ANT: Analysis of Network Traffic

Christos Papadopoulos
John Heidemann
Antonio Ortega
Urbashi Mitra

Alefiya Hussain, Xinming He
Usman Riaz, Connie Kung
Gen Bartlett, Rishi Sinha

http://www.isi.edu/ant

Outline

- **Modeling packet-capture systems**
- Detecting saturated links by looking at the aggregate
- Measuring bandwidth fluctuations

Modeling Packet Capture Systems

- Question: how do we represent a packet capture system in traditional signal processing terms?
  - what sampling frequencies are required?
  - what kinds of error should we expect?
- Components of interest:
  - standard Ethernet card vs. DAG packet capture card
  - effects of interrupt rate, basic network speed, tcpdump/PC clock resolution, etc.

Idealized Network

- 100Mb/s => 1 bit every 10ns => 100MHz

but we care about *packets*

- 40B minimum packet size
  320k pkts/s => pkt every 3.2us => 320kHz
- or 1500B max pkt size
  => pkt every 120us => 8400Hz

but network card, OS, measurement s/w effects?

Measurement System Model

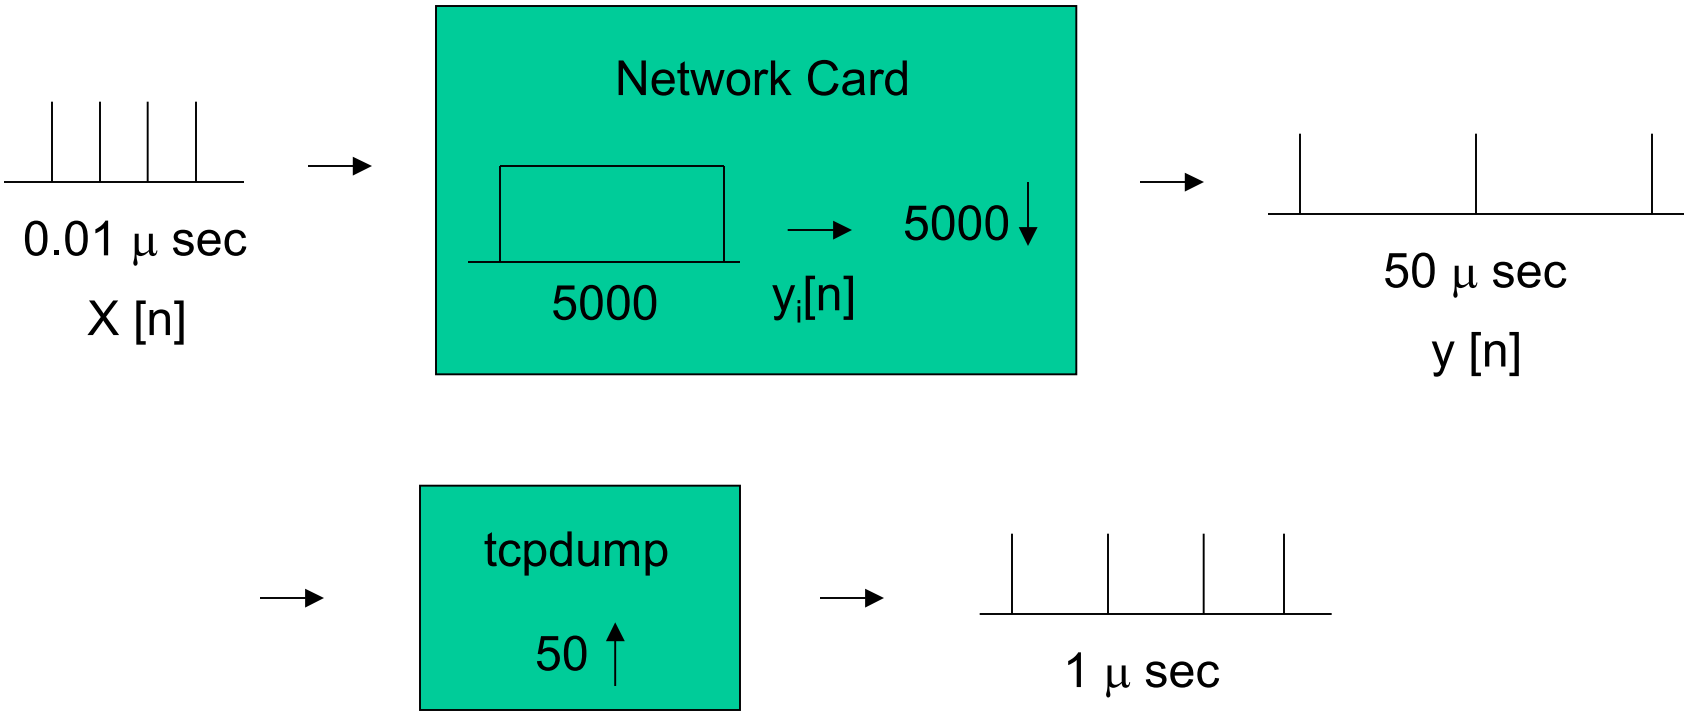

networking card counts number of arrivals in 50us interval => downsampling
tcpdump timestamps at lower resolution, it time stamps when it finishes interaction with card => upsampling

Outline

- Modeling packet-capture systems
- **Detecting saturated links by looking at the aggregate**
- Measuring bandwidth fluctuations

Spectral Characteristics of Saturated Links

- A saturated (bottleneck) link will clock packets out at a regular interval, depending on:
  - Link speed
  - Packet size
- Question: can we detect the presence of a saturated link by examining the aggregate?
- Work done by Xinming He

Spectrum of a Saturated 100 Mbps Ethernet

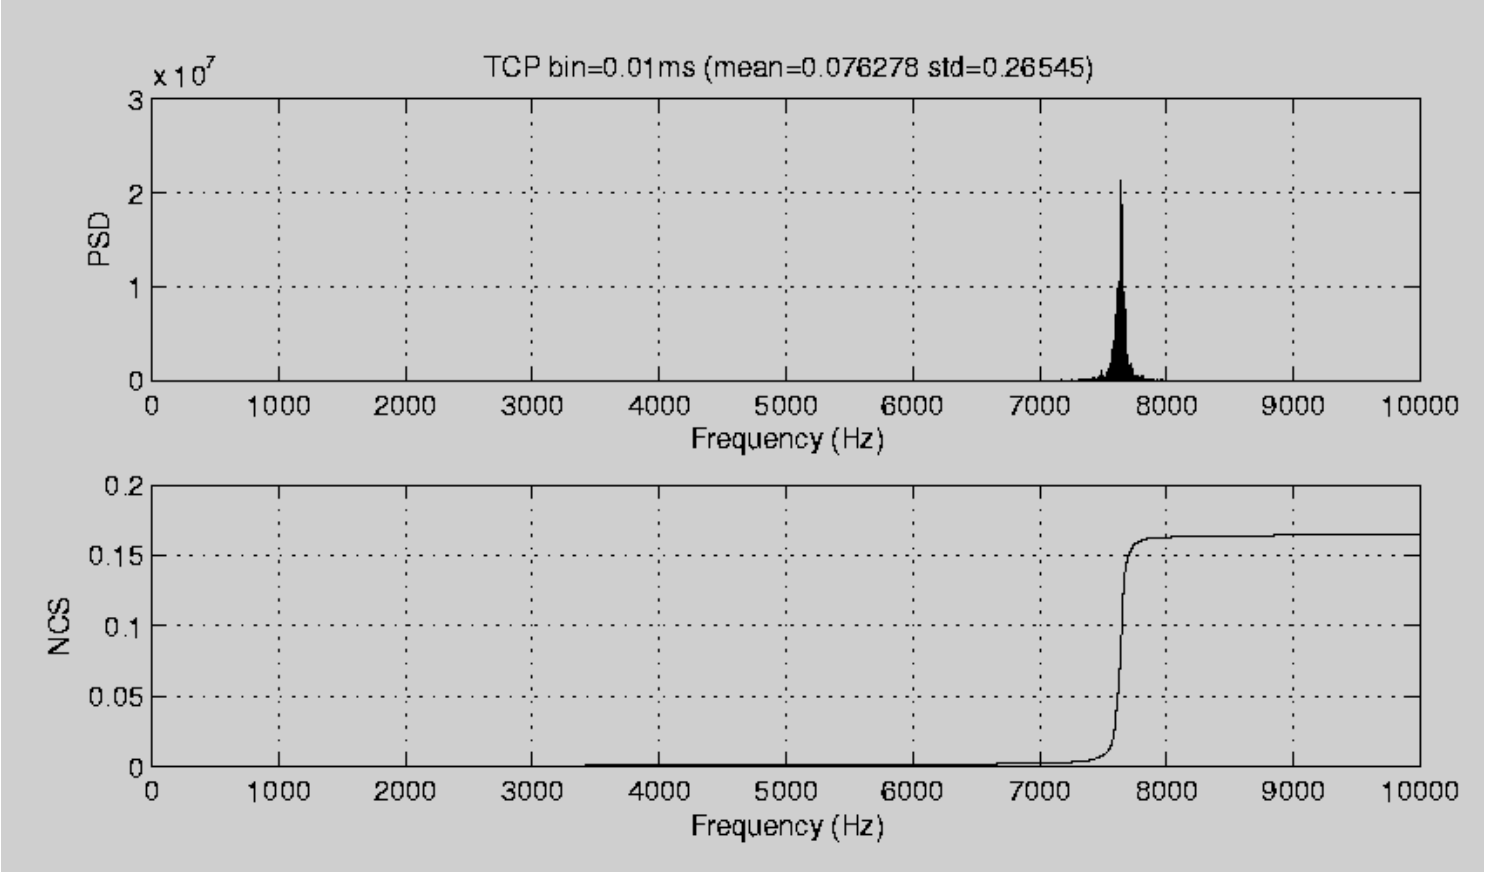

Signal at ~8KHz corresponding to back-to-back max-size packets

Spectrum of a Saturated 10Mbps Ethernet

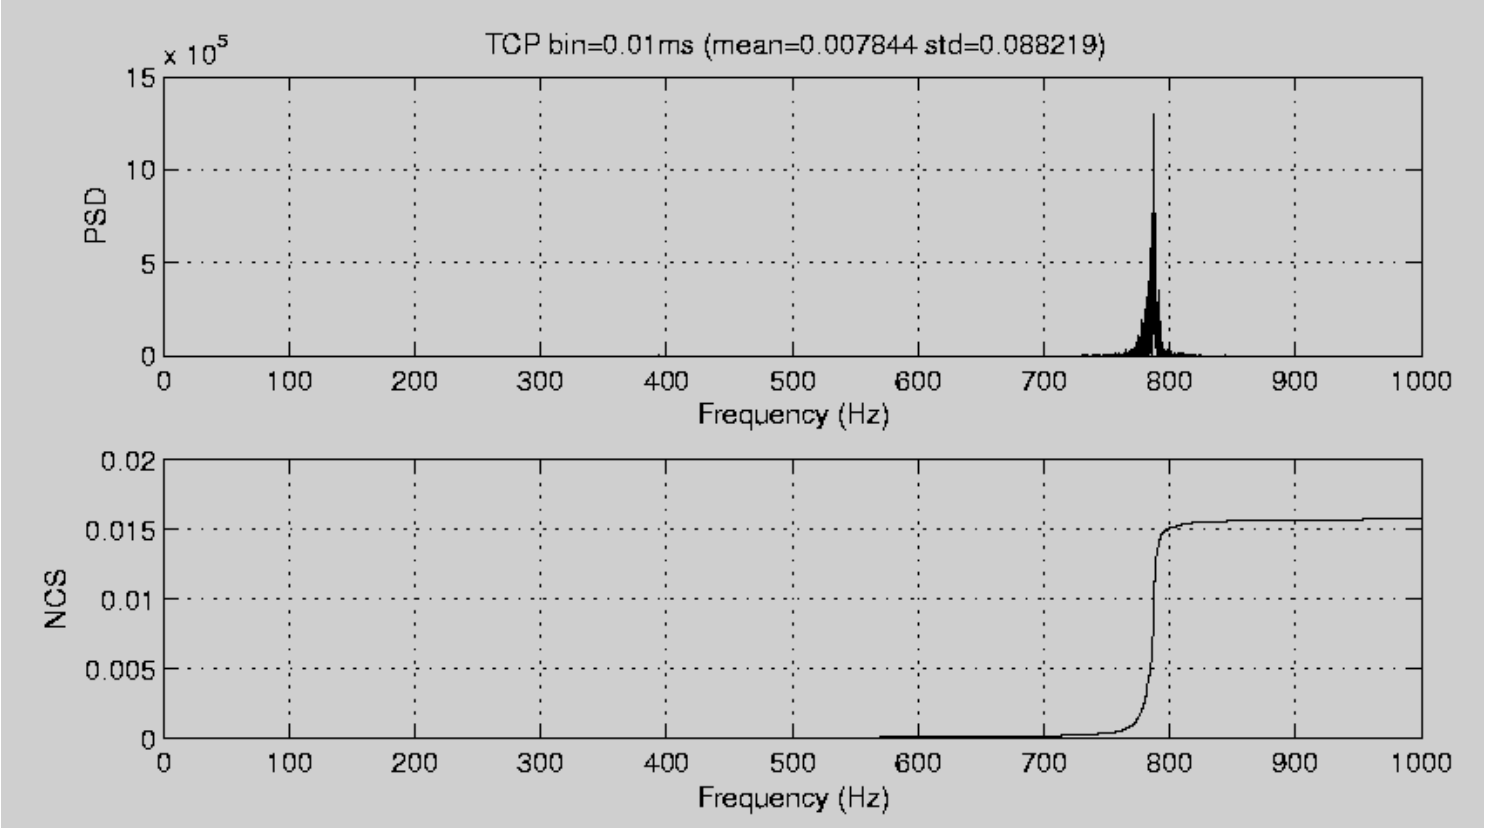

Signal at ~800Hz corresponding to back-to-back max-size packets

Does the Bottleneck Signal Survive?

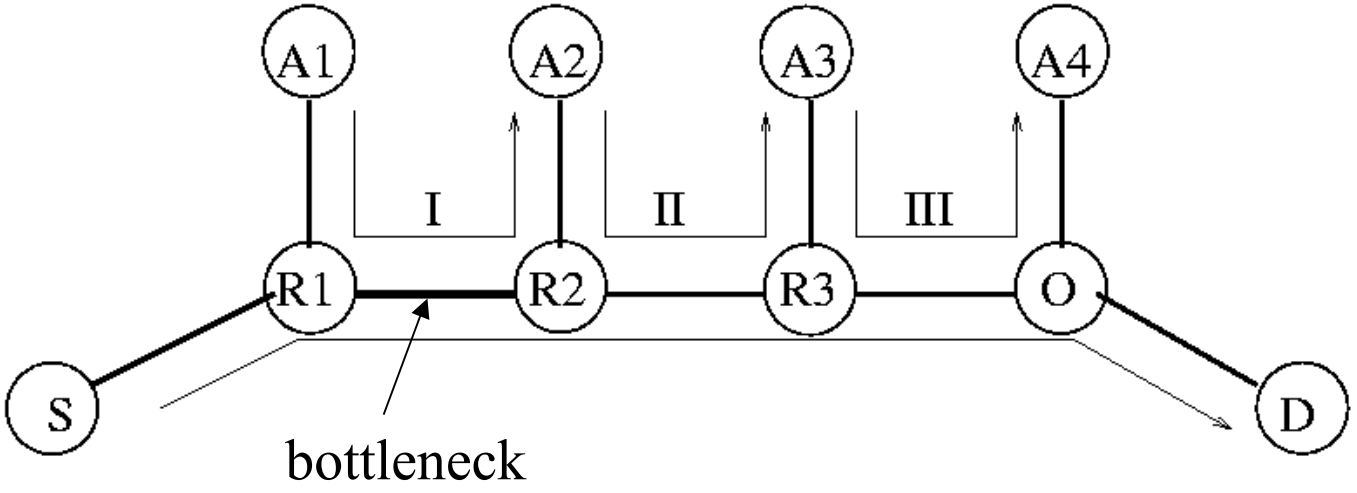

Experiment with three types of cross-traffic:

- Type I: shares bottleneck
- Type II: does not share bottleneck, not visible at observation point
- Type II: does not share bottleneck, visible at observation point

Cross traffic generated with surge (web traffic generator)

Experiment with Type I cross traffic (shares bottleneck)

Result: signal survives

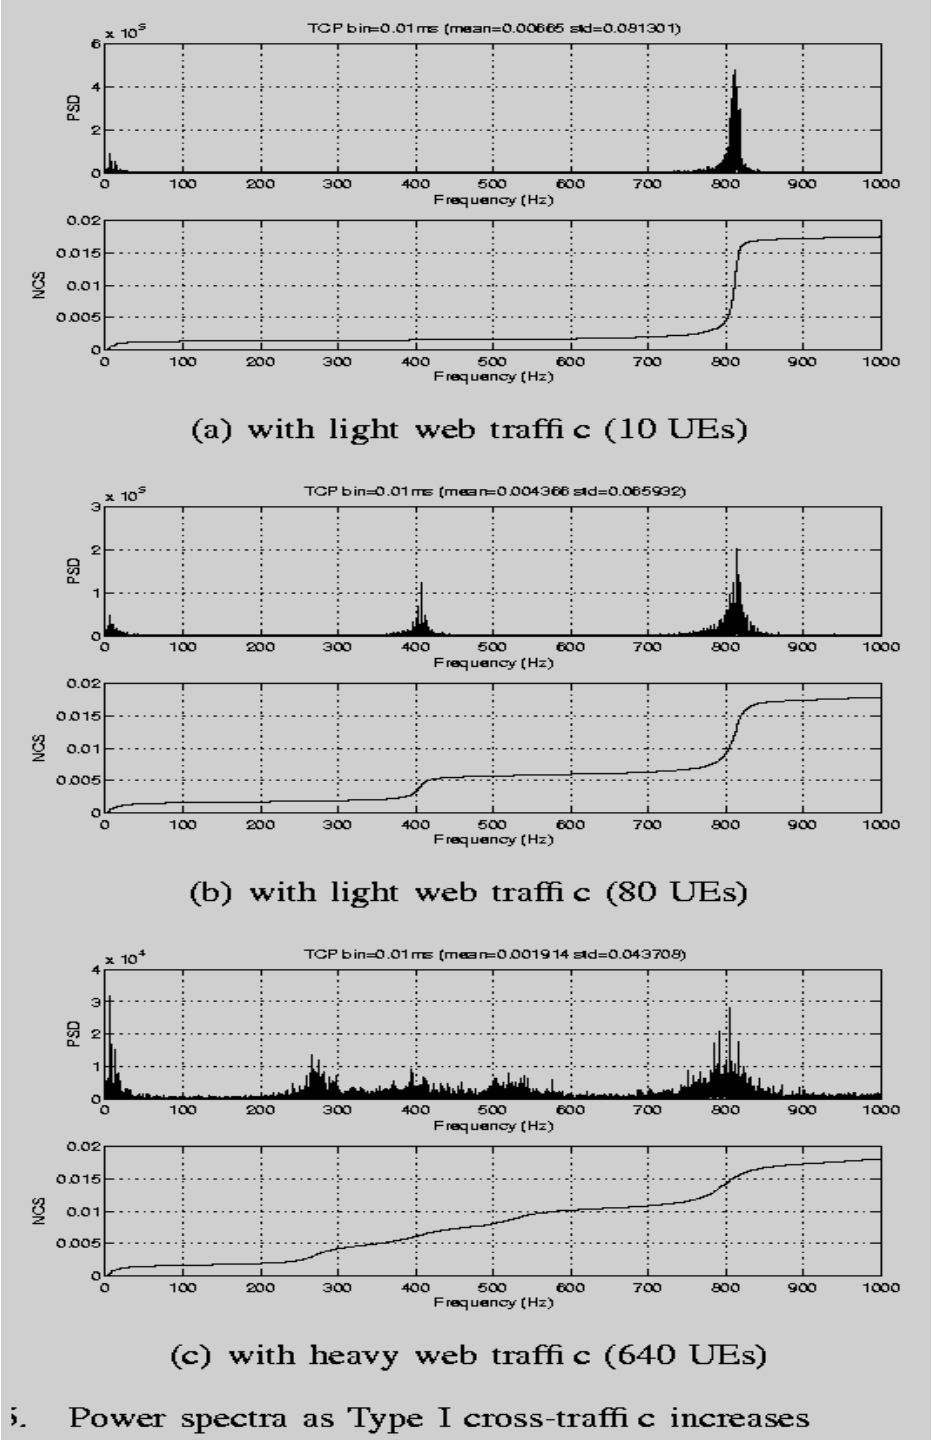

5. Power spectra as Type I cross-traffic increases

Experiment with Type II cross Traffic (does not share bottleneck, not visible at monitoring point)

Result:
Signal survives

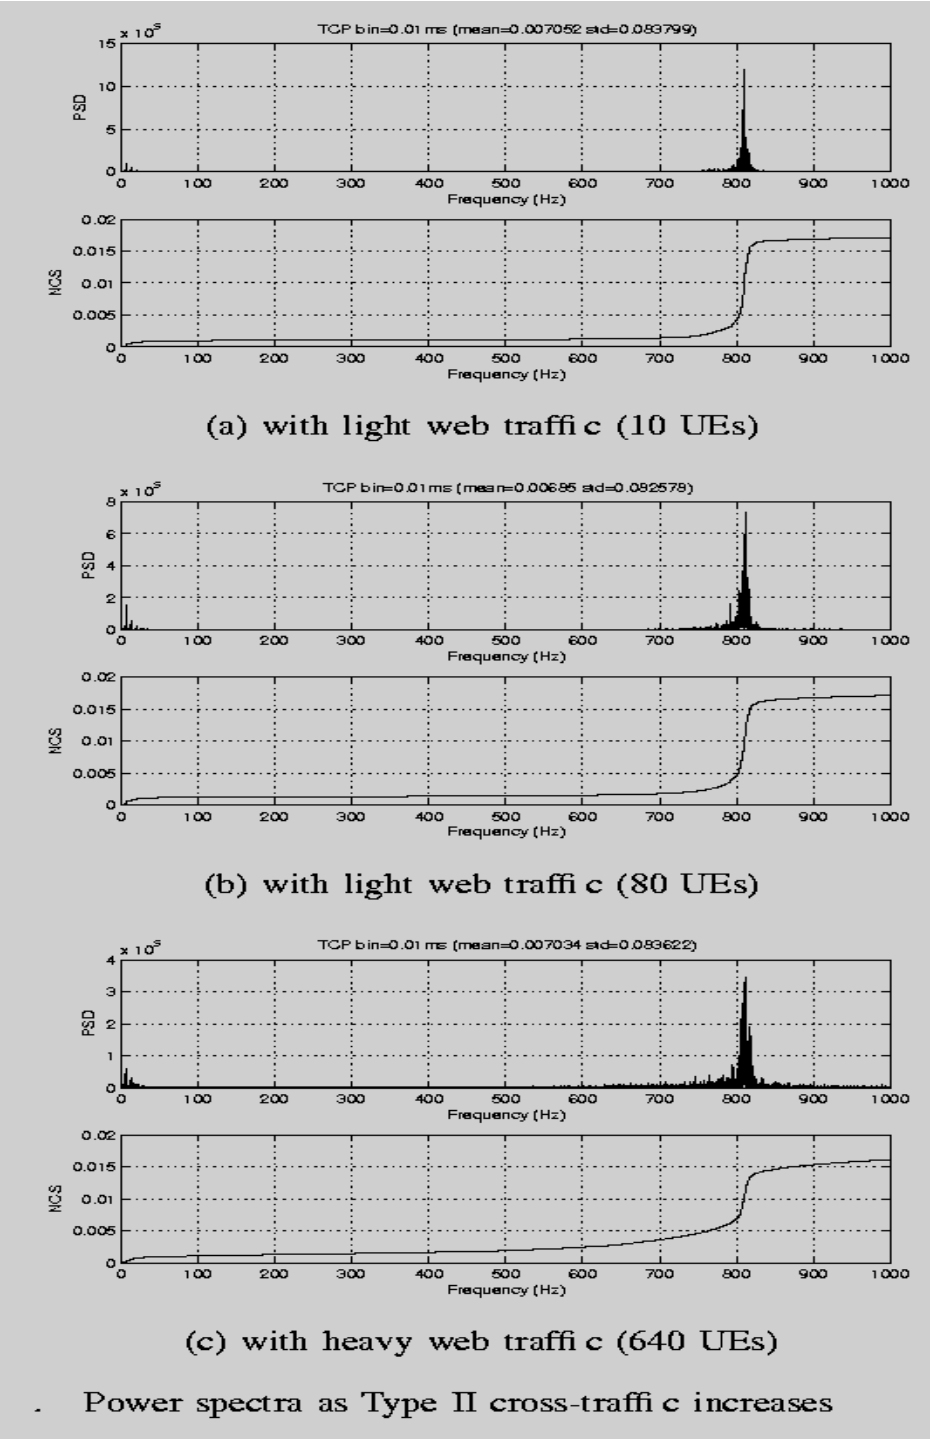

. Power spectra as Type II cross-traffic increases

Experiment with Type III cross Traffic (does not share bottleneck, visible at monitoring point)

Result:
Signal still detectable

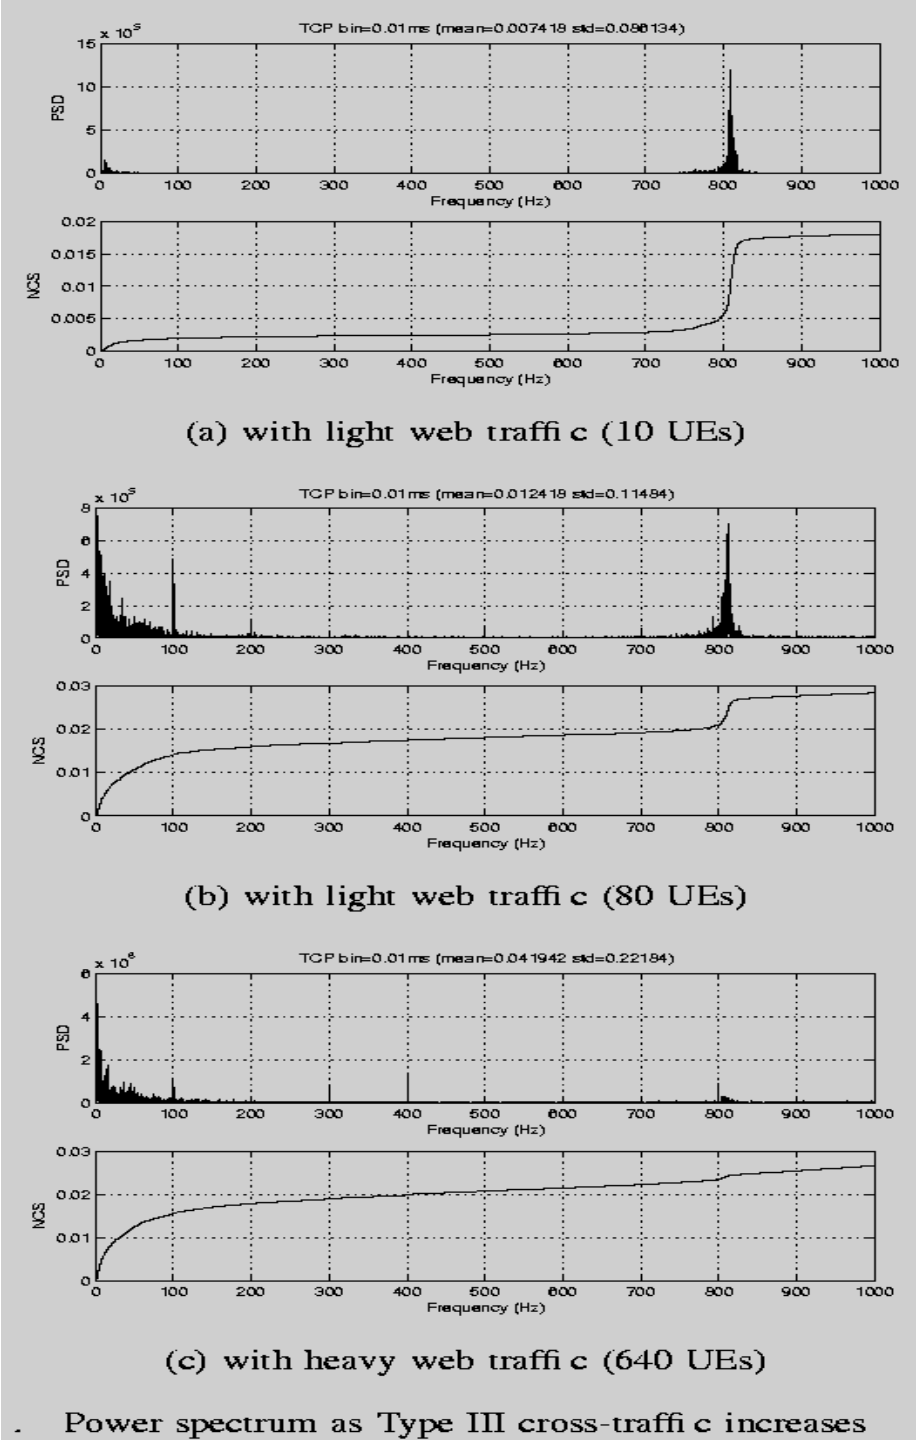

Internet II Experiments

Monitoring at our ISP
Artificially saturated a link on path from UCSB

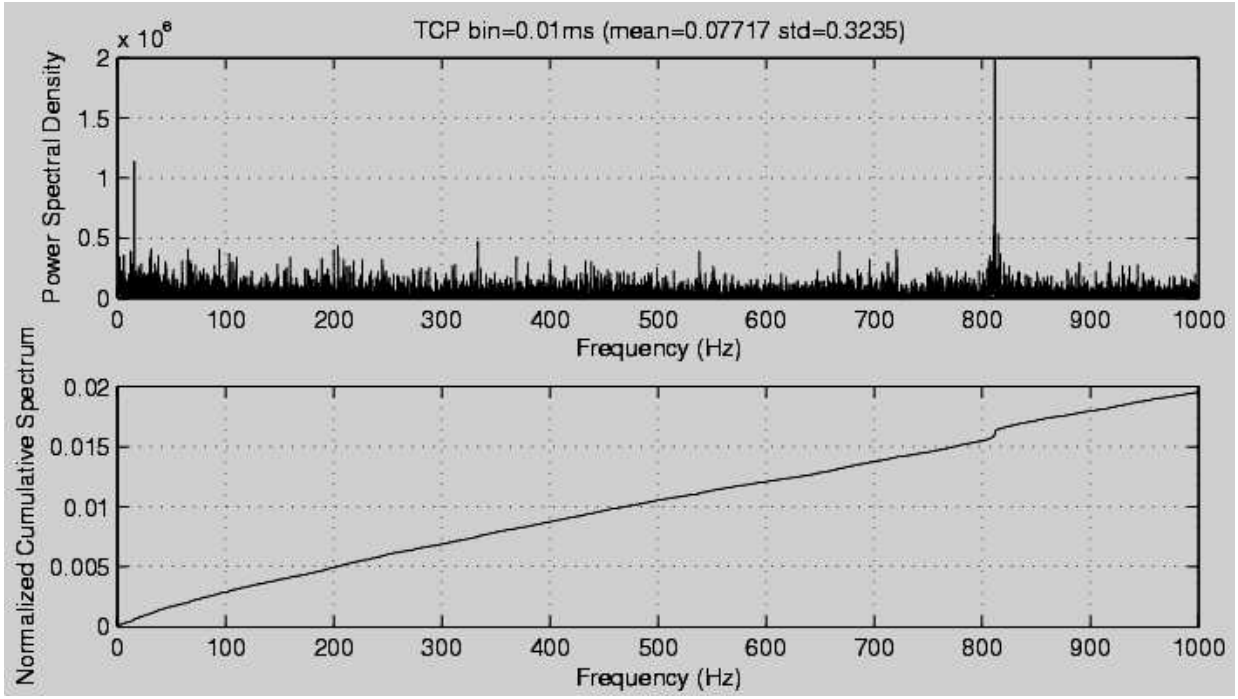

Our Bottleneck Detection Methods

- Leverage off existing techniques for signal detection in wireless transmission.
- Approach:
  - Train on traffic with and without bottleneck
  - Compare distributions:
    - Amplitude distribution of a single frequency
    - Top amplitude in a frequency band
    - All amplitudes in a frequency band
    - Top M amplitudes in a frequency band
- How about detection algorithms with no training?

Outline

- Detecting saturated links by looking at the aggregate
- Modeling packet-capture systems
- **Measuring bandwidth fluctuations**

Detecting BW Fluctuations in a Path

Goal: characterize periodic fluctuations in available BW in a path at small timescales (~10Hz)

- Characterize paths for demanding applications
- Diagnostic tool for network operators
- Investigate and characterize any transient and/or persistent phenomena
- Characterize Internet paths
- Work done by Rishi Sinha

Approach

- Sample available bandwidth continuously, using poison low rate packet-pair dispersion measurements (2-4% of link capacity).
- Create a timeseries of packet dispersion values.
- Average the timeseries using moving average window to eliminate high frequencies
- Determine any periodicities by applying FFT on the averaged time series.
- Work in progress..

Internet Validation

- Create artificial background traffic fluctuating at a known frequency.
- Attempt to detect the frequency.
- Experiments in the lab and on the Internet
- Results seem good – can detect frequencies up to 10 Hz by detecting peak amplitude in the FFT.

Periodicities are Sometimes Obvious

Experiment between USC and Umass
Duration: 60mins

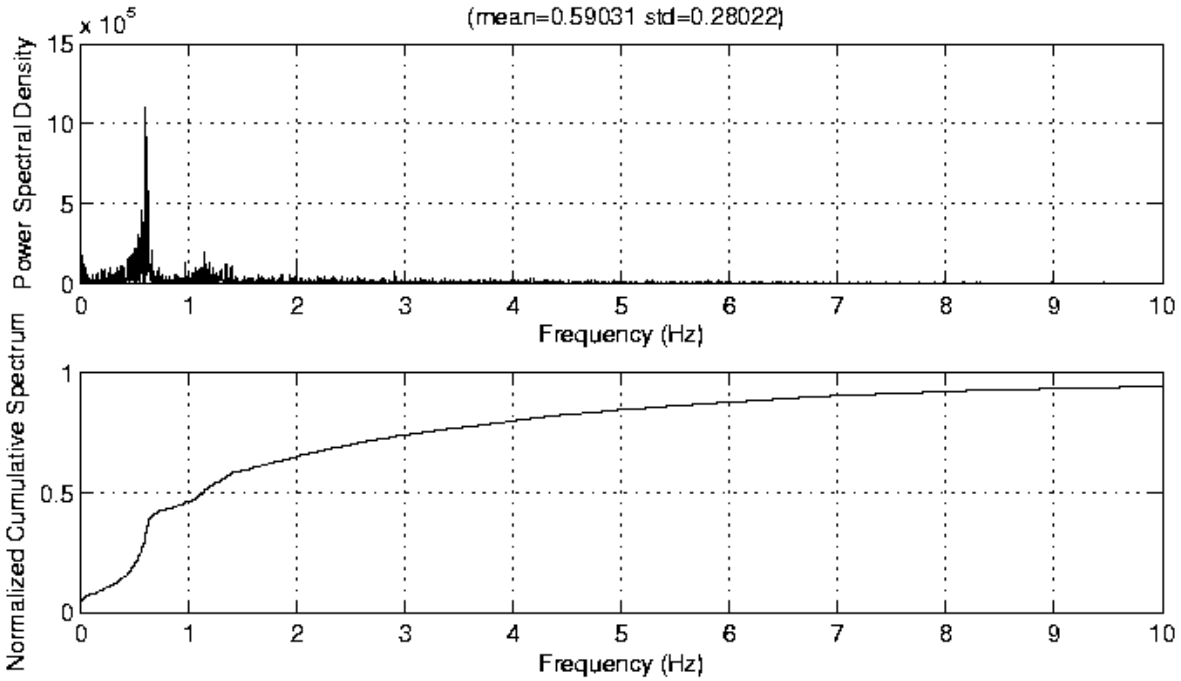

Future Directions

- Refine measurement modeling
- Refine our bottleneck detection algorithms and develop a tool
- Run BW fluctuation experiments over PlanetLab to characterize a larger set of paths

http://www.isi.edu/ant