCTION, CHALLENGES, AND OPPORTUNITIES

Jorge Crichigno
Department of Integrated Information Technology
College of Engineering and Computing
University of South Carolina

Presentation at Universidad Catolica de Asuncion
Asuncion, Paraguay
August 14, 2019

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed ‘science’ network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers’ buffer size
- Resources online

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed ‘science’ network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers’ buffer size
- Resources online

University of South Carolina

- Founded in 1801
- Flagship state institution
- 350+ programs (BSc, MSc, PhD)
- 50,000 students, over 34,000 in Columbia campus

University of South Carolina

- Founded in 1801
- Flagship state institution
- 350+ programs (BSc, MSc, PhD)
- 50,000 students, over 34,000 in Columbia campus

College of Engineering and Computing

- 3222/570+ undergraduate/graduate students
- 135 TTT faculty
- Research awards
  - Federal agencies, foundations, industry
- Industry partnerships
  - IBM, Boeing, Siemens, Samsung
  - Cisco, Palo Alto Networks, Juniper Networks, Barefoot Networks, VMware, etc.

University of South Carolina

- The College of Engineering and Computing includes:
  - Integrated Information Technology (IIT)
  - Computer Science
  - Electrical Engineering
  - Mechanical Engineering
  - Aerospace Engineering
  - Biomedical Engineering
  - Chemical Engineering
  - Civil and Environmental

Information Technology

- More practical than theoretical in nature
- Promote applied research using professional tools and platforms
- Research agenda emerges from the practice
- Laboratory experiences with workplace relevance

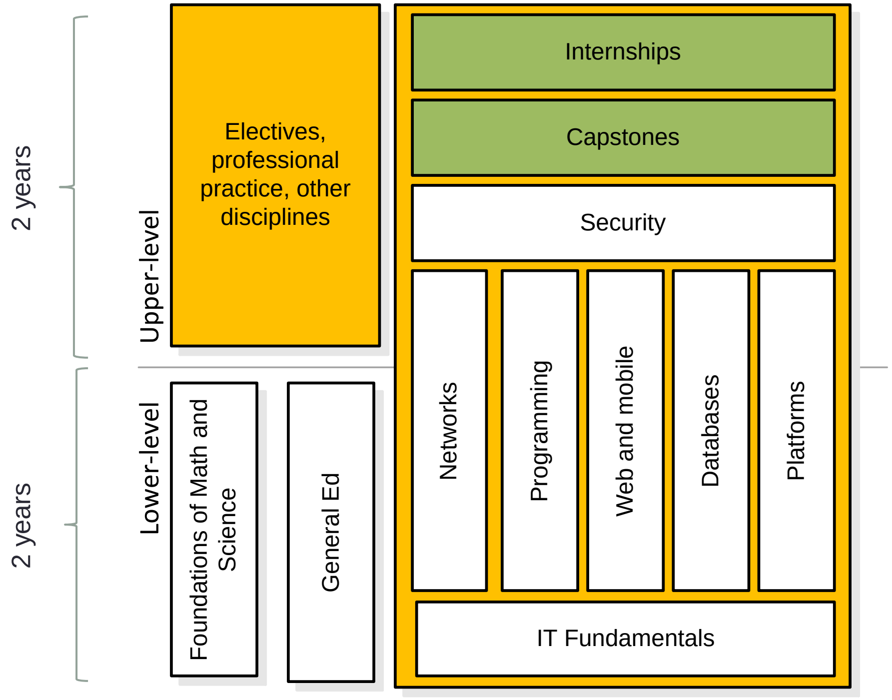

Information Technology

Re: Palo Alto Networks Sales Career Opportunity for Cybersecurity Academy Students!

University of South Carolina
Ph: 803-576-6858

From: Kim Yohannan <kyohannan@paloaltonetworks.com>
Sent: Wednesday, April 24, 2019 7:40:50 PM
Subject: Palo Alto Networks Sales Career Opportunity for Cybersecurity Academy Students!

Dear Faculty,

We are excited to share a Palo Alto Networks career opportunity for recent college and university graduates starting in September 2019. Students who have taken one or more of our Cybersecurity candidates to apply for this position, Business Development Representative – Academy Program if they are interested in a career in sales. This role is part of Inside Sales at Palo Alto Networks and is this role open in the following locations:

- Santa Clara, CA
- Plano, TX
- New York, NY (Empire State Building)

Savannah River NUCLEAR SOLUTIONS℠
FLUOR • NEWPORT NEWS NUCLEAR • HONEYWELL

Savannah River National Laboratory™
OPERATED BY SAVANNAH RIVER NUCLEAR SOLUTIONS

Job details

Auto req ID	4471BR
Job Abbreviation Title	SRNL Industrial Control Systems Security Intern
Job Description	Savannah River National Laboratory (SRNL) is a multi-program laboratory applying state of the art science and practical, high-value, cost-effective solutions to complex techn Department of Energy's (DOE) Savannah River Site (SRS), the laboratory develops and deploys innovative technologies to address some of the nation's environmental, ener Intern will participate in the development of a virtual network which simulates known environments to research vulnerabilities of Industrial Control Systems. The intern may ass through scanning and patching industrial controllers and generating documentation to ensure each system meets SRS cyber security requirements. Additionally the intern ma environments and robotics systems.
Major	Computer Science Other
Other Major	Cyber Security, Industrial Systems, Virtual Reality, Industrial Controls/Robotics
Basic Qualifications (Quantifiable; e.g. Three Years Experience, Bachelors Degree)	Junior or Senior Knowledge and skill in basic computer applications and coding. Pursuing degree in Computer Science, Cyber Security, Industrial Systems, Virtual Reality, Industrial Controls/Robotics or related degree Minimum overall GPA of 2.5 on a 4.0 GPA scale
Preferred Qualifications (e.g. Masters Degree)	Preferred courses: Introduction to Computer Networks Advanced Computer Networks IT Security
Removal Date	22-May-2019

USC's Cyberinfrastructure (CI) Lab

- Information online at http://ce.sc.edu/cyberinfra/
- Development of custom protocols using programmable switches
- TCP rate control using pacing
- Entropy-based intrusion detection
- IoT traffic analysis
- Collaboration in the above topics with
  - University of Texas at San Antonio (UTSA)
  - University of South Florida (USF)
  - U.S. Department of Energy and National Laboratories
  - The Energy Science Network (ESnet)
  - Brandon University (Canada)
  - Brno University (Czech Republic)

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed ‘science’ network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers’ buffer size
- Resources online

Motivation for a High-Speed Science Architecture

- Science and engineering applications are now generating data at an unprecedented rate
- From large facilities to portable devices, instruments can produce hundreds of terabytes in short periods of time
- Data must be typically transferred across high-throughput high-latency Wide Area Networks (WANs)

Applications

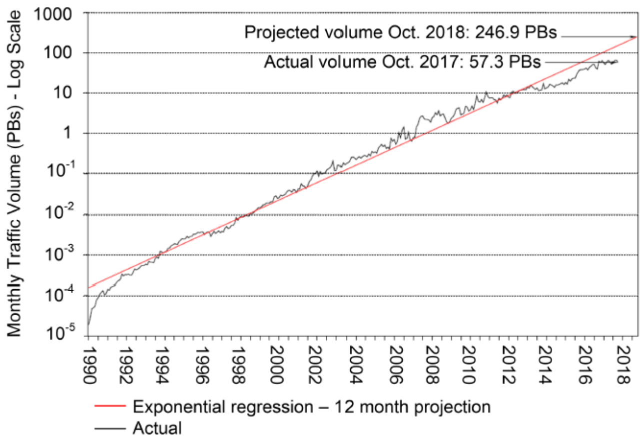

ESnet traffic

The Energy Science Network (ESnet) is the backbone connecting U.S. national laboratories and research centers

Motivation for a High-Speed Science Architecture

- A biology experiment using the U.S. National Energy Research Scientific Computing Center (NERSC) resources

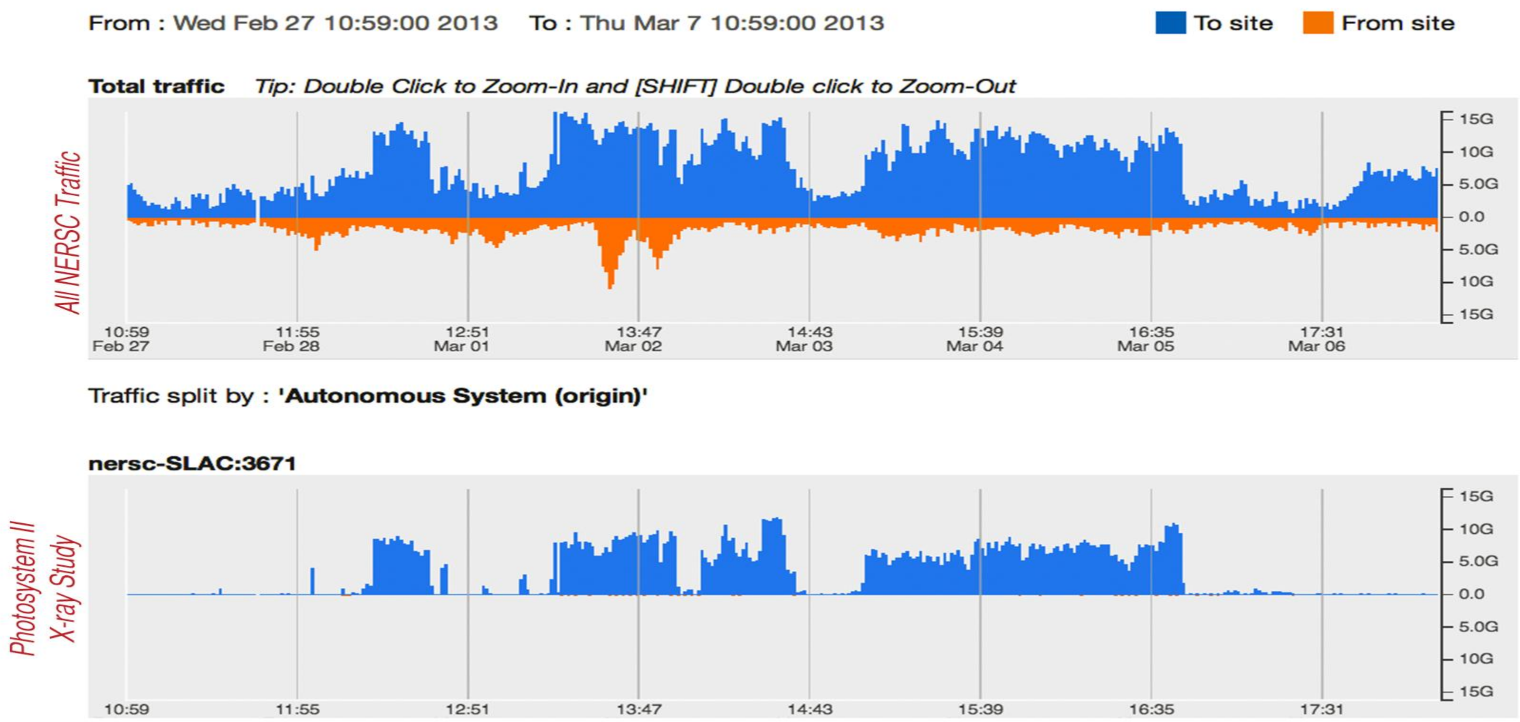

Motivation for a High-Speed Science Architecture

- A biology experiment using the U.S. National Energy Research Scientific Computing Center (NERSC) resources

SnapChat Data produced per day worldwide by millions of people

= 38 TB

One Biology experiment by a team of nine scientists:

= 114 TB

(Photosystem II X-Ray Study)

http://www.nature.com/articles/ncomms5371

Motivation for a High-Speed Science Architecture

Enterprise network limitations:

- Security appliances (IPS, firewalls, etc.) are CPU-intensive
- Inability of small-buffer routers/switches to absorb traffic bursts

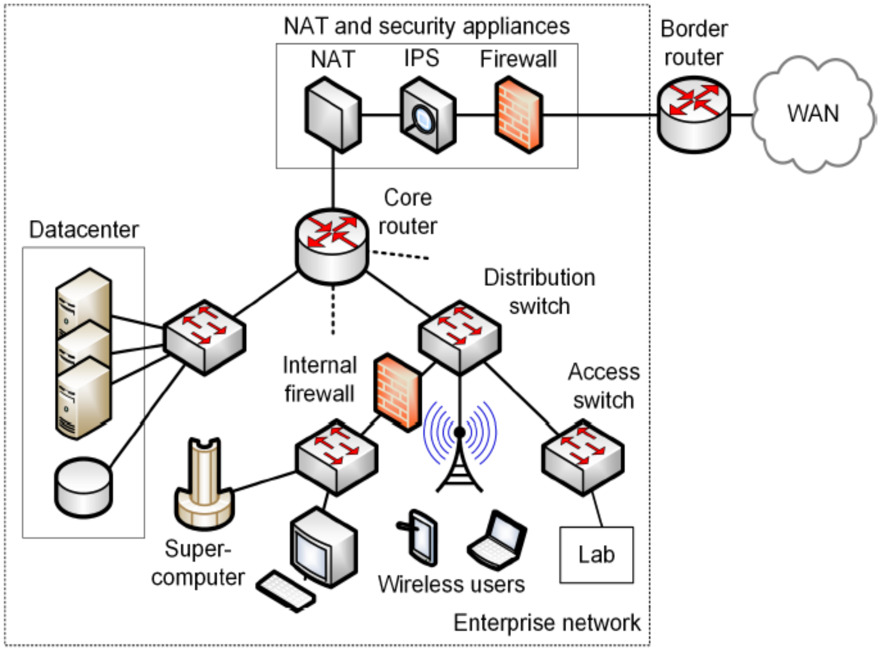

Motivation for a High-Speed Science Architecture

Enterprise network limitations:

- Security appliances (IPS, firewalls, etc.) are CPU-intensive
- Inability of small-buffer routers/switches to absorb traffic bursts
- At best, transfers of big data may last days or even weeks

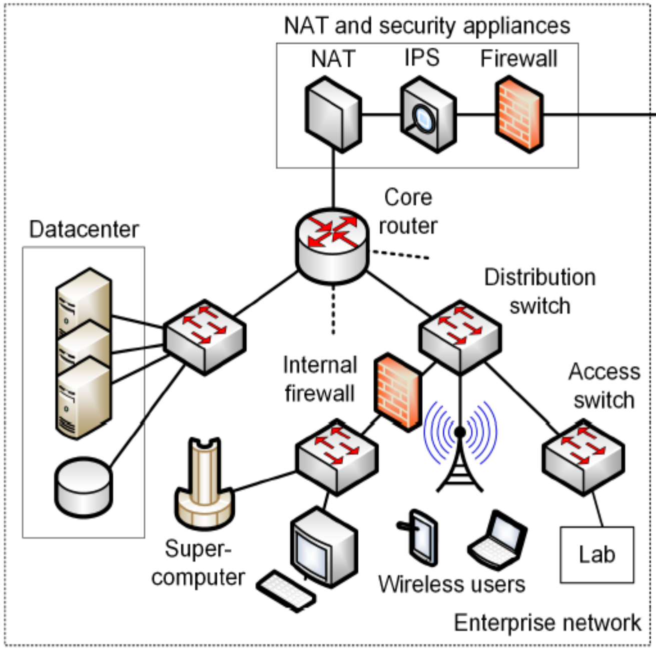

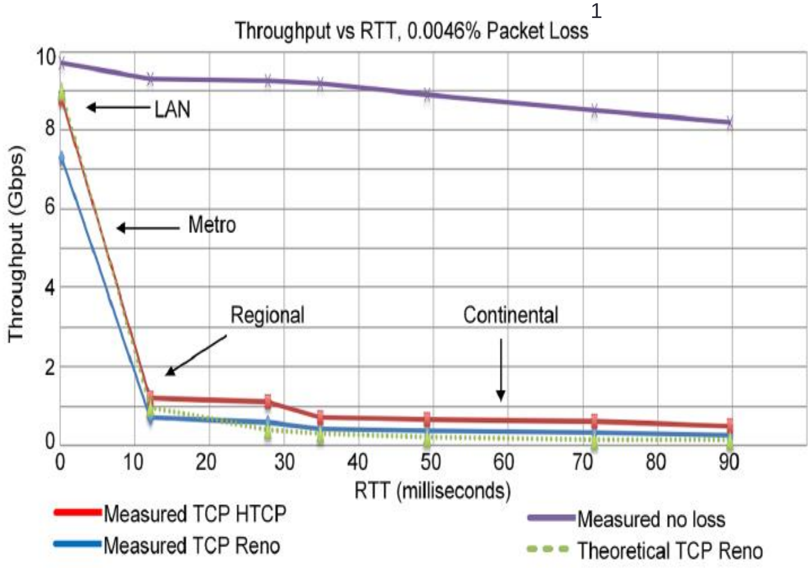

Two devices exchanging data on a 10 Gbps network
Packet loss rate is 1/22,000, or 0.0046%

[1] E. Dart, L. Rotman, B. Tierney, M. Hester, J. Zurawski, "The science dmz: a network design pattern for data-intensive science," *International Conference on High Performance Computing, Networking, Storage and Analysis*, Nov. 2013.

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed ‘science’ network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers’ buffer size
- Resources online

Science DMZ

- The Science DMZ is a network designed for big science data[1, 2]
- Main elements
  - High throughput, friction free WAN paths (no inline security appliances; routers / switches w/ large buffer size)
  - Data Transfer Nodes (DTNs)
  - End-to-end monitoring = perfSONAR
  - Security = access-control list + offline appliance/s (IDS)

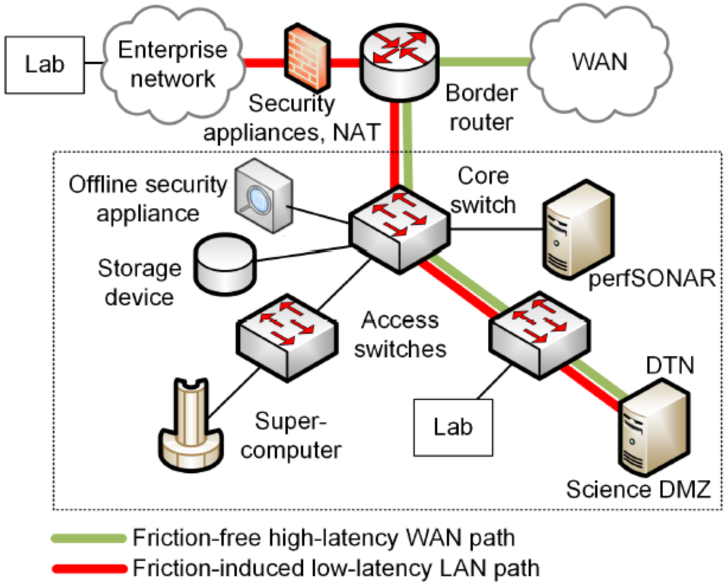

[1]J. Crichigno, E. Bou-Harb, N. Ghani, "A comprehensive tutorial on science DMZ," *IEEE Communications Surveys and Tutorials*, Vol. 21, Issue 2, 2nd quarter 2019.
2. [1]E. Dart, L. Rotman, B. Tierney, M. Hester, J. Zurawski, "The science dmz: a network design pattern for data-intensive science," *International Conference on High Performance Computing, Networking, Storage and Analysis*, Nov. 2013.

Science DMZ

- The Science DMZ is a network designed for big science data[1, 2]
- Main elements
  - ➢ High throughput, friction free WAN paths (no inline security appliances; routers / switches w/ large buffer size)
  - ➢ Data Transfer Nodes (DTNs)
  - ➢ End-to-end monitoring = perfSONAR
  - ➢ Security = access-control list + offline appliance/s (IDS)

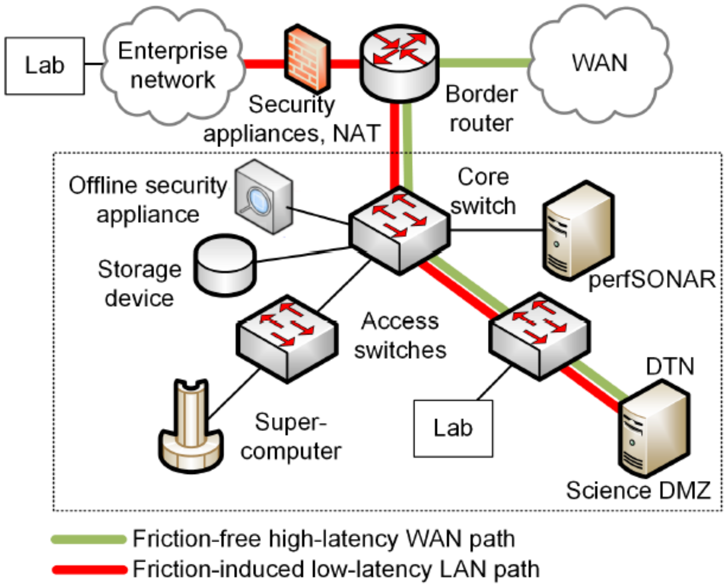

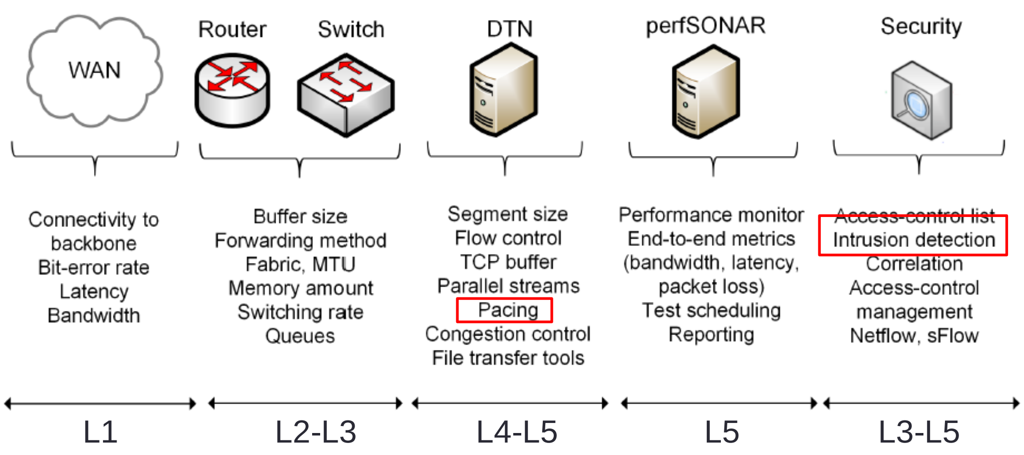

Science DMZ Needs at USC

Researchers	Topic	Current support	Requirements
Gothe Ilieva Strauch	Experimental nuclear physics (ENP)	NSF: 1505615 ($1.2M), 1614773 ($610K), 1812382 ($350K); Brookhaven National Laboratory (BNL) 218624 ($15K); Jefferson Science Associates / DOE ($11K)	100 Gbps throughput to PSI, JLab. High throughput to other collaborators (Brookhaven, Argonne)
Heyden Lauterbach	Chemical engineering	NSF: 1254352 ($400K), 1534260 ($840K), 1565964 ($300K), 1832809 ($160K), 1632824 ($3M), 1805307 ($75K)	High throughput (at least 10 Gbps) to XSEDE (SDSC, TACC), PNNL
Bayoumi	Aerospace, predictive maintenance	Siemens ($628M in-kind [44]), Boeing ($5M [45]), DOD hq017-17-c-7110 ($240K), Missile Def. Ag. HQ0147-16-C-7606 ($35K), Boeing SSOW-BRT-W0915-0001 ($275K)	High throughput with encryption (10 Gbps) to internal and external HPCs, XSEDE, SDSC, TACC
Baalousha Lead	Environment nanoscience	NSF: 1828055 ($635K), 1738340 ($286K), 1655926 (4K), 1553909 ($510K), 1437307 ($300K), 1508931 ($390K), 1834638 ($380K); DOD 450388-19545 ($380K); NIEH 1P01ES028942-01 ($6M), NIH R03ES027406-01 ($144K).	High throughput (5 Gbps) connection from TOF-ICP-MS instrument to Internet2
Sutton Xiaomin Kidane	Digital image correlation (DIC)	NASA C15-2A38-USC ($1.2M), NSF 1537776 ($165K), Boeing SSOW-BRT-W0915-0003 ($140K)	High throughput from USC's DIC laboratory to HPCs (SDSC, TACC) running ABAQUS, ANSYS
Porter	Ntl. Estuarine Research Reserve System	NOAA: NA18NOS4200120 ($760K), NA17NOS4200104 ($980K), OOS.16 (028)USC.DP.MOD.1 ($100K), U. Mich. 3003300692 ($340K), FL Env. Protection CM08P ($92K), NIEHS 1P01ES028942-01 ($6M), USDA ($43K).	High throughput from NOAA's NERRS repository (located at USC) to Internet2 (large datasets downloads worldwide)
Avignone Guiseppe	Particle astrophysics	NSF 1614611 ($900K), NSF 1307204 ($1M), NSF 1808426 ($306K)	100 Gbps connection to MAJORANA (SD), CUORE (Italy), NERSC (CA)
Chandra	Semiconductor material	NSF: 1810116 ($371K), 1711322 ($370K), 1553634 ($695K); NIBIB 1R03EB026813-01 ($136K), DOD W911NF-18-1-0029 ($585K), SRNL/DOE UC150 ($24K), DOE DE-SC0019360 ($666K), RCSA 23976 ($100K)	High throughput (at least 10 Gbps) from X-ray photoelectron spectroscopy instrument and storage to Internet2 (SRNL, INL, Sandia, other institutions)

Science DMZ Needs at USC

Chandra Shustova	Semiconductor material	NSF: 1810116 ($371K), 1711322 ($370K), 1553634 ($695K); NIBIB 1R03EB026813-01 ($136K), DOD W911NF-18-1-0029 ($585K), SRNL/DOE UC150 ($24K), DOE DE-SC0019360 ($666K), RCSA 23976 ($100K)	High throughput (at least 10 Gbps) from X-ray photoelectron spectroscopy instrument and storage to Internet2 (SRNL, INL, Sandia, other institutions)
Richardson Myrick	Phytoplankton spectroscopy	NSF 1542555 ($2M) and DXP Supply Chain Services ($40K)	High throughput (10 Gbps) from image photometer, storage to internal and external HPC
Norman	Genomics data mining	NSF 1149447 ($850K), NIEH 1P01ES028942-01 ($6M), NSF SC EPSCoR 2031-231-2022570 ($100K)	100 Gbps throughput from genomics seq. instrument/storage to USC's HPC; 10+ Gbps connection to Frederick, Argonne, Oak Ridge Ntl. Laboratories, XSEDE resources
Pinckney Benitez	Estuarine ecology	NSF 1736557 ($1M), NOAA R/ER-49 ($130K), NSF 1829519 ($265K), NSF 1458416 ($593K), NSF 1433313 ($362K), NASA 23175500 ($167K)	High throughput from USC's estuarine database to HPCs and Internet2 (datasets downloads)
Dudycha	Genomics, aquatic biology	NSF 1556645 ($1.2M), SC Sea Grant Consortium/NOAA/DOC N250 ($40K), DOD W81XWH1810088 ($287K)	100 Gbps connection to USC's HPC; 10+ Gbps connection to transport DNA / RNA-seq. datasets to XSEDE
Vasquez	Math, genome dynamics	NSF: 1751339 ($290K), 1410047 ($210K)	100 Gbps connection from genomics laboratory to USC's HPC, XSEDE
Brooks Hikmet Schooley	Mathematical models for patient treatment	SC Department of Commerce ($300K), Duke Endowment Child Care Division 1971-SP ($646K), American Cancer Society IRG-17-179-04 ($30K), Patient-Centered Outcomes Research Institute ME-1303-6011 ($960K)	100 Gbps connection from engineering storage to USC's HPC
Ramstad Shervette Ghoshroy	Other USC campuses, genomics	NOAA/DOC NA18NMF4330239 ($503K), NOAA/DOC NA18NMF4270203 ($230K), NOAA NA17NMF4540137 ($153K), NOAA 719583-712683 ($189K), NOAA NA15NMF4330157 ($466K).	10 Gbps connection to move datasets between USC Aiken - Internet2
Crichigno	Cyberinfrast.	NSF 1822567 ($420K), NSF 1829698 ($500K)	100 Gbps programmable network

USC's Science DMZ

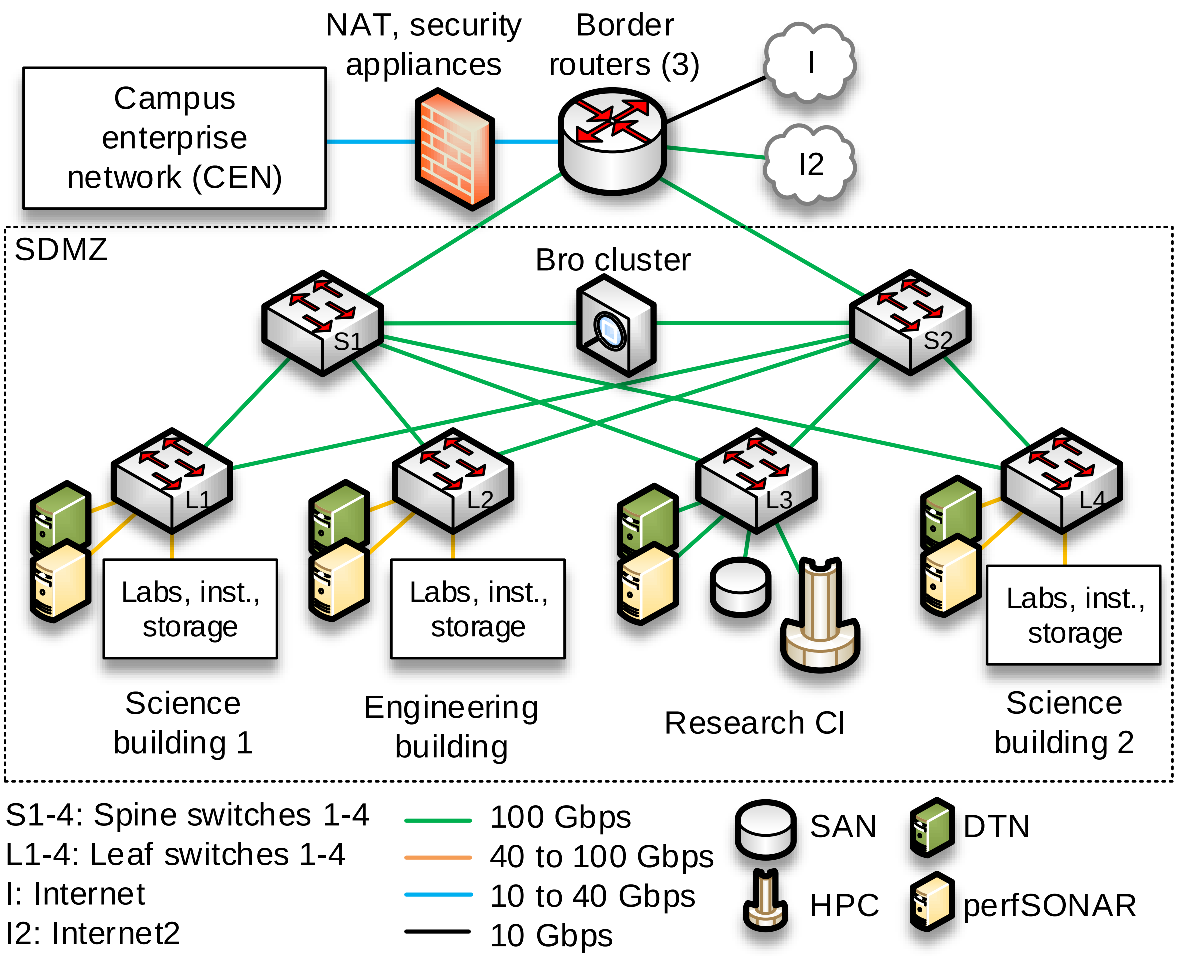

CC* Networking Infrastructure: Building a Science DMZ for Data-intensive Research and Computation at the University of South Carolina, NSF Award # 1925484. Available online at https://nsf.gov/awardsearch/showAward?AWD_ID=1925484&HistoricalAwards=false

U.S. Backbones: Internet2 and ESnet

Internet2

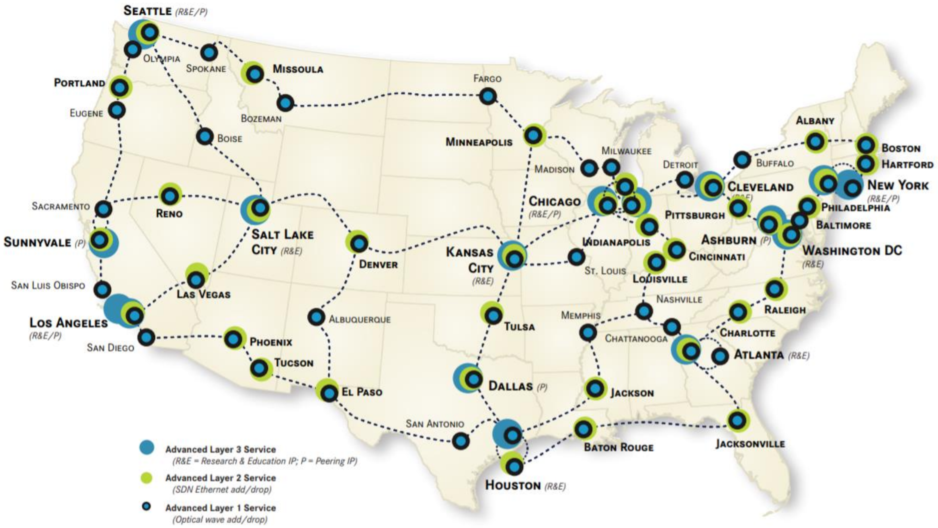

ESnet

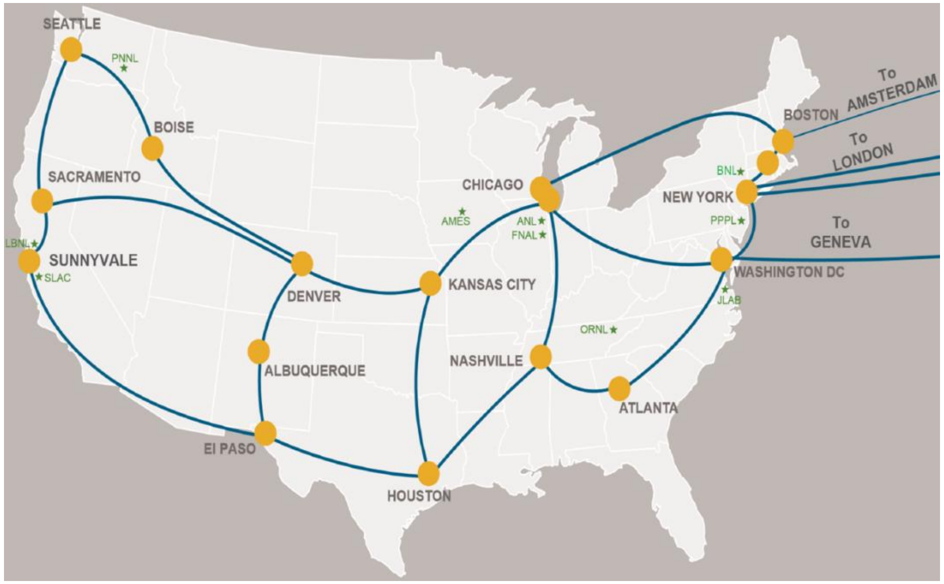

Science DMZs in the U.S.

- Science DMZ deployments, U.S.

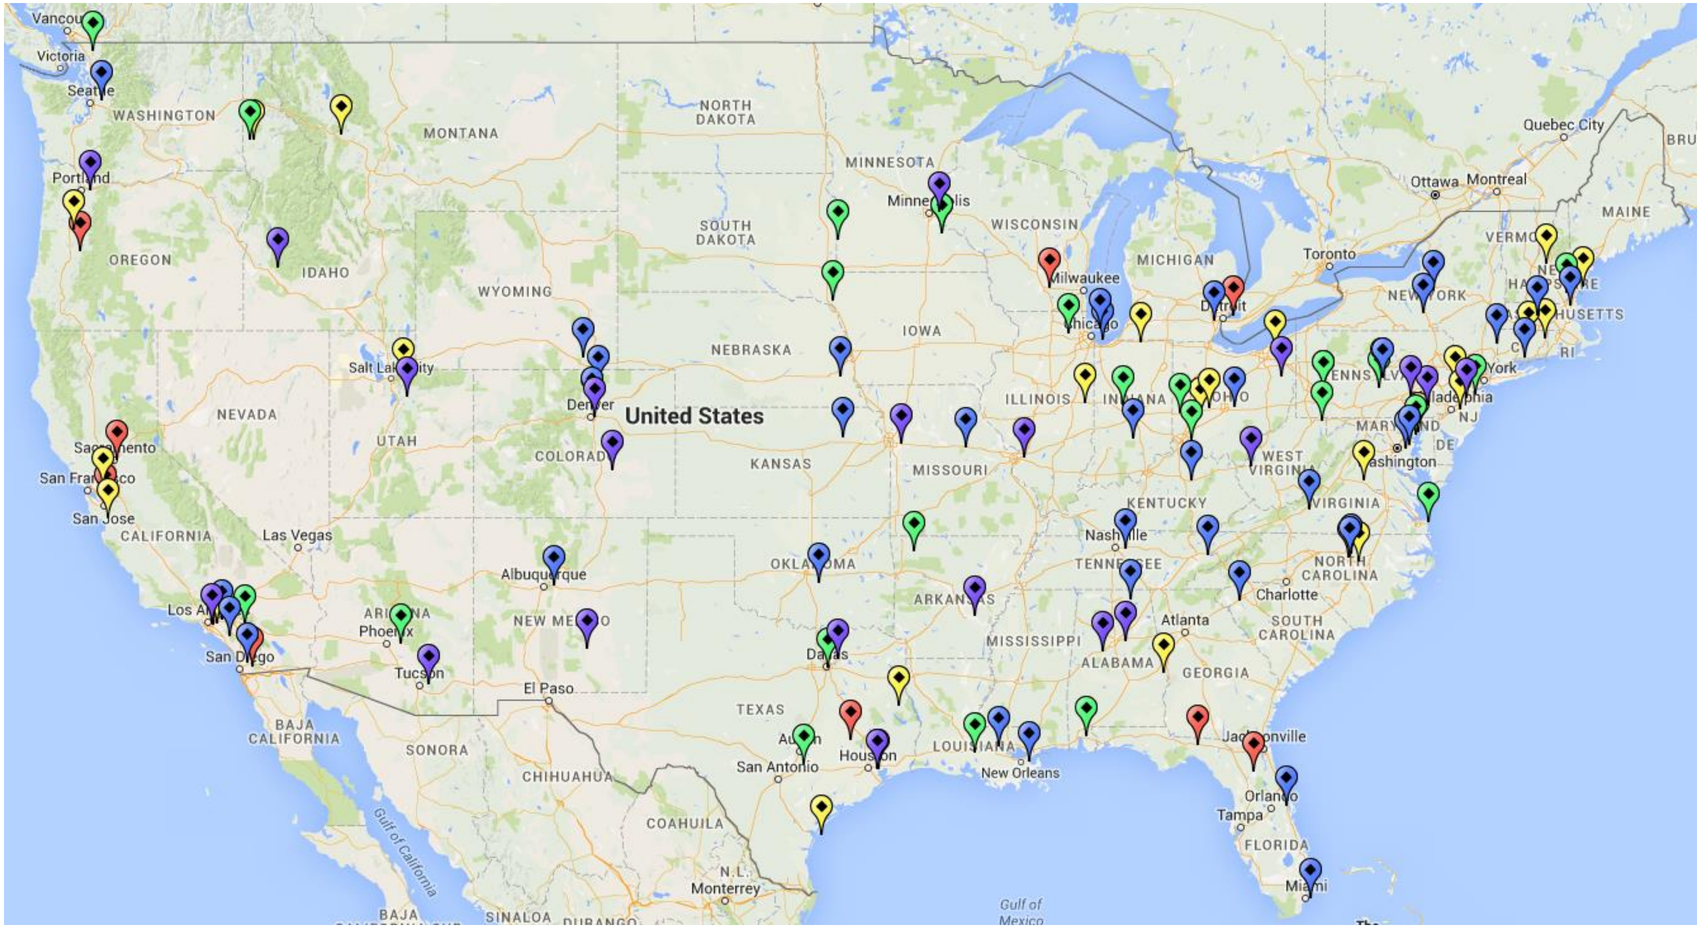

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed 'science' network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers' buffer size
- Resources online

Research Opportunities – Pacing

- Packet loss is expensive in high-throughput high-latency networks

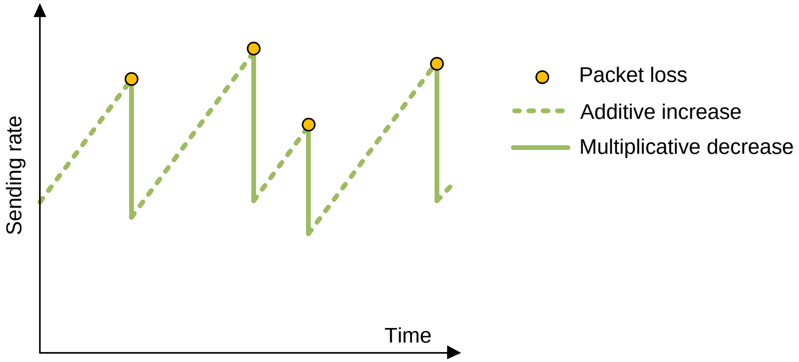

(a) Sawtooth behavior

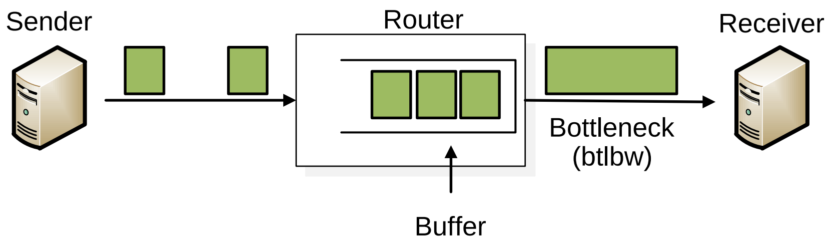

(b) TCP view of a connection

M. Mathis, J. Semke, J. Mahdavi, T. Ott, "The macroscopic behavior of the tcp congestion avoidance algorithm," *ACM Computer Communication Review*, vol. 27, no 3, pp. 67-82, Jul. 1997.

Research Opportunities – Pacing

- Packet loss is expensive in high-throughput high-latency networks

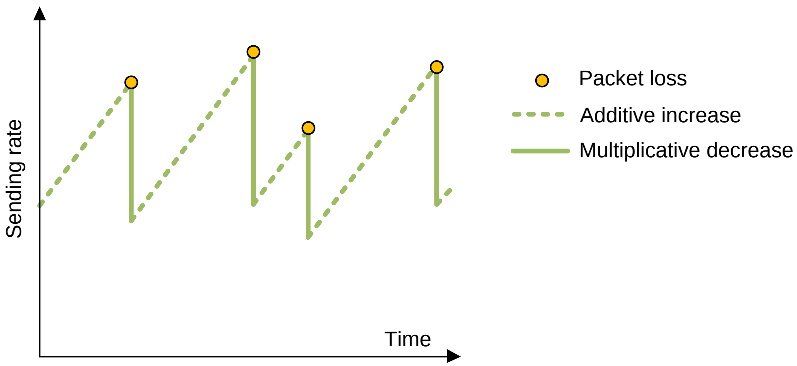

(a) Sawtooth behavior

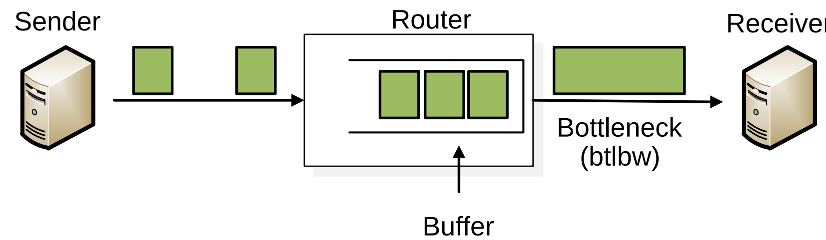

(b) TCP view of a connection

$$\text{TCP throughput} = \frac{c \cdot MSS}{RTT \cdot \sqrt{p}}$$

MSS: maximum segment size
RTT: round-trip time
p: loss rate
c: constant

(c) Average throughput

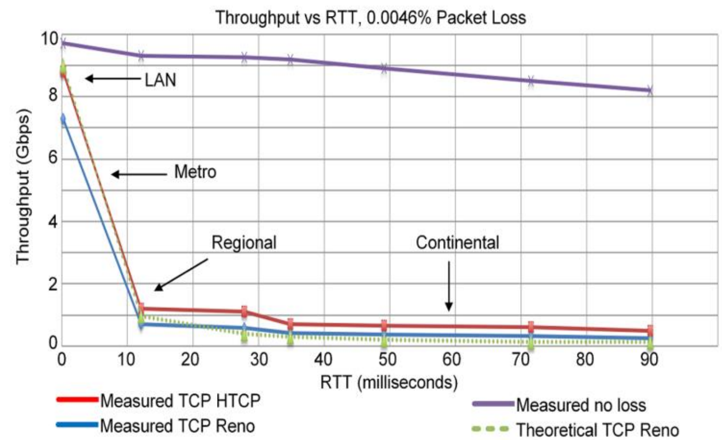

(d) Impact of packet loss

M. Mathis, J. Semke, J. Mahdavi, T. Ott, “The macroscopic behavior of the tcp congestion avoidance algorithm,” *ACM Computer Communication Review*, vol. 27, no 3, pp. 67-82, Jul. 1997.

Pacing

- With TCP pacing, a transmitter evenly spaces or paces packets at a pre-configured rate
  - helps to mitigate transient bursts
  - improves fairness
  - challenge: how to discover the bottleneck bandwidth?

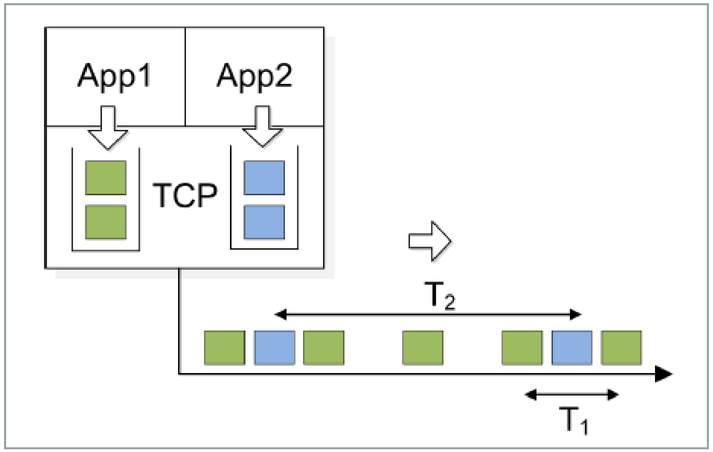

Pacing

- Consider the following test[1]
  - 100 Gbps network, 92 msec RTT
  - Four concurrent flows

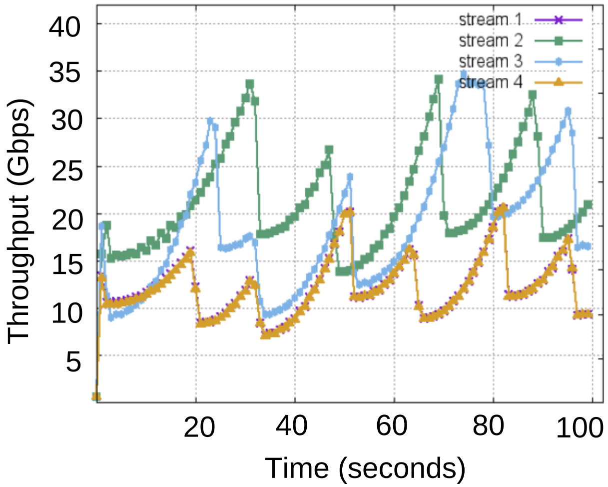

(a) Regular TCP

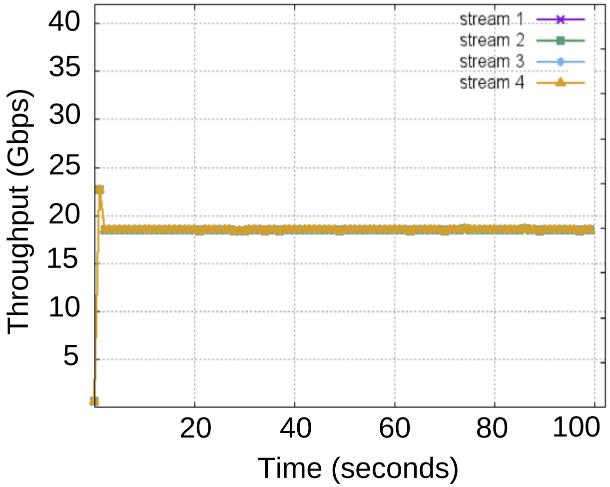

(b) Operator sets rates manually

1. https://meetings.internet2.edu/media/medialibrary/2016/10/24/20160927-tierney-improving-performance-40G-100G-data-transfer-nodes.pdf

ENABLING TCP PACING USING PROGRAMMABLE DATA PLANE SWITCHES

Elie Kfoury, Jorge Crichigno
College of Engineering and Computing
University of South Carolina

IEEE Telecommunications and Signal Processing Conference (TSP’19)
Budapest, Hungary
July 1, 2019

Overview P4 Switches

- Programming Protocol-Independent Packet Processors (P4) is a programming language for switches
- SDN is used to program the control plane
- P4 switches permit operators to program the data plane
    - Add proprietary features: invent, *develop custom protocols*

```
/**************************************************************************
********************** P A R S E R  ***********************************
*************************************************************************/

    state parse_ethernet {
        packet.extract(hdr.ethernet);
        transition select(hdr.ethernet.etherType) {
            TYPE_IPV4: parse_ipv4;
            default: accept;
        }
    }

    state parse_ipv4 {
        packet.extract(hdr.ipv4);
        verify(hdr.ipv4.ihl >= 5, error.IPHeaderTooShort);
        transition select(hdr.ipv4.ihl) {
            5             : accept;
            default       : parse_ipv4_option;
        }
    }
```

P4 code

Barefoot's Tofino (2016)

Overview P4 Switches

- Programming Protocol-Independent Packet Processors (P4) is a programming language for switches
- SDN is used to program the control plane
- P4 switches permit operators to program the data plane
  - Add proprietary features: invent, *develop custom protocols*

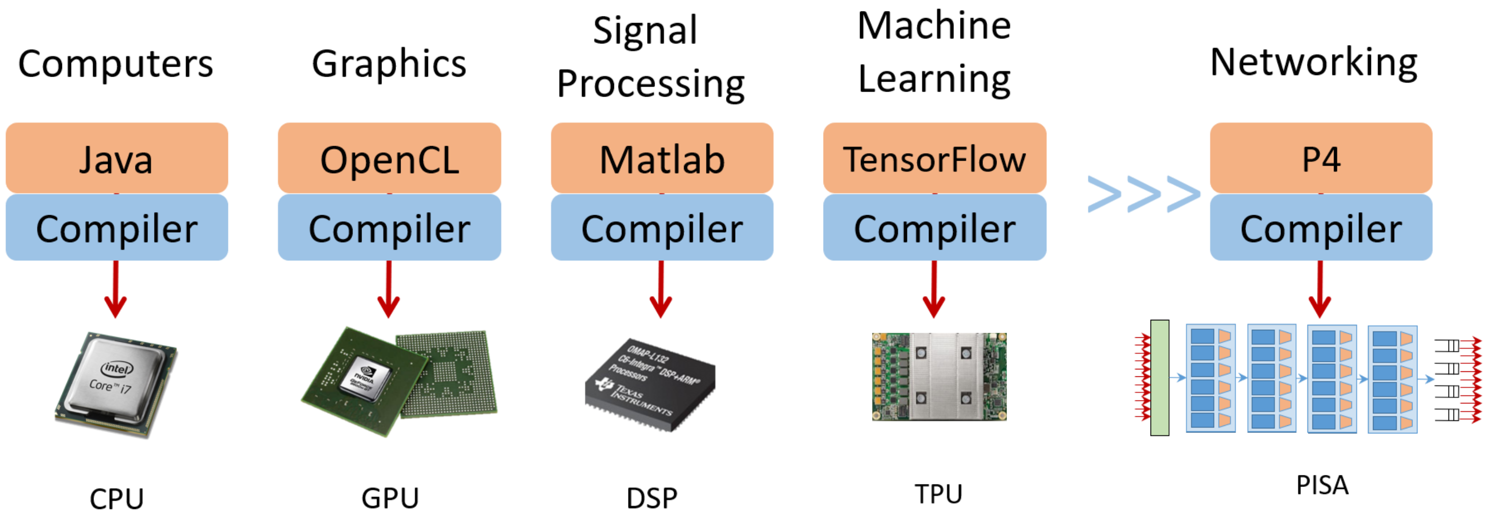

N. McKeown, "Software Defined Networking: How it has transformed networking and what happens next," Future Forum Summit, Beijing, Nov. 2018. Available online at http://yuba.stanford.edu/~nickm/talks.html.

Pacing using Programmable Switches

- What if a sender's rate is adjusted based on feedback provided by a P4 switch?
- Feedback includes number of large flows and more

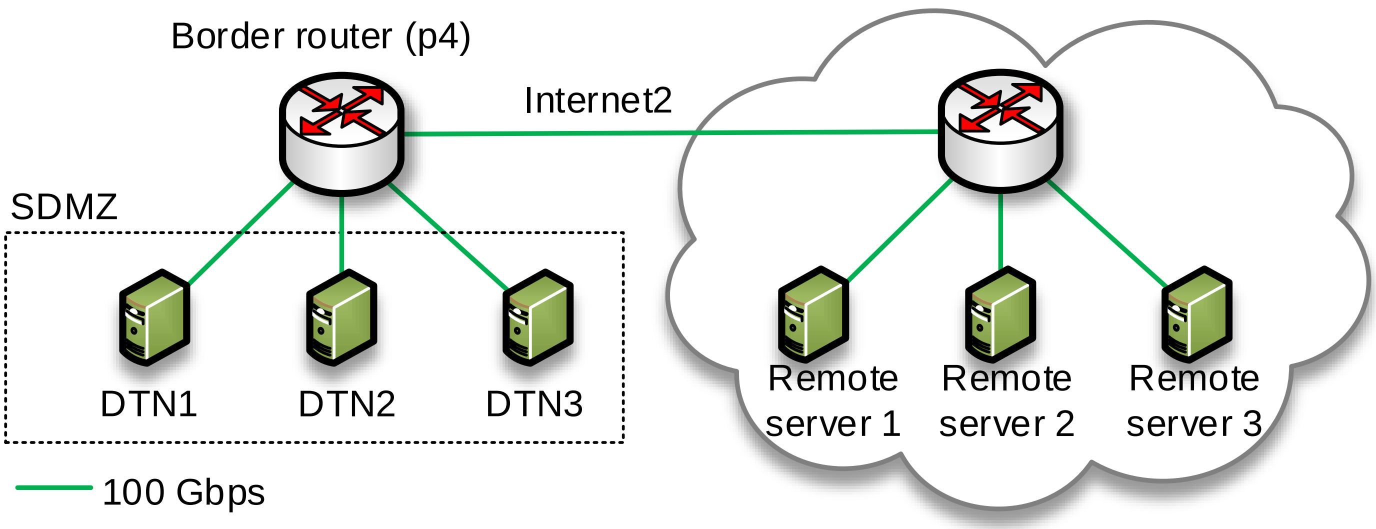

Pacing using Programmable Switches

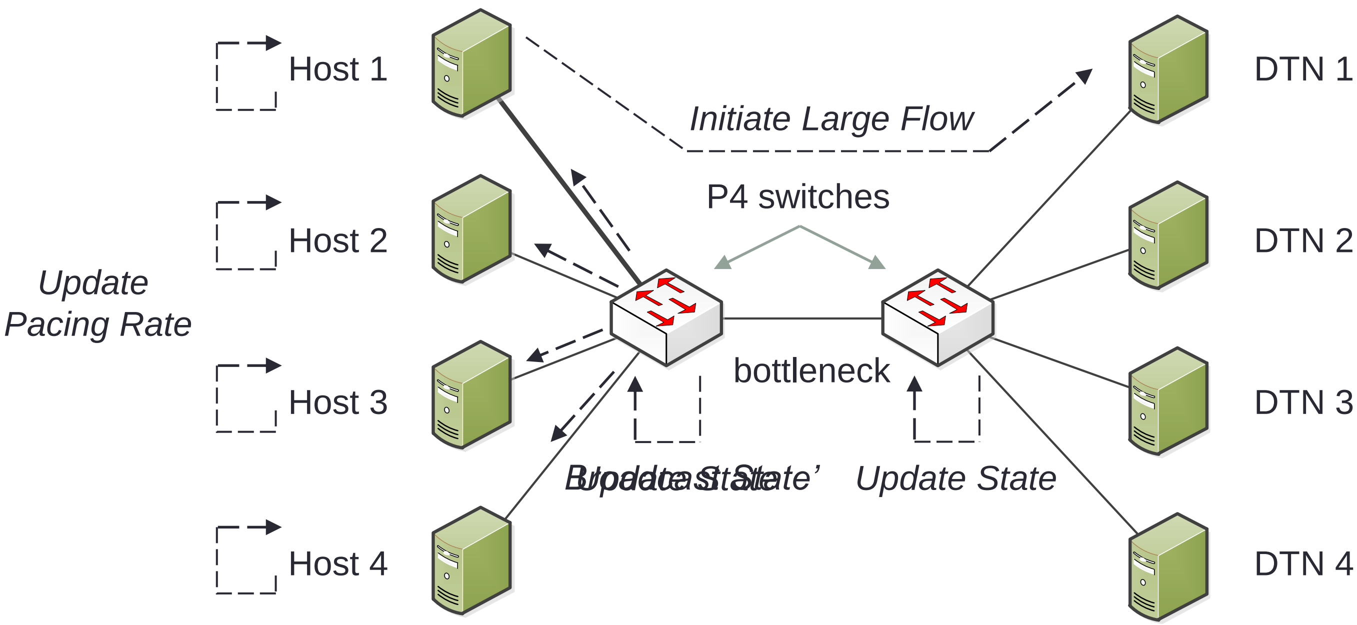

Pacing using Programmable Switches

- Switches store network's state (number of large flows)
- To initiate a large flow, a DTN inserts a custom header during the TCP 3-way handshake, using the IP options field
- Switches parse custom header, update number of large flows
- Number of large flows is returned in the SYN-ACK message, and sent to all DTNs. DTNs update their *pacing* rate

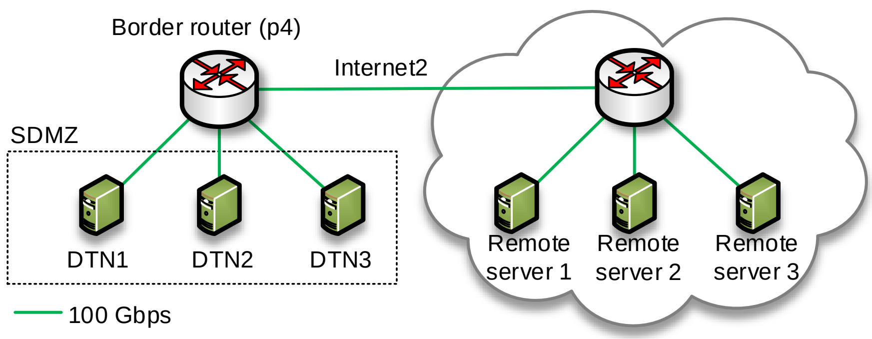

Sample topology

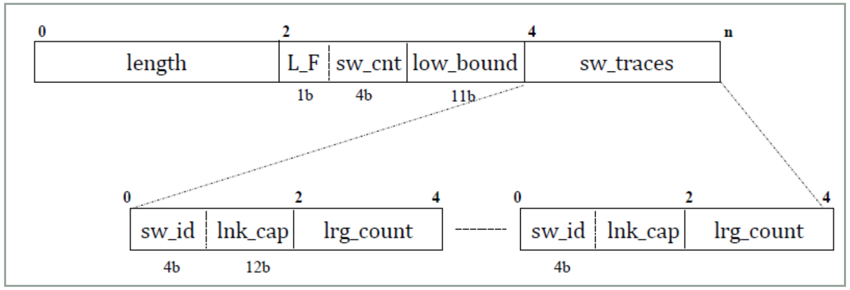

Custom protocol built using IP options field

Emulation Results

- The custom protocol was implemented in Mininet
- The P4 switch is the BMv2 from P4.org
- Four hosts (DTNs) generating flows; 100 Mbps, 20ms RTT
- Hosts adjusted their pacing rate using two pacing disciplines
  - Fair Queue (FQ)
  - Hierarchical Token Bucket (HTB)

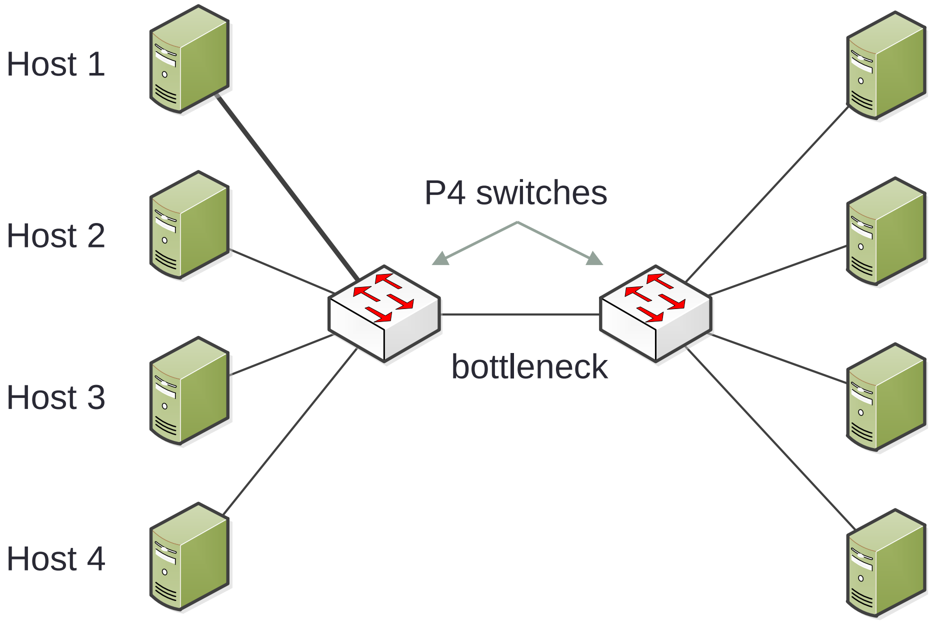

Emulation Results

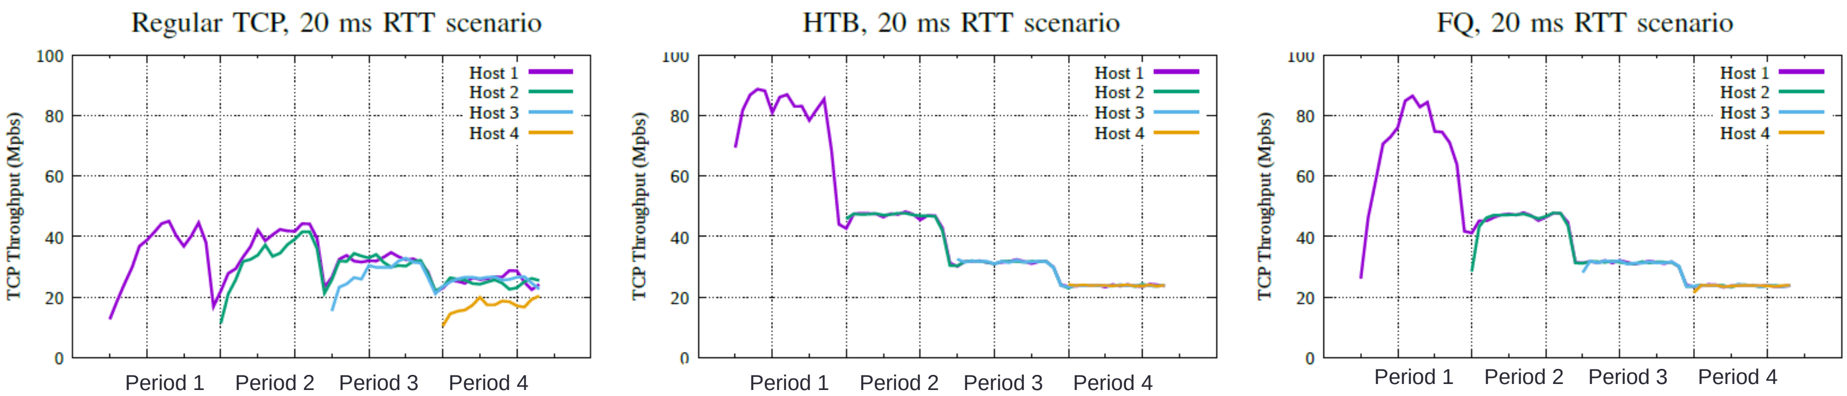

Throughput

Period	Regular TCP $\sum T_i$	T_1	T_2	T_3	T_4	HTB $\sum T_i$	T_1	T_2	T_3	T_4	FQ $\sum T_i$	T_1	T_2	T_3	T_4
P_1 (01-15 sec)	**33.62**	33.62	N/A	N/A	N/A	**81.25**	81.25	N/A	N/A	N/A	**66.59**	66.59	N/A	N/A	N/A
P_2 (16-30 sec)	**67.27**	36.06	31.21	N/A	N/A	**93.1**	46.40	46.70	N/A	N/A	**89.91**	45.85	44.06	N/A	N/A
P_3 (31-45 sec)	**88.83**	31.27	30.61	26.95	N/A	**94.42**	31.40	31.37	31.65	N/A	**93.72**	31.40	31.36	30.96	N/A
P_4 (46-60 sec)	**91.86**	25.32	24.63	25.32	16.59	**95.12**	23.78	23.75	23.73	23.86	**94.52**	23.71	23.71	23.67	23.43

Coefficient of variation and Jain's fairness

Period	Regular TCP F	CV_1	CV_2	CV_3	CV_4	HTB F	CV_1	CV_2	CV_3	CV_4	FQ F	CV_1	CV_2	CV_3	CV_4
P_1 (01-15 sec)	1.00	32.32	N/A	N/A	N/A	1.0000	8.188	N/A	N/A	N/A	1.0000	28.427	N/A	N/A	N/A
P_2 (16-30 sec)	.994	22.63	30.08	N/A	N/A	.99998	3.773	2.998	N/A	N/A	.99960	4.351	14.142	N/A	N/A
P_3 (31-45 sec)	.994	9.349	10.90	19.69	N/A	.99998	2.065	2.081	1.985	N/A	.99960	1.618	1.317	3.879	N/A
P_4 (46-60 sec)	.974	7.806	5.260	6.447	17.27	.99999	1.168	1.138	.755	.684	.99997	1.022	1.020	.996	3.336

Work in progress

- Implement proposed protocol using a real P4 switched network
- Support for more complex topologies
- Extend the sharing bandwidth scheme for scenarios where an uneven allocation is desirable (priorities)
- Use proposed protocol in the production Science DMZ at USC

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed ‘science’ network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers’ buffer size
- Resources online

A FLOW-BASED ENTROPY CHARACTERIZATION OF A NATED NETWORK AND ITS APPLICATION ON INTRUSION DETECTION

Jorge Crichigno
College of Engineering and Computing
University of South Carolina

IEEE International Conference on Communications (ICC’19)
Shanghai, China
May 22, 2019

Motivation

- Offline scalable security appliances are required in high-speed networks such as Science DMZs
- There are two approaches to characterize traffic:
  - Flow-based: information collected from header fields
  - Payload-based: information collected from payload (deep inspection)
- The amount of processing of payload-based approaches may become excessive at very high rates[1, 2]

1. R. Hofstede, P. Celeda, B. Trammell, I. Drago, R. Sadre, A. Sperotto, A. Pras, "Flow monitoring explained: from packet capture to data analysis with netFlow and ipfix," *IEEE Communications Surveys and Tutorials*, vol. 16, no. 4, 2014.
2. A. Gonzalez, J. Leigh, S. Peisert, B. Tierney, A. Lee, J. Schopf, "Monitoring big data transfers over international research network connections," in *Proceedings of the IEEE International Congress on Big Data,*, Jun. 2017.

Motivation

- Most networks use Network Address Translation (NAT)
- Although NAT has been used since early 2000s, traffic behind NAT has not been characterized
- One approach for flow characterization is to measure the *randomness* or *uncertainty* of elements of a flow
- E.g., entropy of IP addresses, ports, and combinations
- Goal: characterize normal traffic behavior (entropy) by using flow information

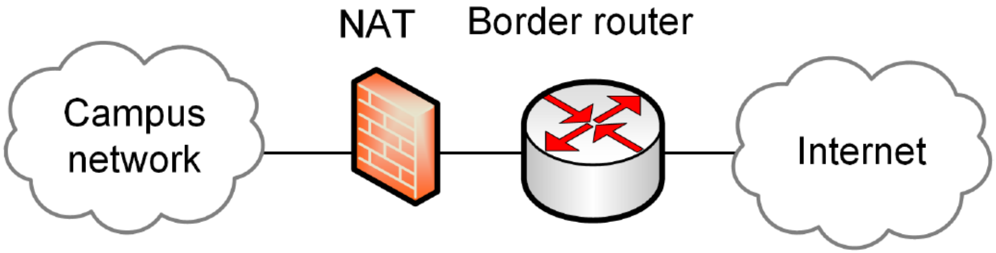

Methodology

- A flow is uniquely identified by the external IP, campus IP, external port, campus port, protocol
- Measure flow-element entropies

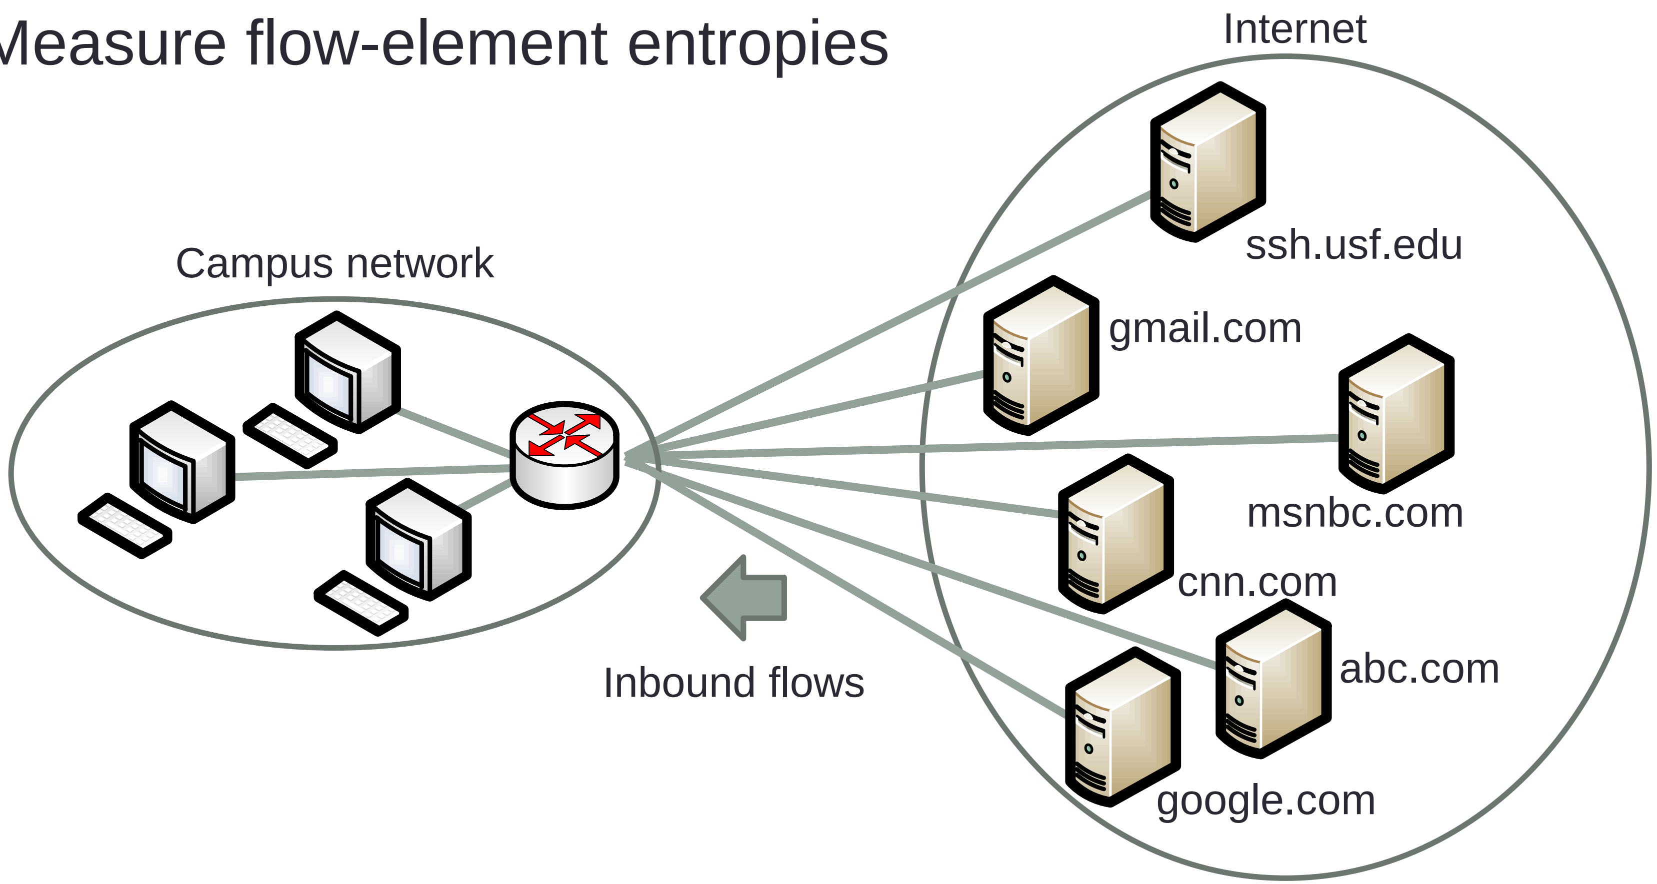

Methodology

- A flow is uniquely identified by the external IP, campus IP, external port, campus port, protocol
- Measure flow-element entropies

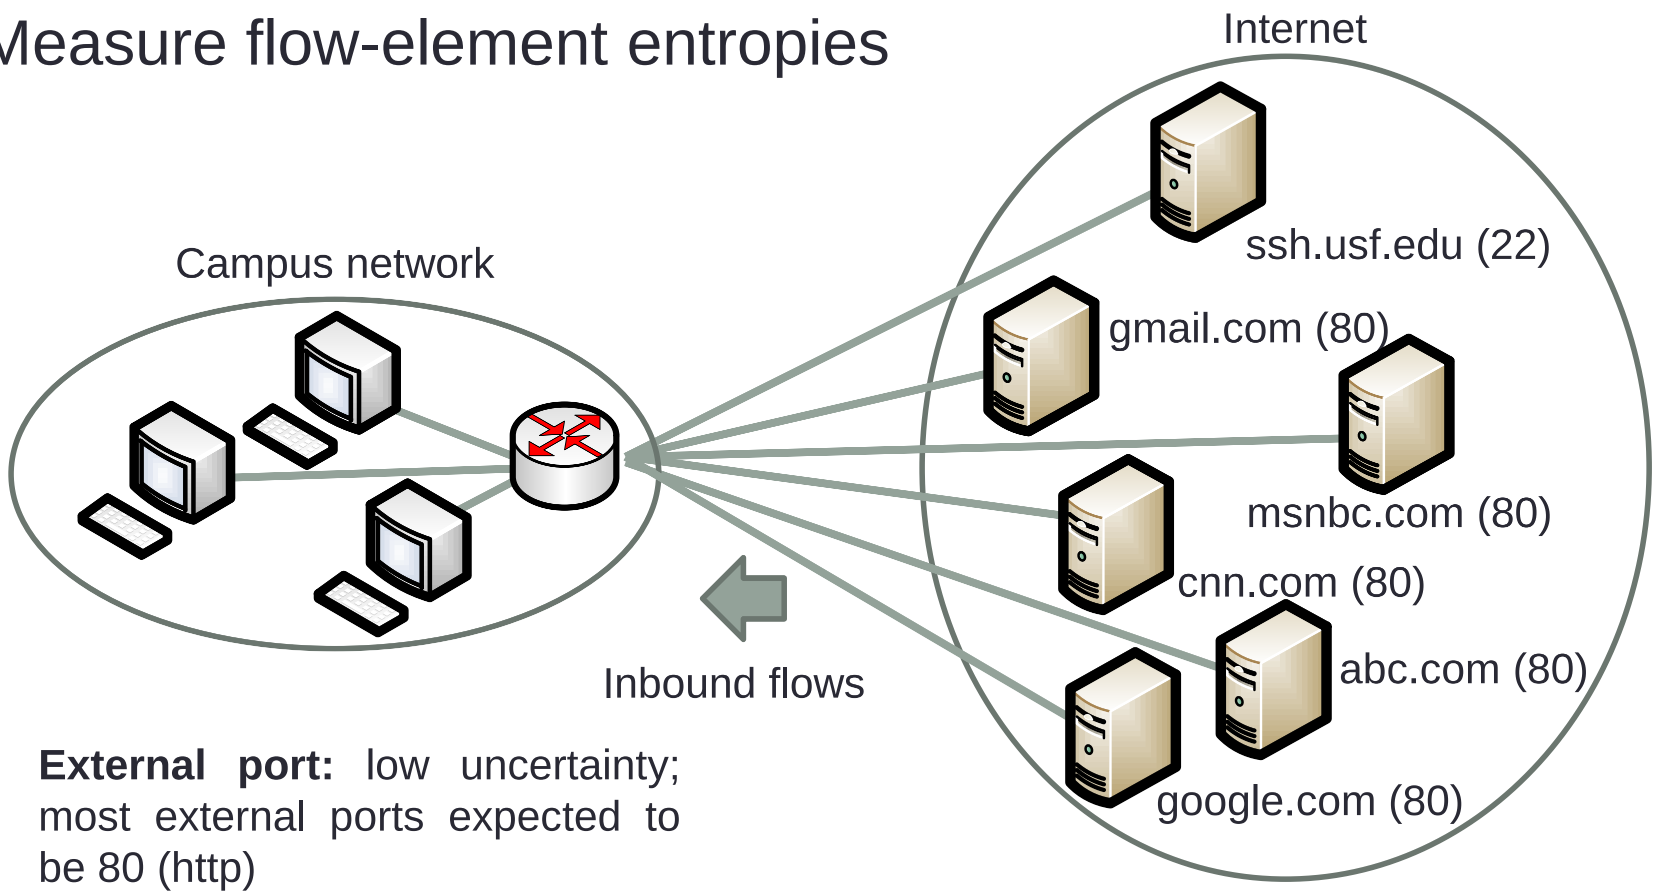

External port: low uncertainty; most external ports expected to be 80 (http)

Methodology

- Entropy provides a measure of randomness or uncertainty
- For a variable X, entropy of X = $\sum_{x \in X} p_x \log_2 \left(\frac{1}{p_x} \right)$
- For the previous port example, let *X* be the variable indicating the external port

$$X = \begin{cases} 80 \text{ with probability } p_1 = \frac{5}{6} \\ 22 \text{ with probability } p_2 = \frac{1}{6} \end{cases}$$

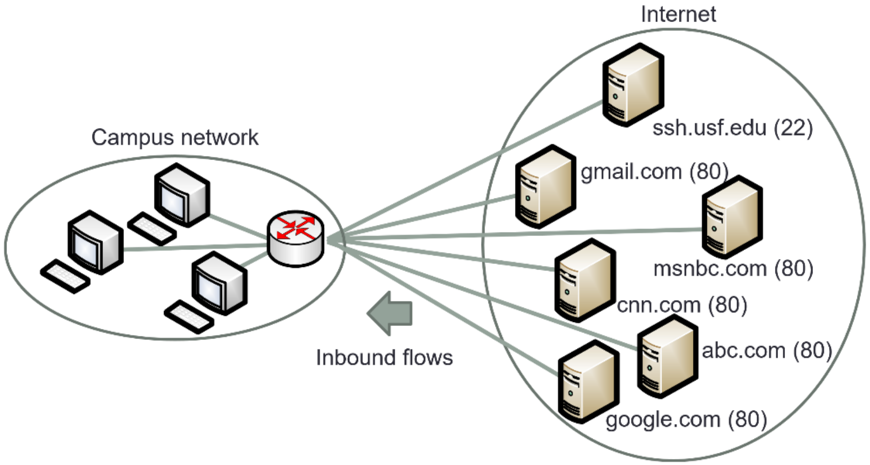

Methodology

- Entropy provides a measure of randomness or uncertainty
- For a variable X, entropy of X = $\sum_{x \in X} p_x \log_2 \left(\frac{1}{p_x}\right)$
- For the previous port example, let *X* be the variable indicating the external port

$$X = \begin{cases} 80 \text{ with probability } p_1 = \frac{5}{6} \\ 22 \text{ with probability } p_2 = \frac{1}{6} \end{cases}$$

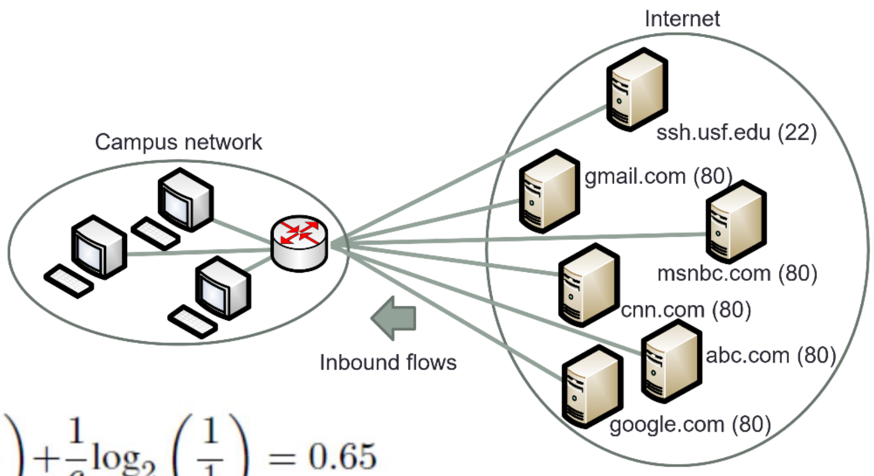

$$\text{Entropy External Port} = \sum_{i=1}^{2} p_i \log_2 \left(\frac{1}{p_i}\right) = \frac{5}{6}\log_2\left(\frac{1}{\frac{5}{6}}\right) + \frac{1}{6}\log_2\left(\frac{1}{\frac{1}{6}}\right) = 0.65$$

Methodology

- Entropy provides a measure of randomness or uncertainty
- For a variable X, entropy of X = $\sum_{x \in X} p_x \log_2 \left(\frac{1}{p_x}\right)$
- For the previous port example, let *X* be the variable indicating the external port

$$X = \begin{cases} 80 \text{ with probability } p_1 = \frac{5}{6} \\ 22 \text{ with probability } p_2 = \frac{1}{6} \end{cases}$$

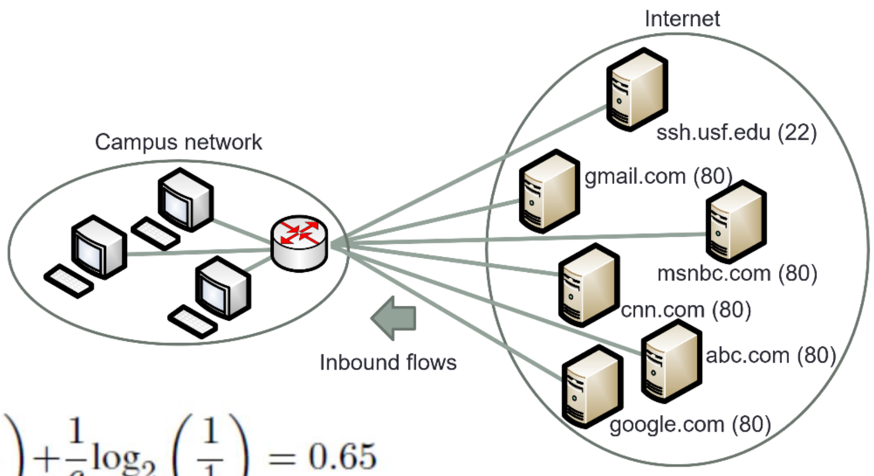

$$\text{Entropy External Port} = \sum_{i=1}^{2} p_i \log_2 \left(\frac{1}{p_i}\right) = \frac{5}{6}\log_2\left(\frac{1}{\frac{5}{6}}\right) + \frac{1}{6}\log_2\left(\frac{1}{\frac{1}{6}}\right) = 0.65$$

- 0 entropy ~ no uncertainty (e.g., all external ports are 80)
- 1 entropy ~ random -> high uncertainty

Methodology

- Campus network with 15 buildings
- Inbound traffic is used as a reference (external IP address is in the Internet, campus IP address is on campus)
- The collector organizes flow data in five-minute time slots
- Traffic data observed during a week is representative of the campus traffic

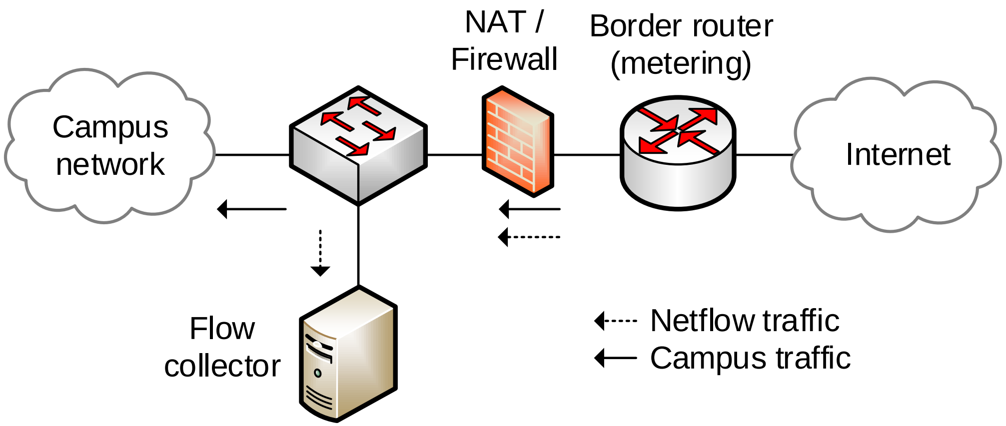

Methodology

- The entropy of a random variable X is:

$$H(X) = \sum_{i=1}^{N} p(x_i)\log_2\left(\frac{1}{p(x_i)}\right),$$

where $x_1, x_2, \ldots x_N$ is the range of values for X, and $p(x_i)$ is the probability that X takes the value x_i

- For each external (campus) IP address (port) x_i, the probability $p(x_i)$ is calculated as

$$p(x_i) = \frac{\text{Flows with } x_i \text{ as external (campus) IP addr. (port)}}{\text{Total number of flows}}$$

- Entropies are normalized

Methodology

- This paper also considers the entropy of the 3-tuple {external IP, campus IP, campus port}
- For a given 3-tuple x_i, the corresponding probability is calculated as

$$p(x_i) = \frac{\text{Flows with } x_i \text{ as 3-tuple}}{\text{Total number of flows}}$$

Results

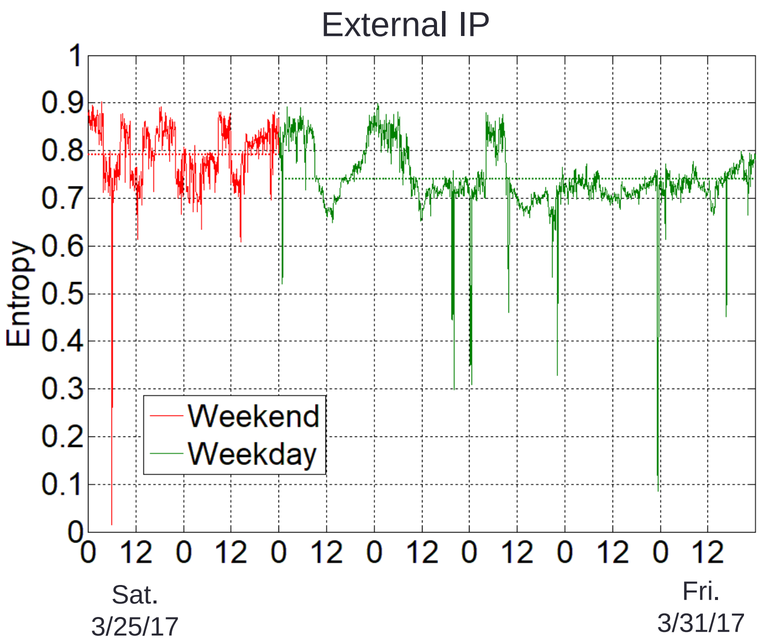

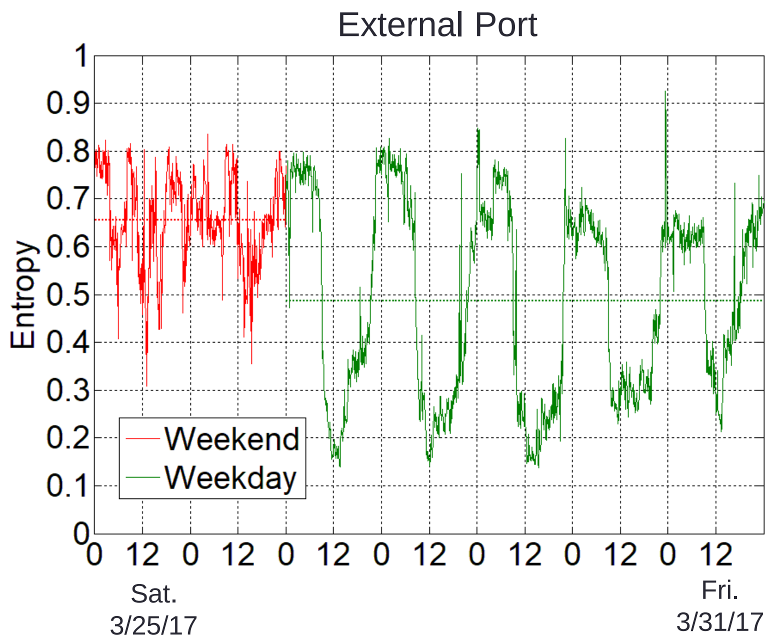

External IP

- In general, high entropy, 'many' external IP addresses
- External IPs dispersed in the Internet
- Abnormal low entropy points
- Entropy near zero (no uncertainty of the external IP address), or 'very low' level (few external IP addresses dominate the distribution)

External port

- Higher entropy during the night, weekends
- Low entropy during the day, noon
- Large volume of http flows when students are on campus (less uncertainty/entropy on external port)
- Abnormal high entropy points
- Entropy widely varies over 'hours' but not over very short time periods

Results

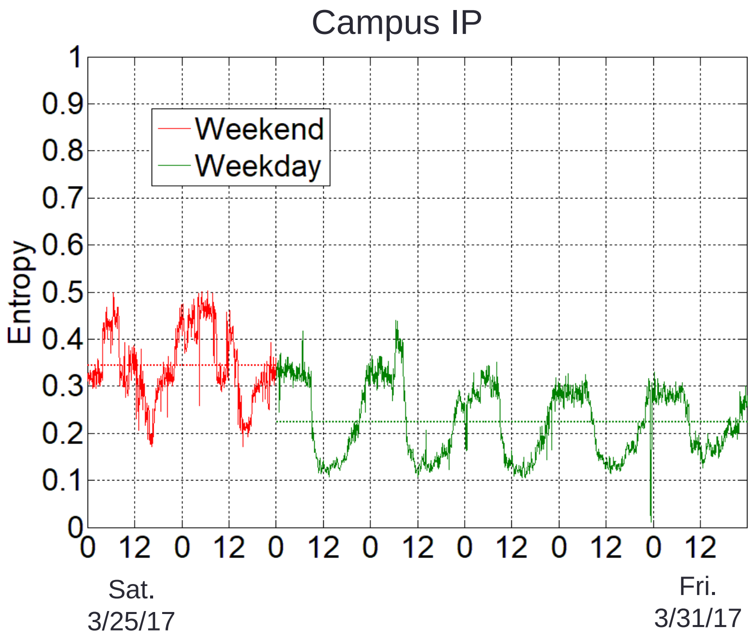

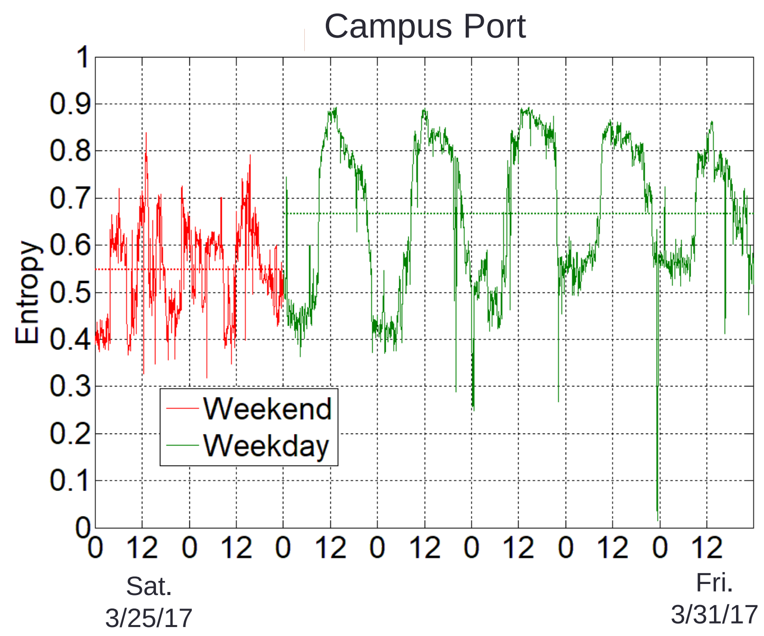

Campus IP

- In general, low entropy, 'few' IP addresses on campus
- Higher entropy on weekends and at night
- Lower entropy when students are on campus
- A handful of public IP addresses used for regular Internet connectivity (NAT operation)
- Entropy varies over 'hours' but not over very short time periods

Campus port

- Lower entropy at night
- High entropy (close to uniform distribution) at noon
- Dynamic ports used by browsers when students connect to the Internet
- Abnormal low entropy points
- Entropy widely varies over 'hours' but not over very short time periods

Results

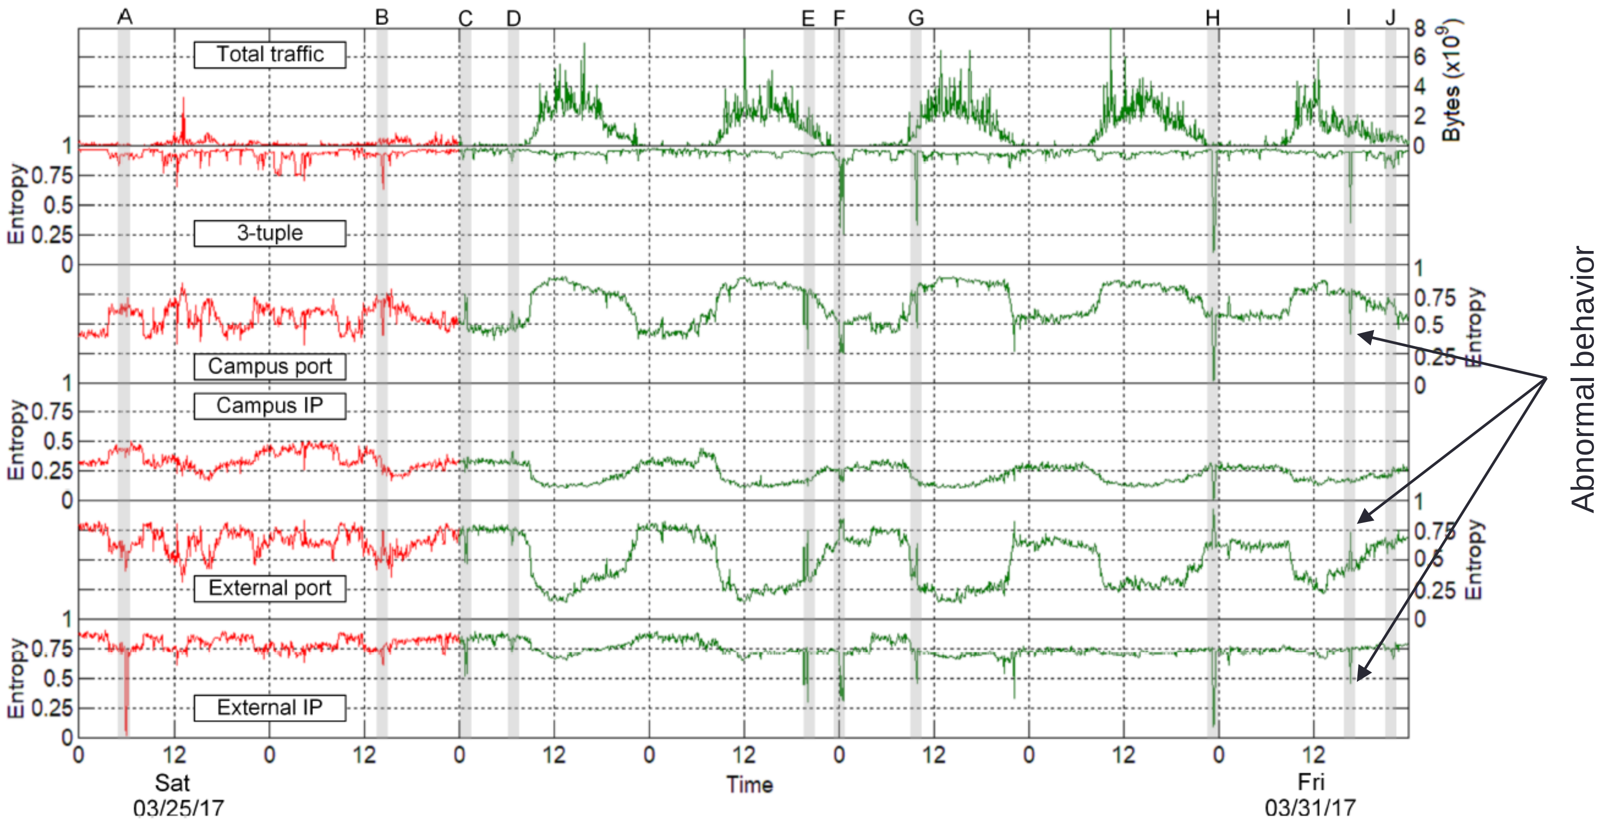

- Anomalies are detected by a single feature or by correlating multiple features
- E.g., event I: low campus port's entropy, high external port's entropy, low external IP's entropy

Results

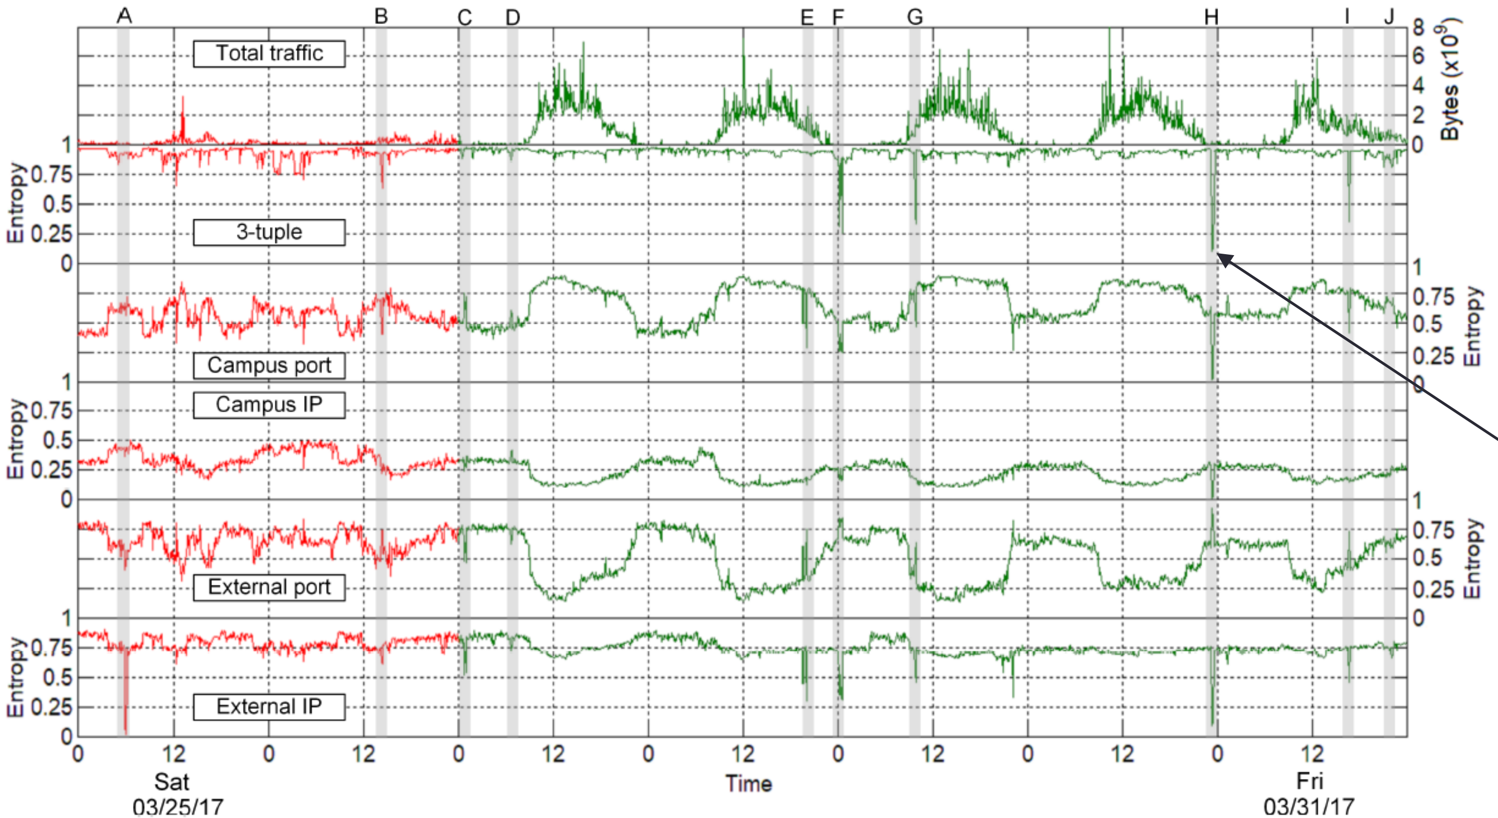

Distributed brute force attack

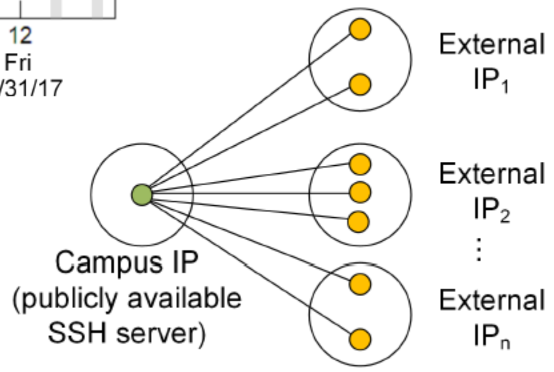

Results

- Correlation of entropy time-series

	Campus IP	Campus port	External IP	External port	Total traffic
Weekday					
3-tuple	0.23	0.1	0.6	-0.02	-0.05
Campus IP		-0.85	0.6	0.89	-0.8
Campus port			-0.37	-0.98	0.78
External IP				0.45	-0.36
External port					-0.81
Weekend					
3-tuple	-0.23	-0.12	0.56	0.06	-0.03
Campus IP		0.15	-0.38	0.06	-0.38
Campus port			-0.48	-0.93	0.31
External IP				0.48	-0.05
External port					-0.39

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed ‘science’ network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers’ buffer size
- Resources online

ROUTERS' BUFFER SIZE

Bufferbloat

- Routers and switches must have enough memory allocated to hold packets momentarily (buffering)
- Rule of thumb:
  - Buffer size = RTT · bottleneck bandwidth[1, 2]

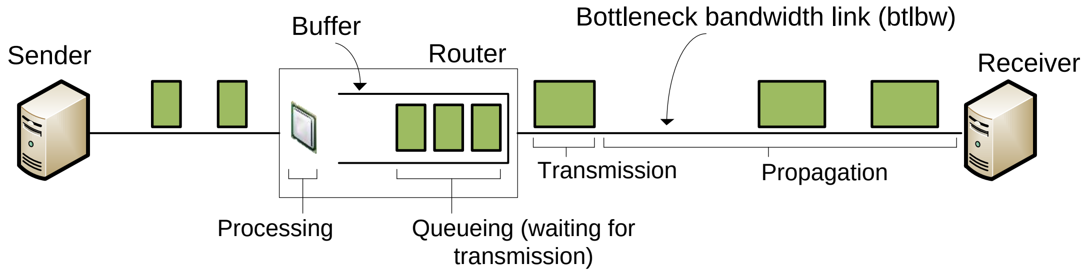

1. C. Villamizar, C. Song, "High performance TCP in ansnet," ACM Computer Communications Review, vol. 24, no. 5, pp. 45-60, Oct. 1994.
2. R. Bush, D. Meyer, "Some internet architectural guidelines and philosophy," Internet Request for Comments, RFC Editor, RFC 3439, Dec. 2003. [Online]. Available: https://www.ietf.org/rfc/rfc3439.txt.

Bufferbloat

- Bufferbloat is a condition that occurs when the router buffers too much data, leading to excessive delays

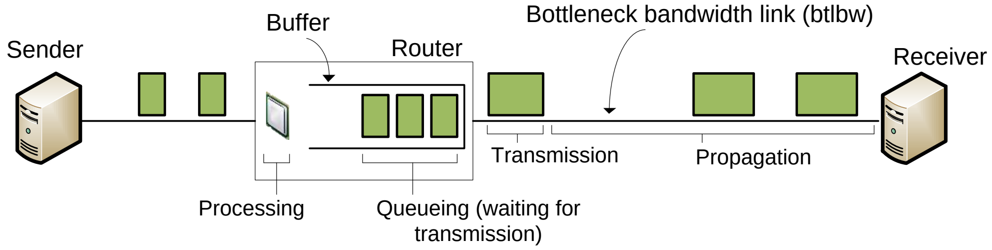

Bufferbloat

- Bufferbloat is a condition that occurs when the router buffers too much data, leading to excessive delays

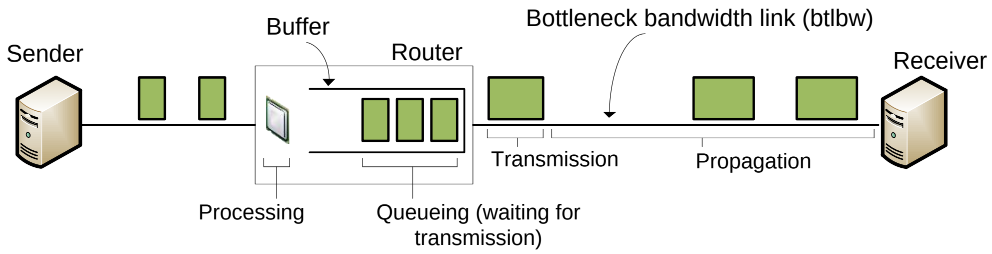

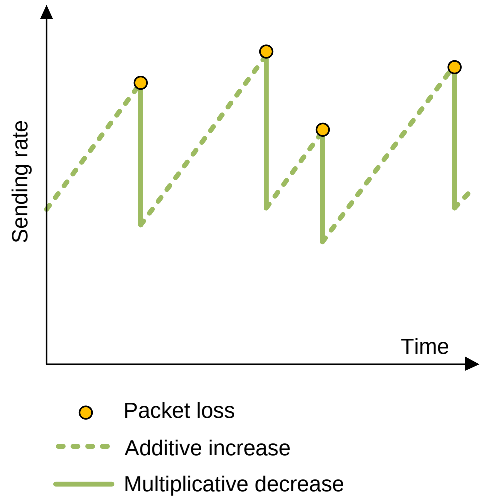

Bufferbloat

- Bufferbloat is a condition that occurs when the router buffers too much data, leading to excessive delays

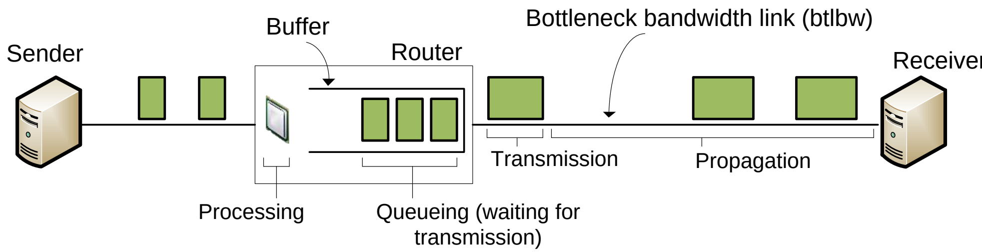

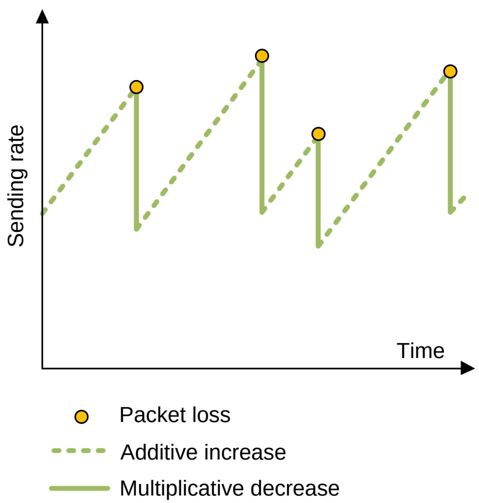

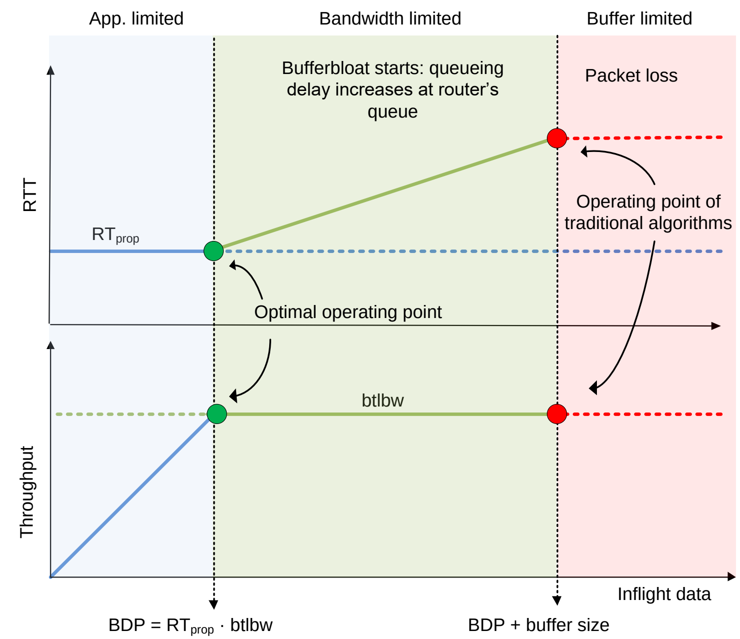

1. N. Cardwell, Y. Cheng, C. Gunn, S. Yeganeh, V. Jacobson, “BBR: congestion-based congestion control,” *Communications of the ACM*, vol 60, no. 2, pp. 58-66, Feb. 2017.

Bufferbloat

- Topology Lab 14
- 1 Gbps, 20ms link s1-h3
  - Measure RTT and throughput h1 > h3
  - Modify buffer size at s1 (interface s1-eth2)
    - ✓ Case 1: buffer size = $(1 \cdot 10^{9}) \cdot (20 \cdot 10^{-3})$ [bits] = 2,500,000 [bytes]
    - ✓ Case 2: buffer size = 25,000,000 [bytes]

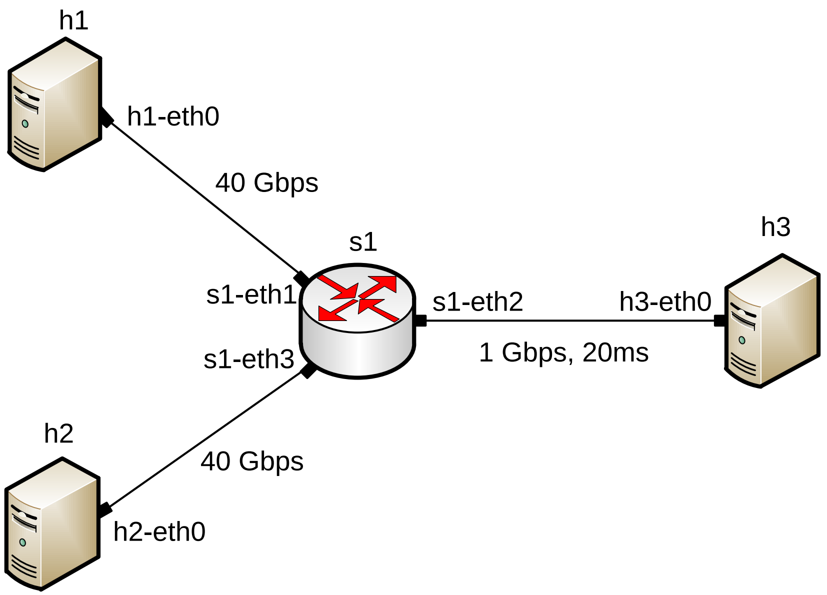

Bufferbloat

Buffer size = 1 BDP

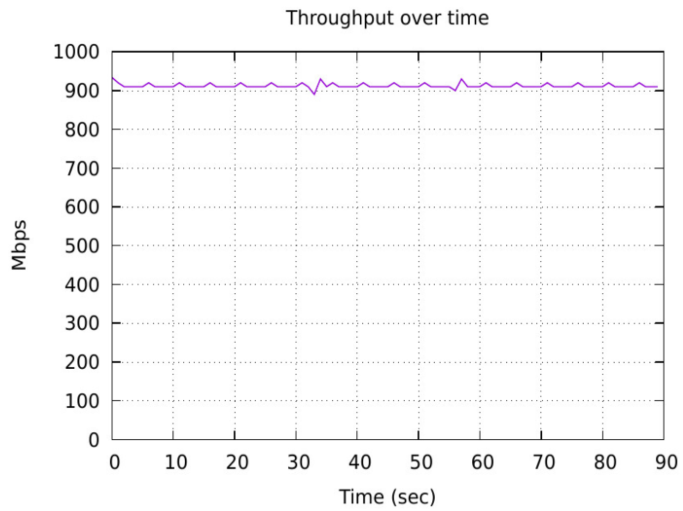

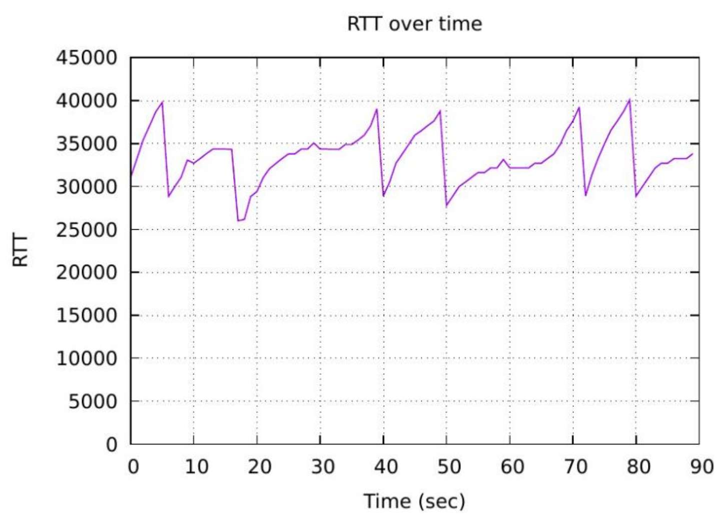

Buffer size = 10 BDP

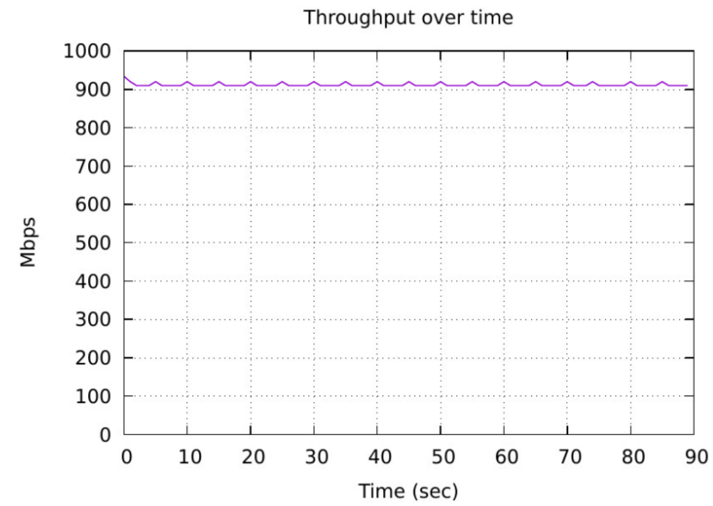

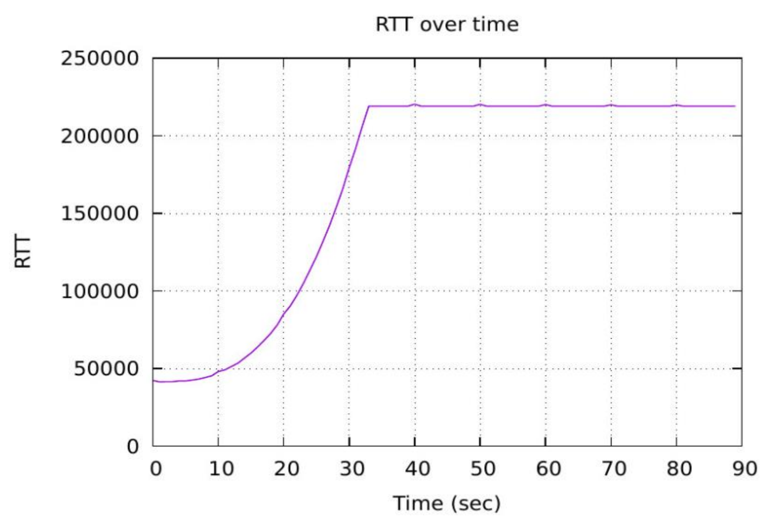

Agenda

- Introduction to University of South Carolina (USC)
- The Science DMZ
  - Motivation for a high-speed ‘science’ network architecture
  - Science DMZ architecture
  - Research opportunities: pacing, entropy-based intrusion detection, routers’ buffer size
- Resources online

Resources Online

- CI Lab website
  - http://ce.sc.edu/cyberinfra/
- A tutorial on Tools and Protocols for High-Speed Networks
  - http://ce.sc.edu/cyberinfra/workshop.html
- University of South Carolina
  - https://sc.edu/

Additional Slide

- Protocol Independent Switch Architecture