"HANDS-ON VLABS: TCP BEST PRACTICES; CONGESTION CONTROL, BUFFERS, PARALLEL STREAMS, MSS, PACING"

J. Crichigno, E. Kfoury
Department of Integrated Information Technology
University of South Carolina

NSF Award 1829698
"CyberTraining CIP: Cyberinfrastructure Expertise on High-throughput Networks for Big Science Data Transfers"

LAB SERIES: NETWORK TOOLS AND PROTOCOLS

LAB 14: ROUTER'S BUFFERBLOAT

Bufferbloat

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

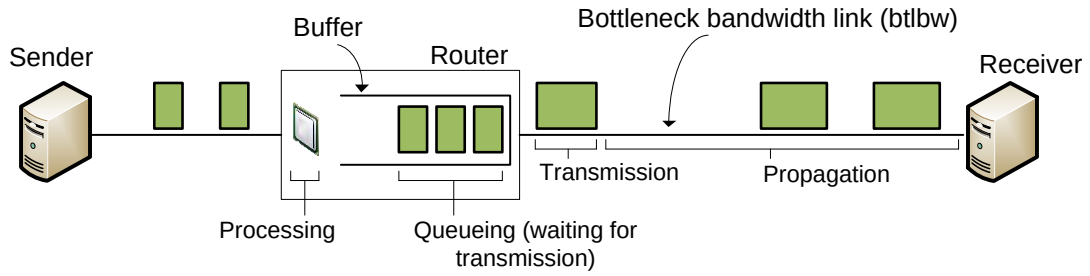

1. C. Villamizar, C. Song, “High performance TCP in ansnet,” ACM Computer Communications Review, vol. 24, no. 5, pp. 45-60, Oct. 1994.
2. R. Bush, D. Meyer, “Some internet architectural guidelines and philosophy,” Internet Request for Comments, RFC Editor, RFC 3439, Dec. 2003. [Online]. Available: https://www.ietf.org/rfc/rfc3439.txt.

Bufferbloat

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

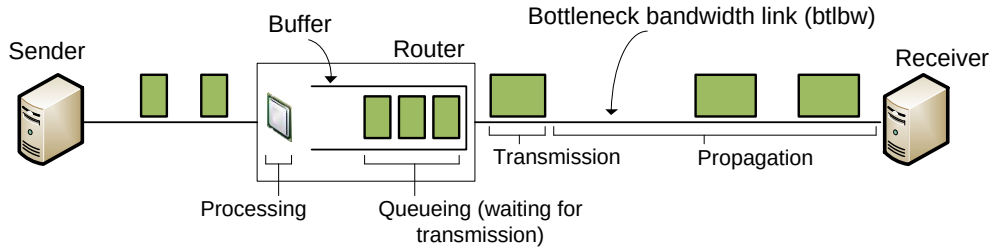

Bufferbloat

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

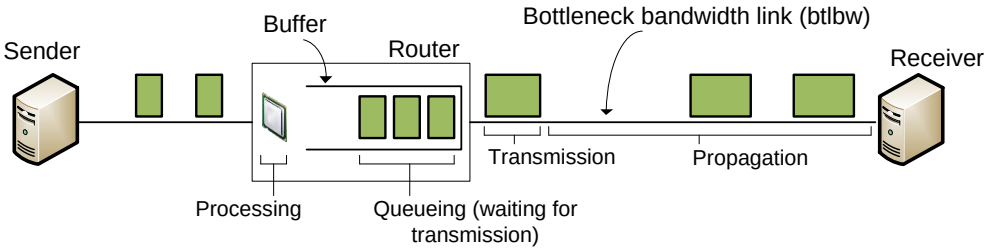

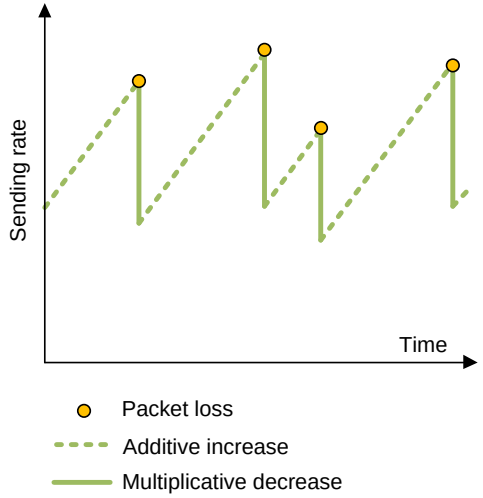

Bufferbloat

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

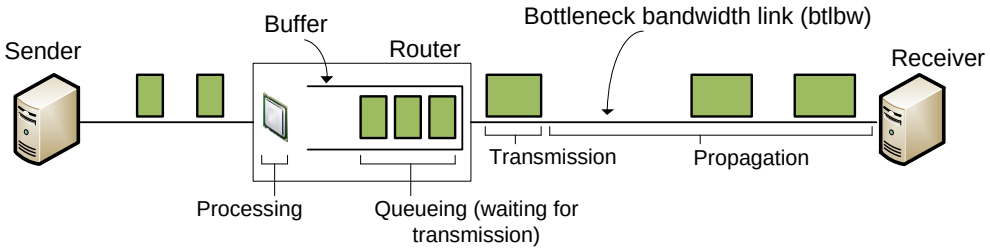

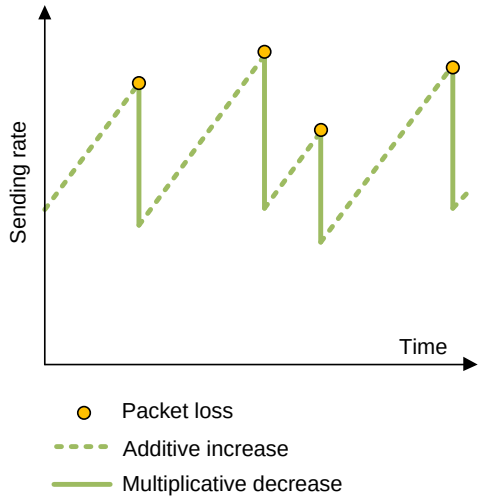

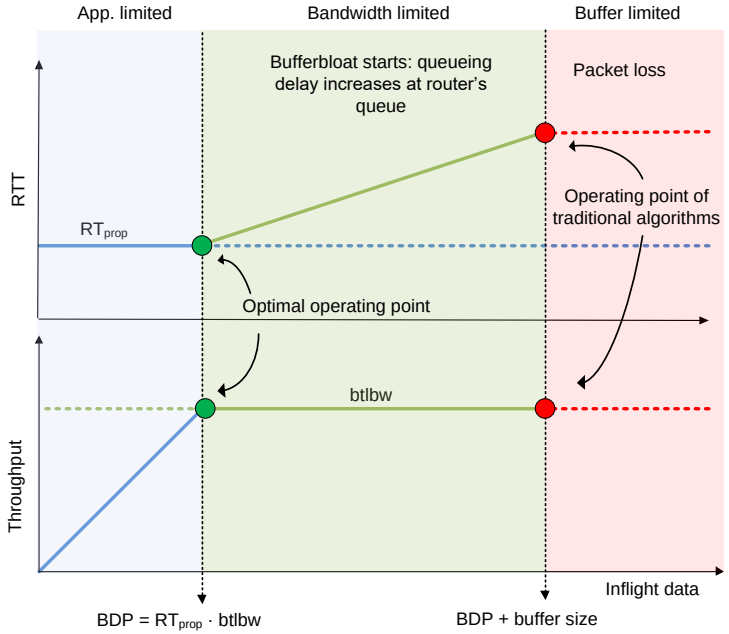

1. N. Cardwell, Y. Cheng, C. Gunn, S. Yeganeh, V. Jacobson, "BBR: congestion-based congestion control," *Communications of the ACM*, vol 60, no. 2, pp. 58-66, Feb. 2017.

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]

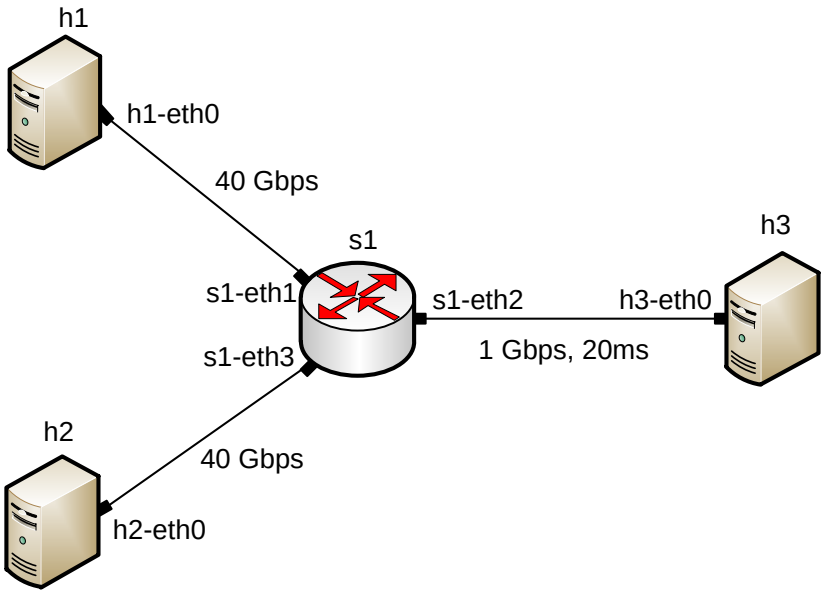

Bufferbloat

Buffer size = 1 BDP

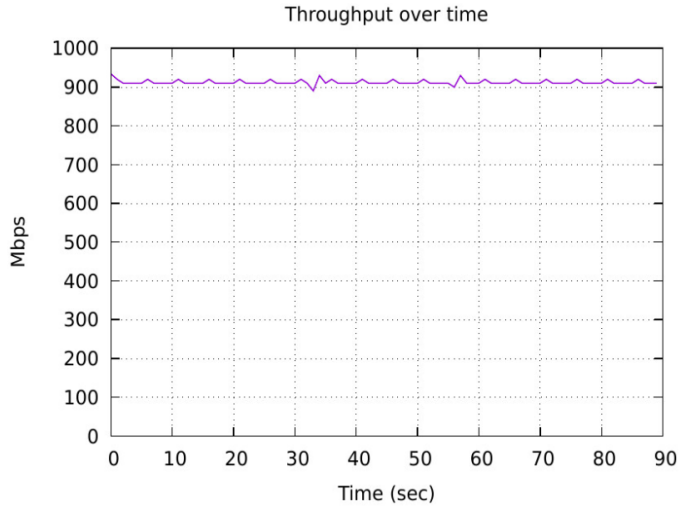

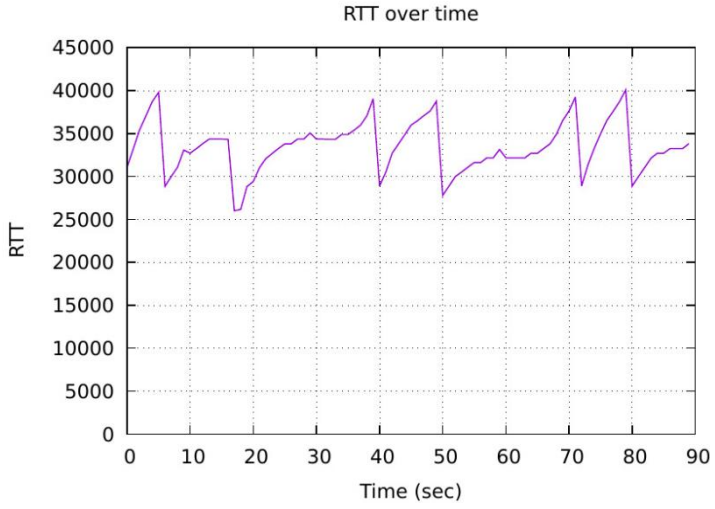

Buffer size = 10 BDP

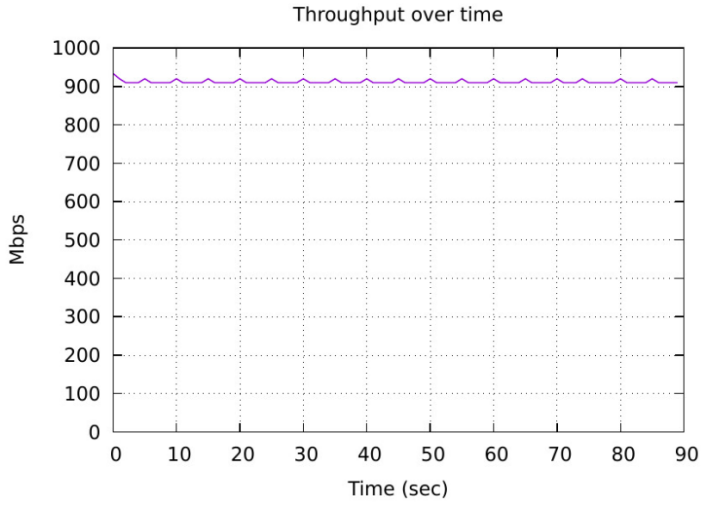

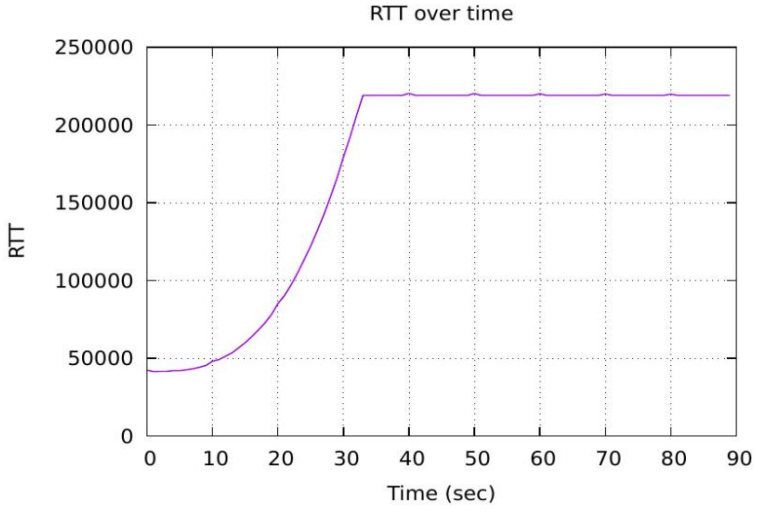

LAB 12: TCP RATE CONTROL WITH PACING

TCP Pacing

- Traditional congestion control algorithms
  - do not provide a time period over which the data should be transmitted
  - do not specify how the data should be spread
- If packets were transmitted at a steady pace, queues formation may be reduced, avoiding packet losses

TCP Pacing

- With TCP pacing, a transmitter evenly spaces or paces packets at a pre-configured rate
  - helps to mitigate transient bursts
  - improves fairness

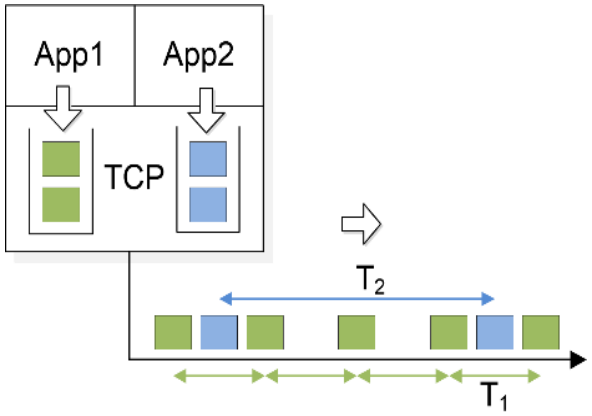

TCP Pacing

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

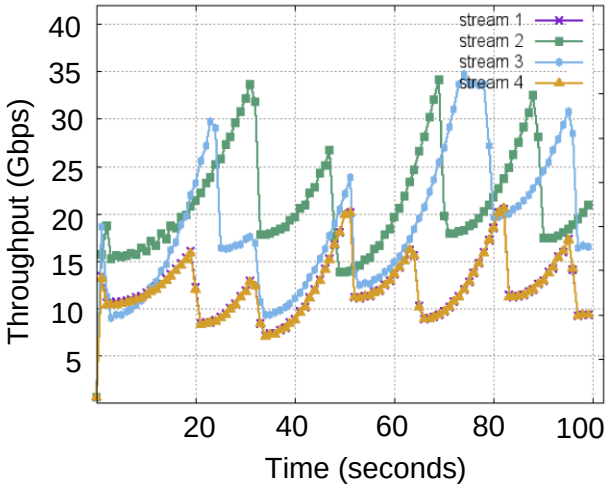

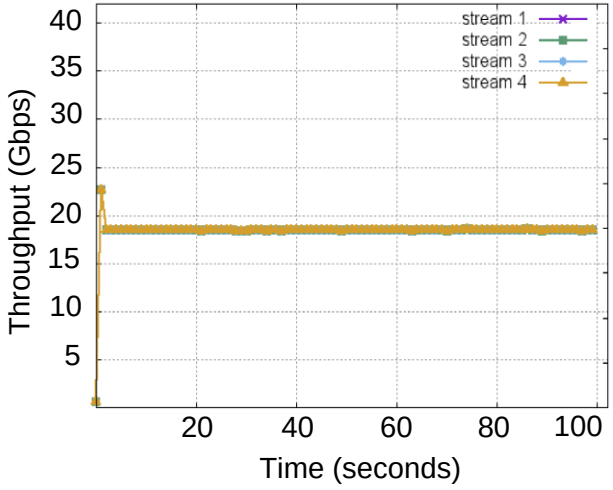

1. https://meetings.internet2.edu/media/medialibrary/2016/10/24/20160927-tierney-improving-performance-40G-100G-data-transfer-nodes.pdf

TCP Pacing

- Topology Lab 12
- 10 Gbps, 20ms link s1-s2
- Simultaneous flows
  - h1 > h2
  - h3 > h4
  - h5 > h6
  - h7 > h8

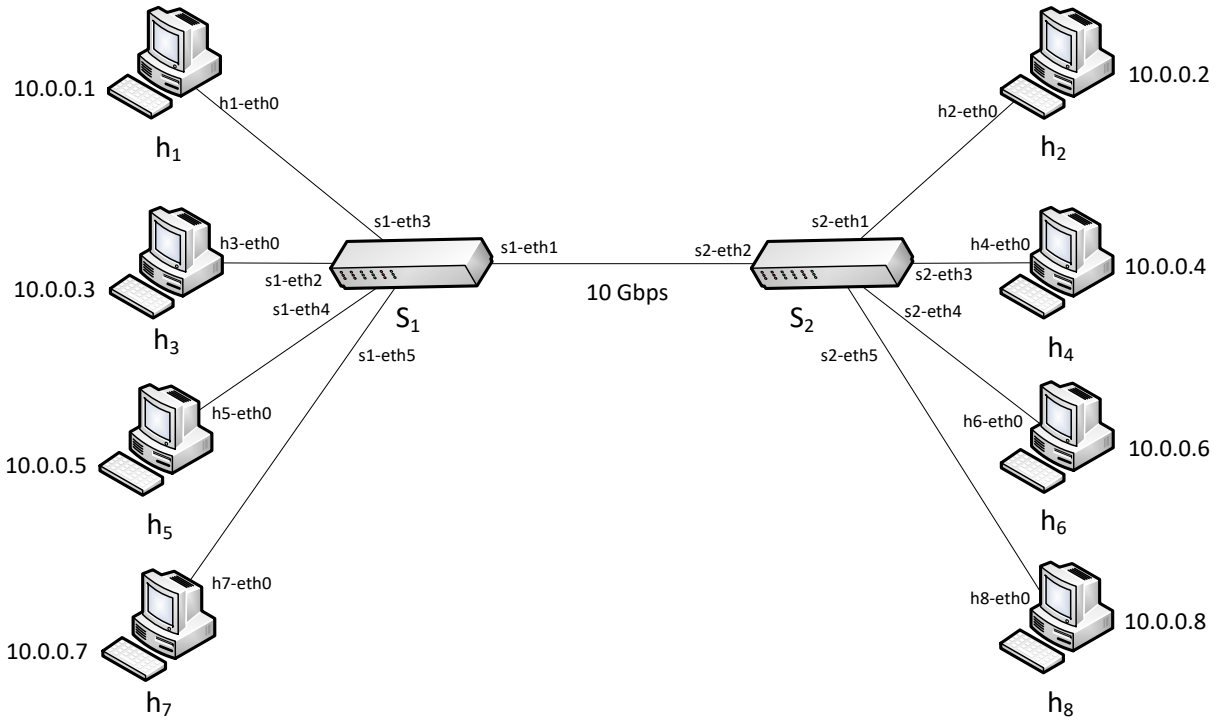

TCP Pacing

- Topology Lab 12
- 10 Gbps, 20ms link s1-s2
- Simultaneous flows
  - h1 > h2
  - h3 > h4
  - h5 > h6
  - h7 > h8

Throughput w/o pacing

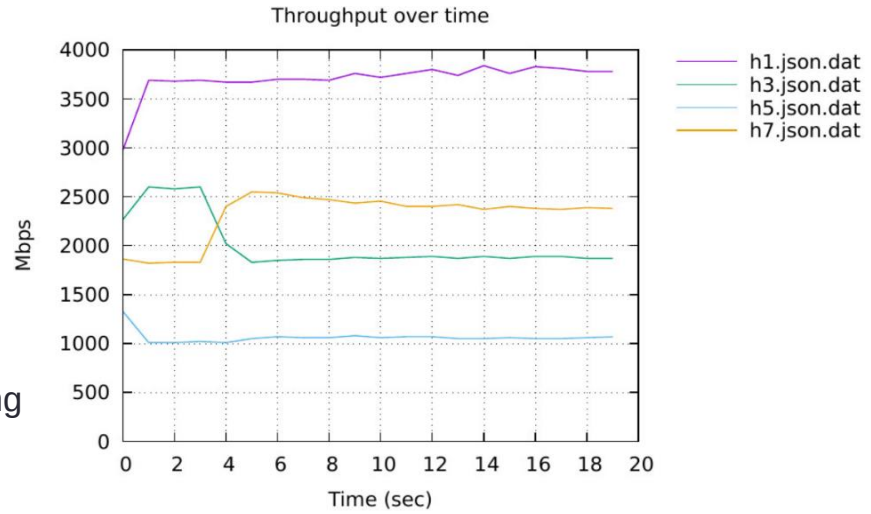

Throughput w/ pacing

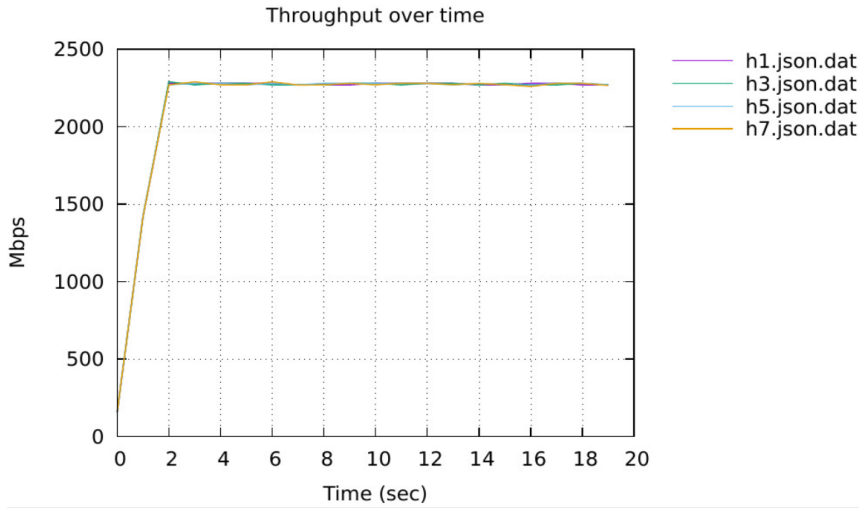

Network Information

“uscguest” wireless network	You must create an account to access Internet
http://ce.sc.edu/cyberinfra/workshop.html	Workshop website
https://netlab.cec.sc.edu/	URL of the virtual lab platform Username: lastname (lowercase letters) Password: nsf2019
https://10.173.78.50	IP address of the virtual lab platform (for users using the students’ wireless network)