Ethernet Adaptive Link Rate: System Design and Performance Evaluation

Chamara Gunaratne
Ken Christensen
Department of Computer Science and Engineering
University of South Florida
Tampa, FL 33620
{pgunarat, christen}@cse.usf.edu

This material is based upon work funded by the National Science Foundation under grant CNS-0520081.

Agenda

- Introduction
- ALR control policies
- Simulation performance evaluation
- Summary and future work

Agenda

- **Introduction**
- **ALR control policies**
- **Simulation performance evaluation**
- **Summary and future work**

Introduction

- **Internet is a major consumer of electricity**
  - Over 6 TWh/yr in 2000 (not including end nodes) [1]
- **One TWh/yr = $80 million**
  - 0.75 million tons of CO_2
- **Ethernet links operate at very low utilization**
  - Typical average utilization is less than 5%
- **Adaptive Link Rate (ALR) proposed in 2005**
  - Matches link rate to utilization
  - Energy saved by operating at a lower link rate most of time
    - Potential to save a large amount of electricity

[1] K. Roth, F. Goldstein, and J. Kleinman, "Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings, Volume I," Arthur D. Little Reference No. 72895-00, January 2002.

ALR overview

- **ALR *mechanism* to switch link rate**
  - For example, from 1 Gb/s to 100 Mb/s
    - Could be a MAC frame handshake
- **ALR *control policy* to determine when to switch**
  - For example, based on queue length thresholds
- **ALR is a trade-off between…**
  - Packet delay
  - Energy saved (time spent in low data rate)

The subject of this talk

- **ALR policy must prevent oscillation between rates**
  - Oscillation causes excessive packet delay

Agenda

- Introduction
- ALR control policies
- Simulation performance evaluation
- Summary and future work

Not in paper

ALR mechanism

- **Mechanism implements rate switching**
  - ALR implemented in both ends of an Ethernet link

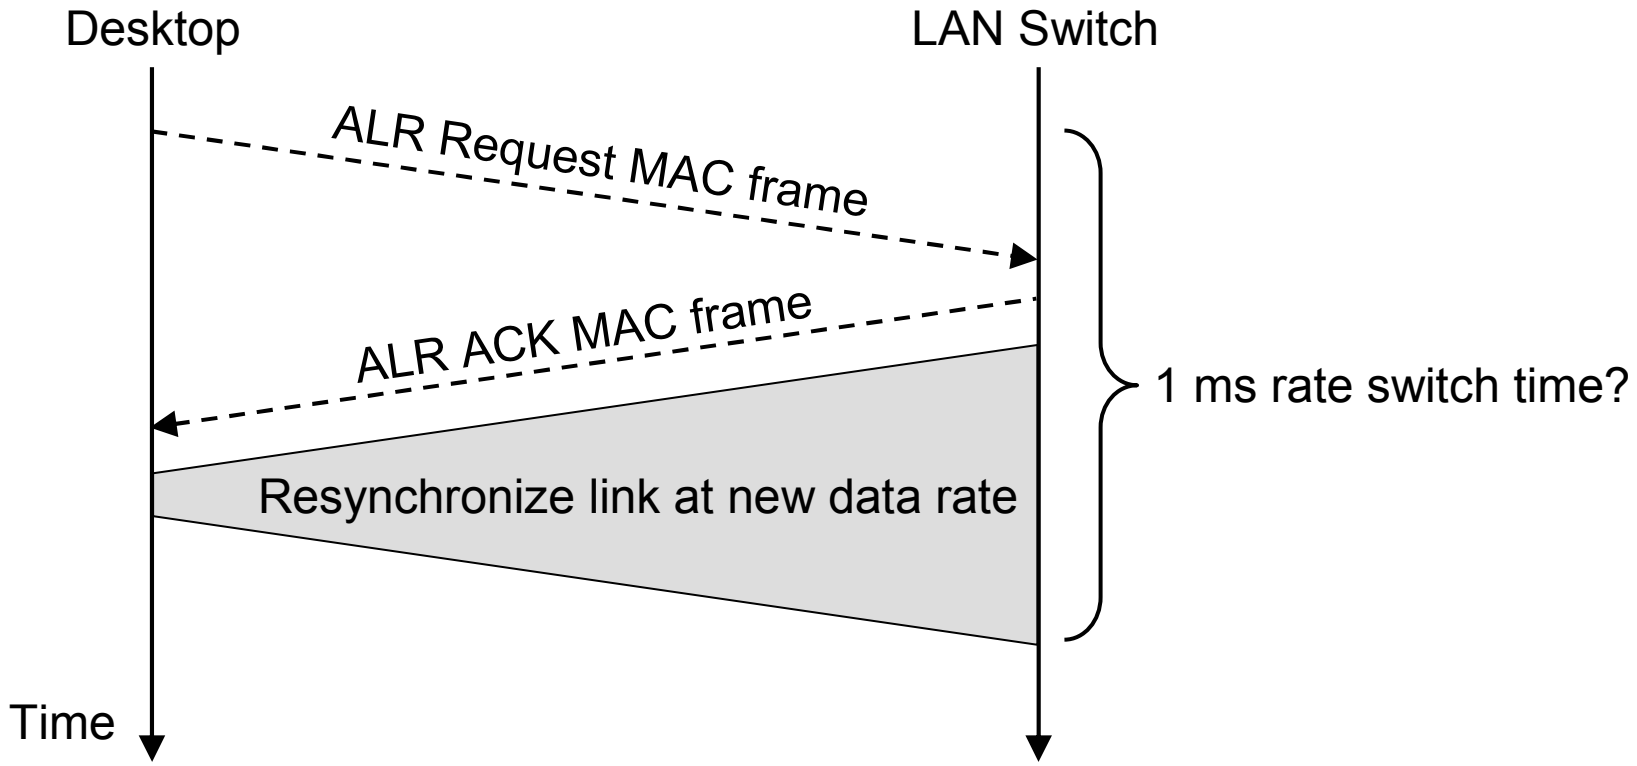

ALR control policy

- **Control policy used to determine when to switch**
  - Can use buffer levels and measured utilization
    - Buffers in PC NIC and LAN switch line card

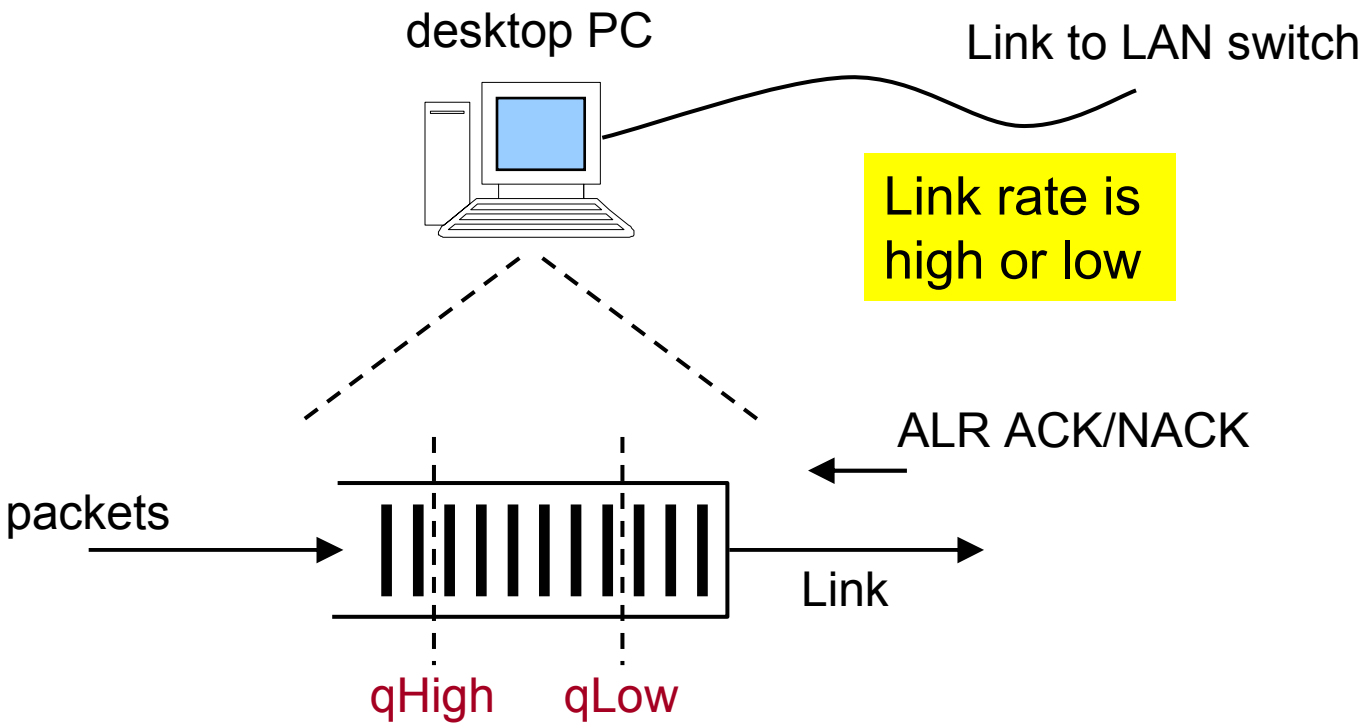

Two ALR control policies

- **Dual-threshold policy**
  - Use thresholds to switch link rate up and down
  - Use timers to stabilize policy

The subject of this talk

- **Utilization-threshold policy**
  - Use high threshold to increase link rate
  - Use explicit utilization measurement to reduce link rate

Performance of dual-threshold policy

- **Behavior of dual-threshold policy**
  - Poisson arrivals, fixed length packets (M/D/1)

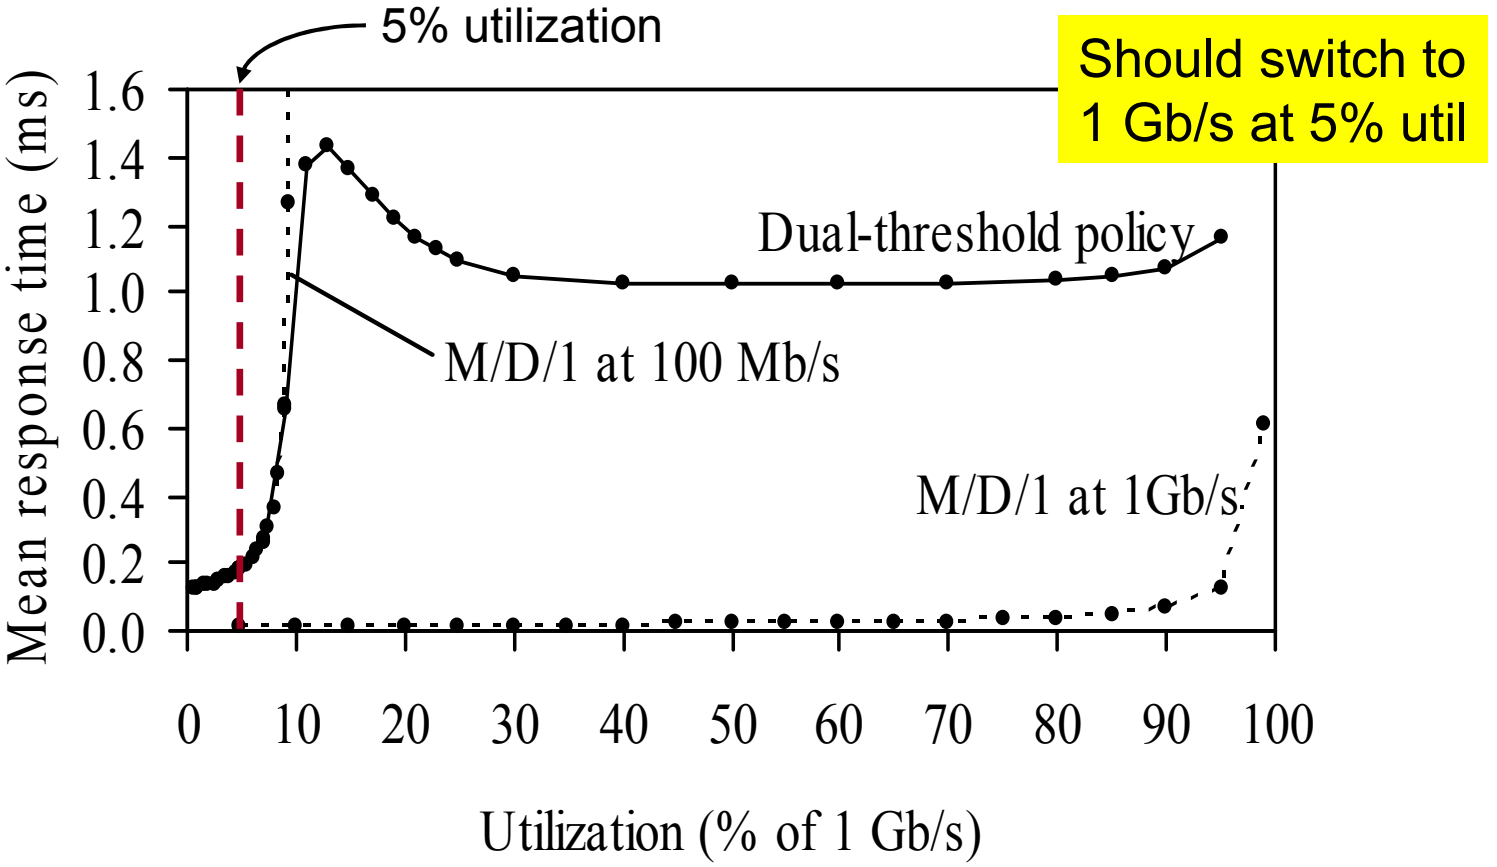

ALR utilization-threshold policy

- **Detailed FSM (see paper)**

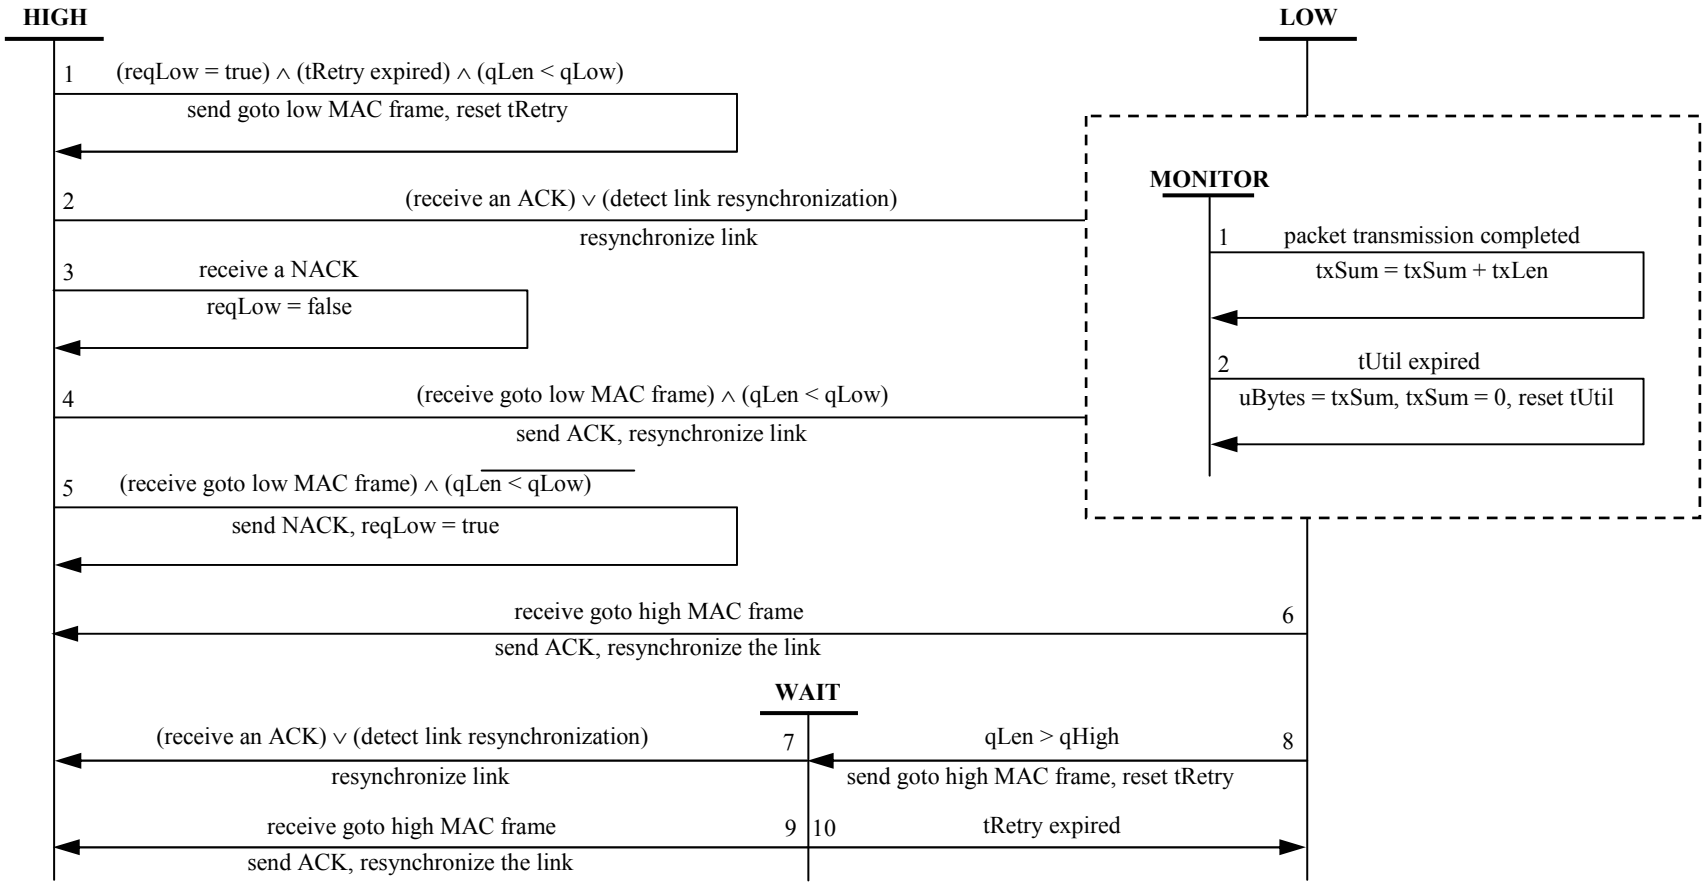

Not in paper

ALR utilization-threshold policy continued

- **Simplified description**

Executes on receiving a frame

```
if (link rate is low)
  if (buffer size exceeds high queue threshold)
    handshake for high link rate
```

Executes periodically at end of dampening time period (tUtil)

```
if (link rate is high)
  if (buffer size less than low queue threshold)
    if (bytes sent less than byte count threshold)
      handshake for low link rate
```

Agenda

- Introduction
- ALR control policy
- Simulation performance evaluation
- Summary and future work

Performance evaluation

- **Model ALR at the NIC using queues**
  - Built a queueing model using CSIM19
  - Realistic traffic model for input
- **Control variables**
  - Queue threshold value (low and high)
  - Utilization sample time (tUtil timer)
  - Utilization byte count threshold value (uThresh)
- **Response variables**
  - Mean packet delay
  - Mean time spent in low data rate

} Trade-off

Traffic model

- **Developed a synthetic traffic generator**
  - Able to match most characteristics of traced traffic
    - Including utilization, burst length, CoV, and Hurst parameter
  - Able to be “tuned” unlike trace traffic
- **Models bursty traffic**
  - Pareto burst size and exponential inter-burst time
- **Key parameters of traffic generator**
  - Data rate (100 Mb/s, 1 Gb/s, and 10 Gb/s)
  - Minimum and maximum burst size
  - Pareto index for bursts
  - Burst intensity
  - Packet length distribution
  - Mean utilization

Traffic model comparison

- **Measurement comparison to traced traffic**
  - **Synthetic versus traced for 100 Mb/s**

Close to the same

Characteristic	Actual	Synthetic
Mean inter-packet time (ms)	1.10	1.06
CoV of inter-packet times	1.76	3.81
Mean packet size (bytes)	577	526
CoV of packet size	1.16	1.15
Hurst parameter of packet counts	0.66	0.64
Utilization (% of 100 Mb/s)	4.11	3.95

Traffic model comparison continued

- **Visual comparison to traced traffic**

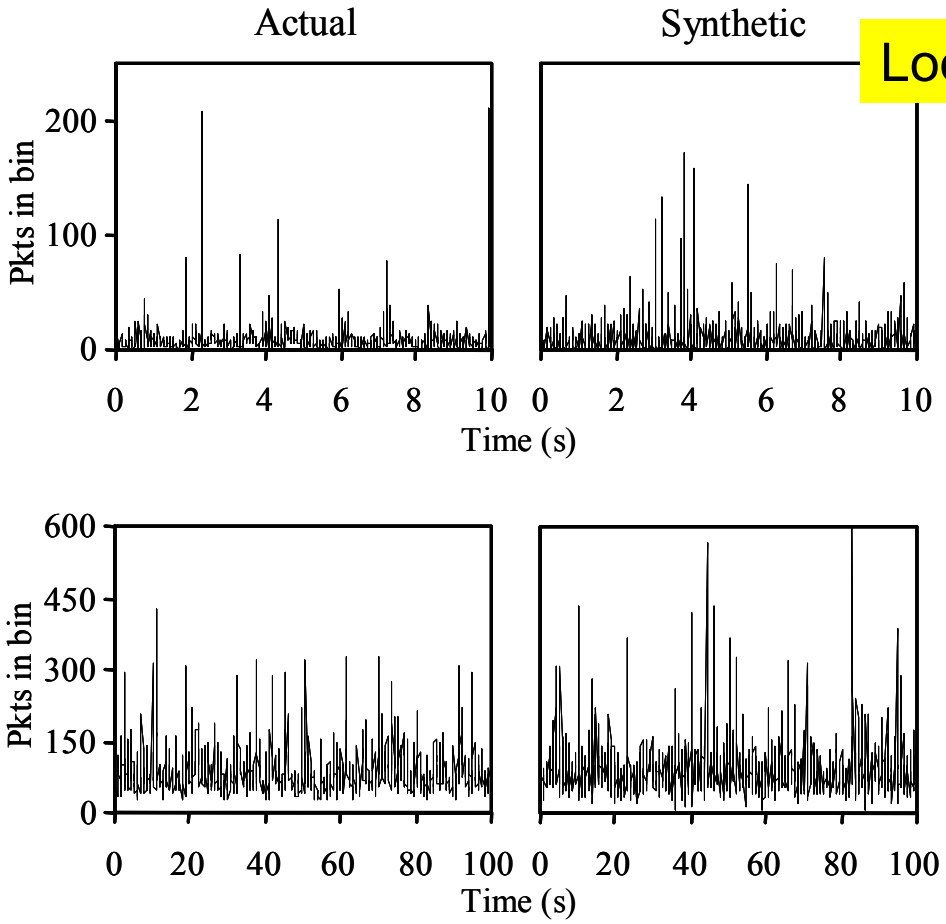

Simulation experiments

Rate switch time is 1 ms in all cases

- **Smooth traffic experiment**
  - Poisson arrivals, fixed length packets, vary the utilization
- **Single burst experiment**
  - Single burst (Poisson) at 80% utilization and 0.4 sec duration
- **Bursty traffic experiment**
  - Bursty (mean burst = 8.4 KBytes), vary the utilization
- **LAN switch experiment**
  - Use bursty traffic as input to a LAN switch model
  - Model used measured power use from Catalyst 2970
    - Base unit with no links = 46 W
    - Each 100 Mb/s link = 0.3 W and 1 Gb/s link = 1.8 W

Smooth traffic experiment results

- **Mean response time (packet delay)**
  - Note that dual-threshold does not come down in delay

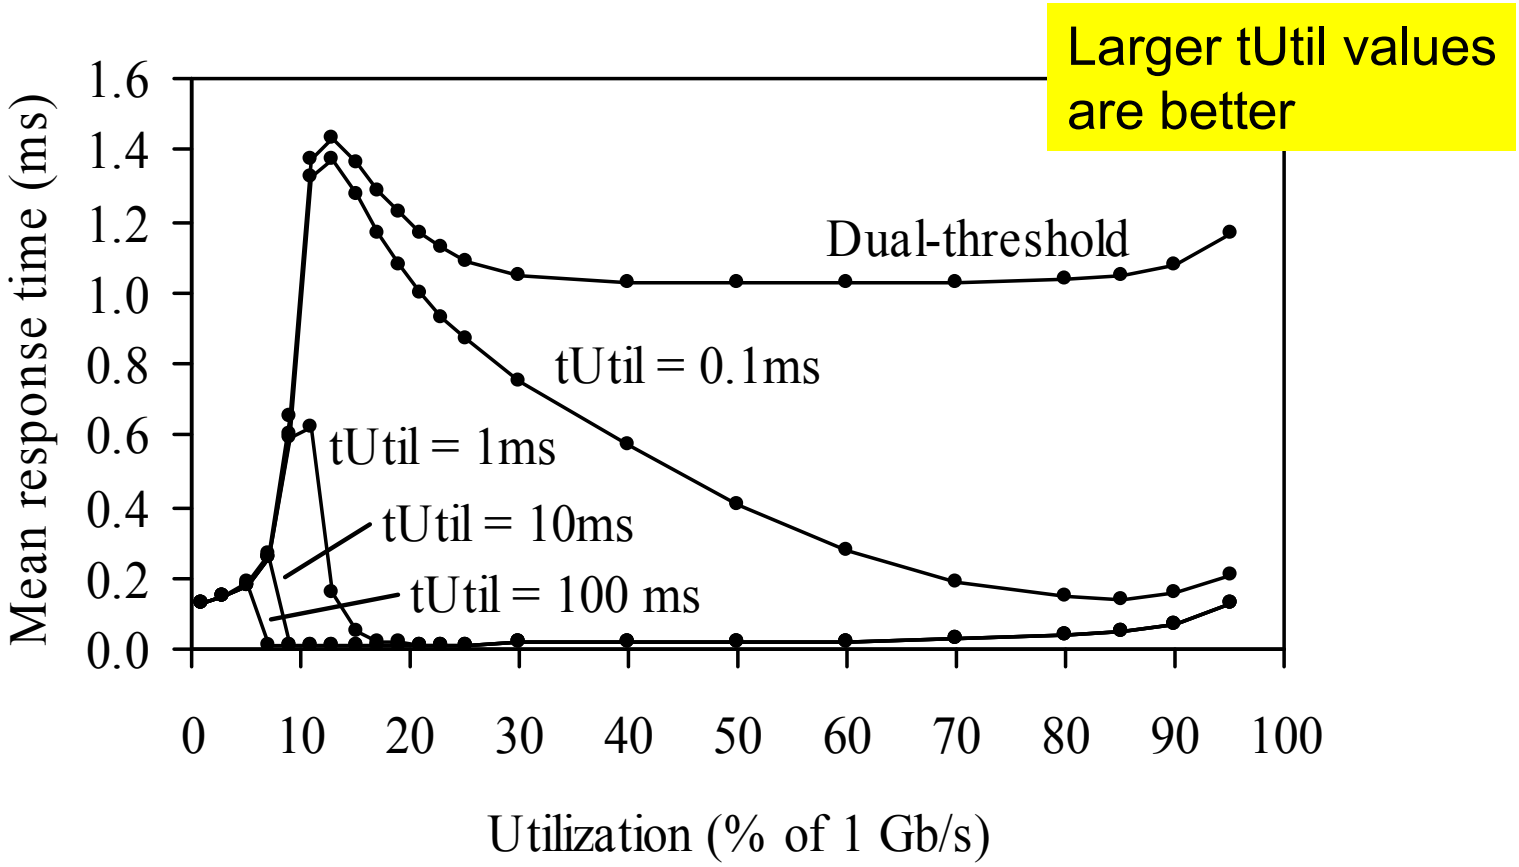

Single burst experiment results

- **Response time (packet delay) – transient behavior**
  - Note the oscillations (increased packet delay)

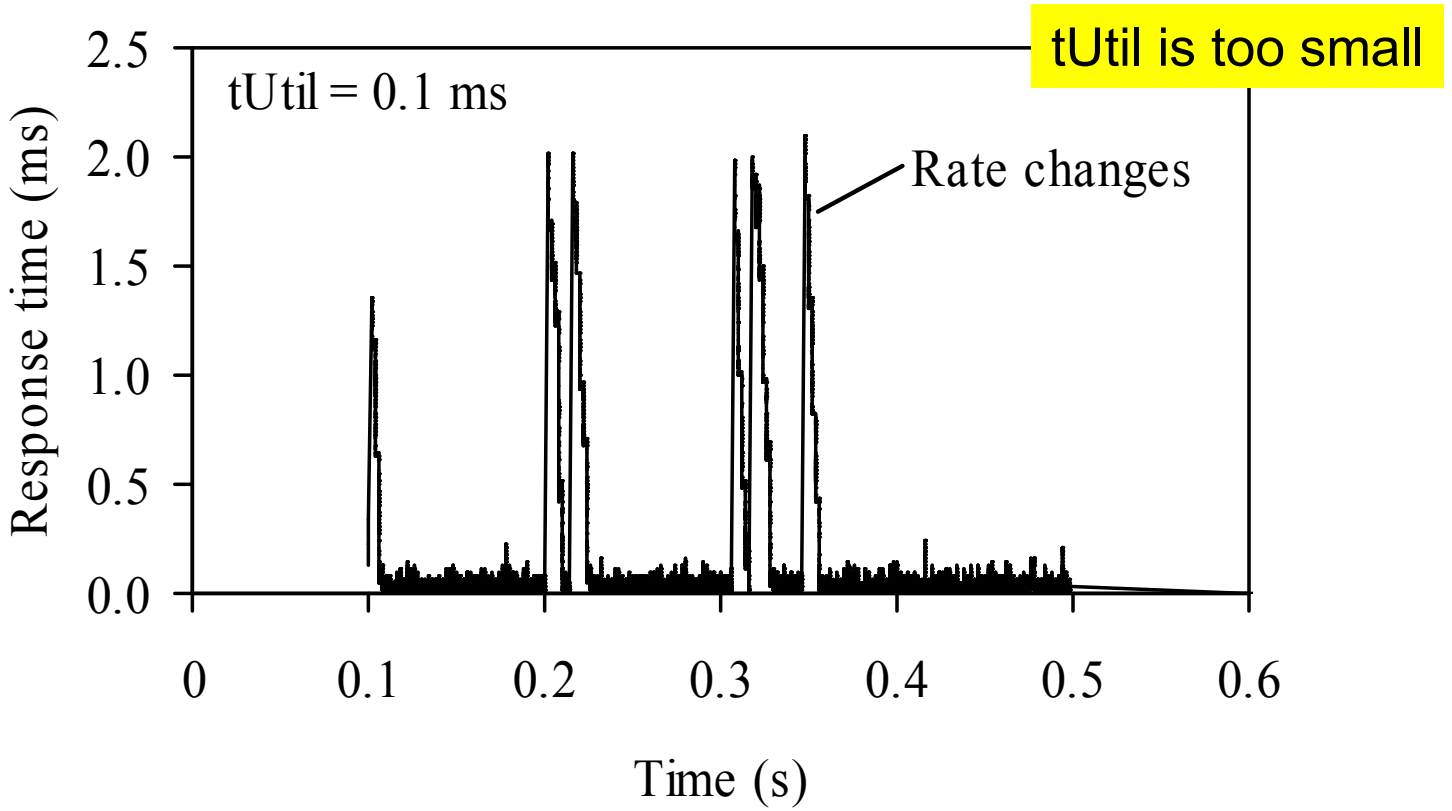

Bursty traffic experiment results

- **Mean response time (packet delay)**

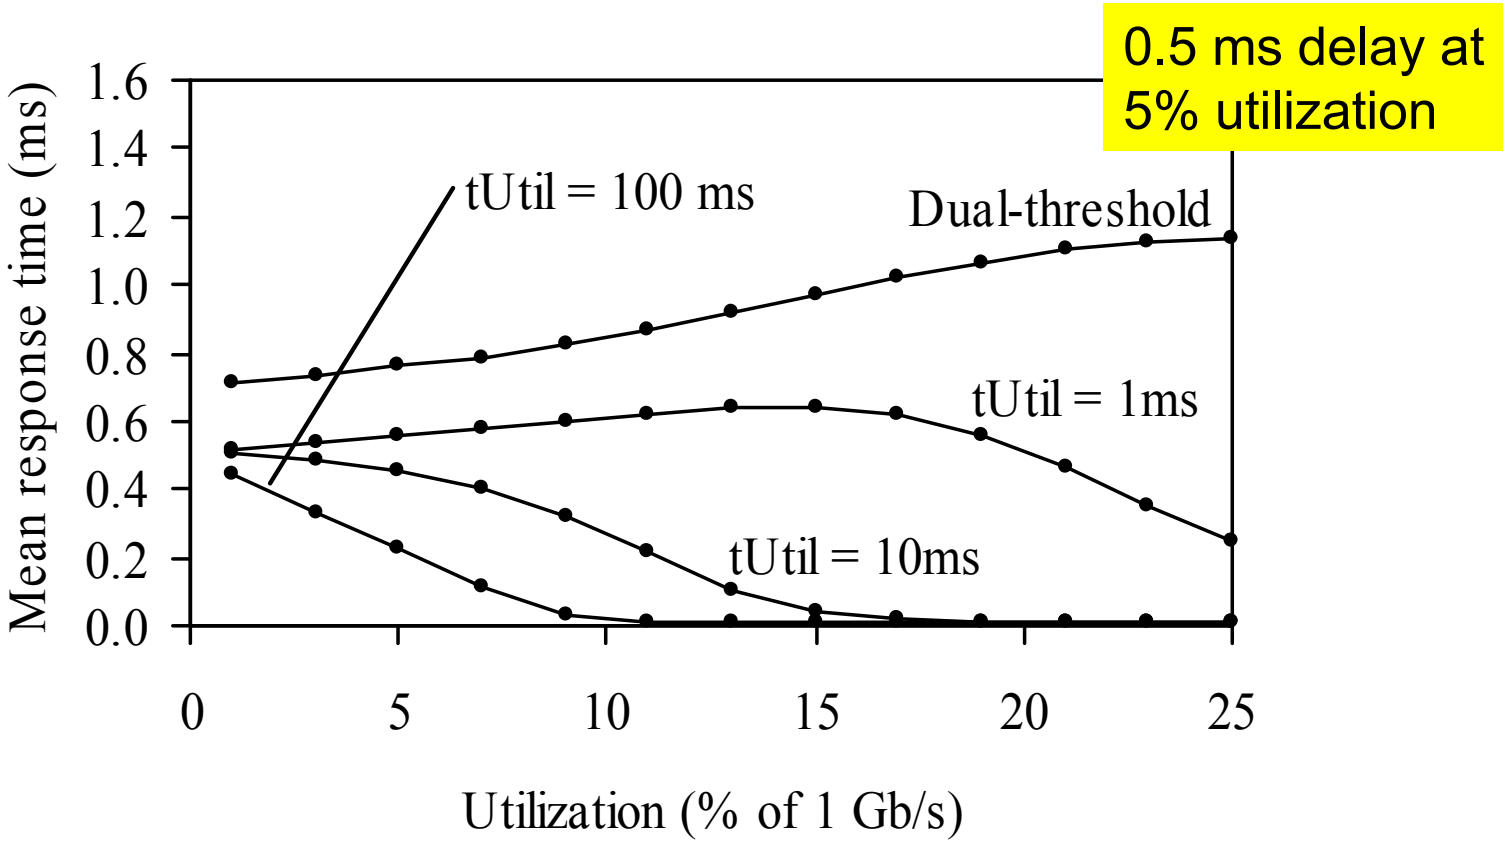

Bursty traffic experiment results continued

- **Time in low data rate (energy saving)**

80% time in low rate at 5% utilization

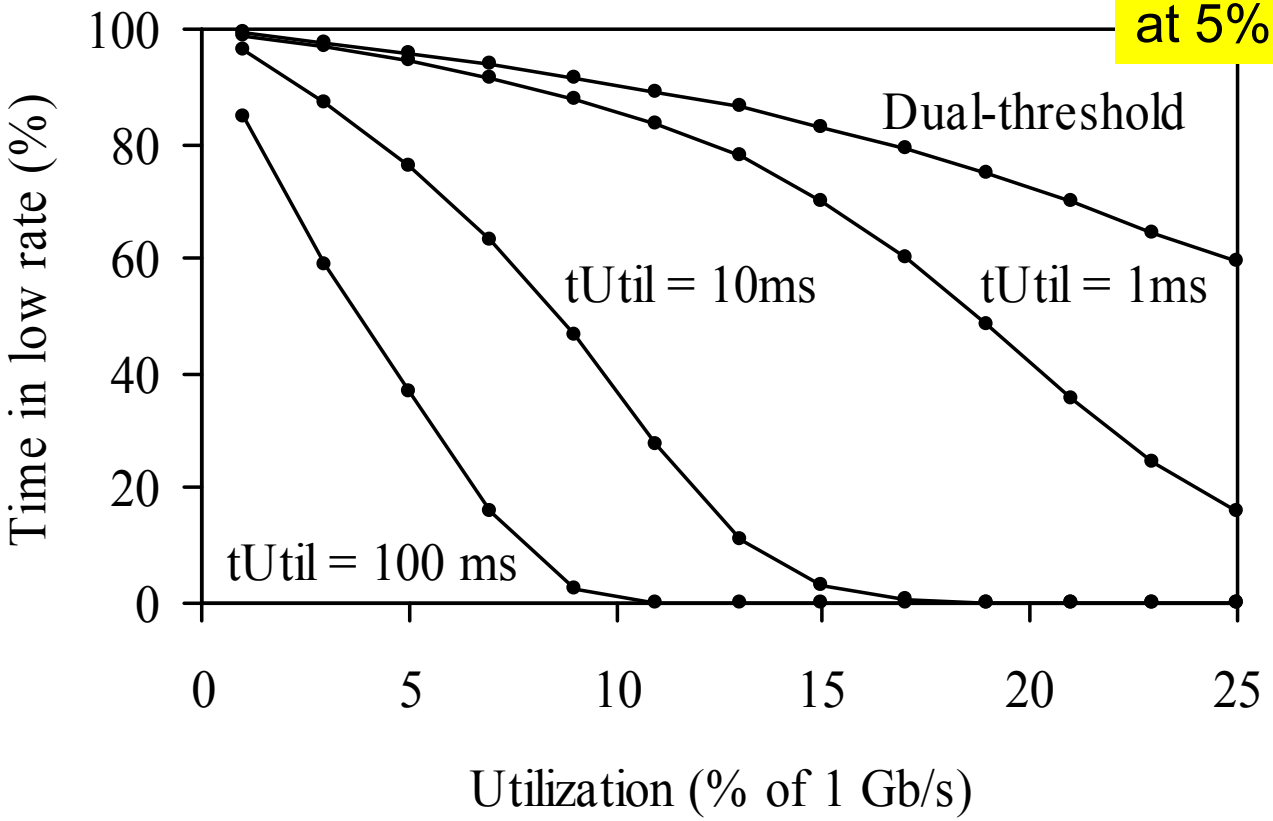

LAN switch experiment results

- **Power use and delay in one graph**
  - Shows trade-off

20% energy savings at 5% utilization with an added delay of less than 0.5ms

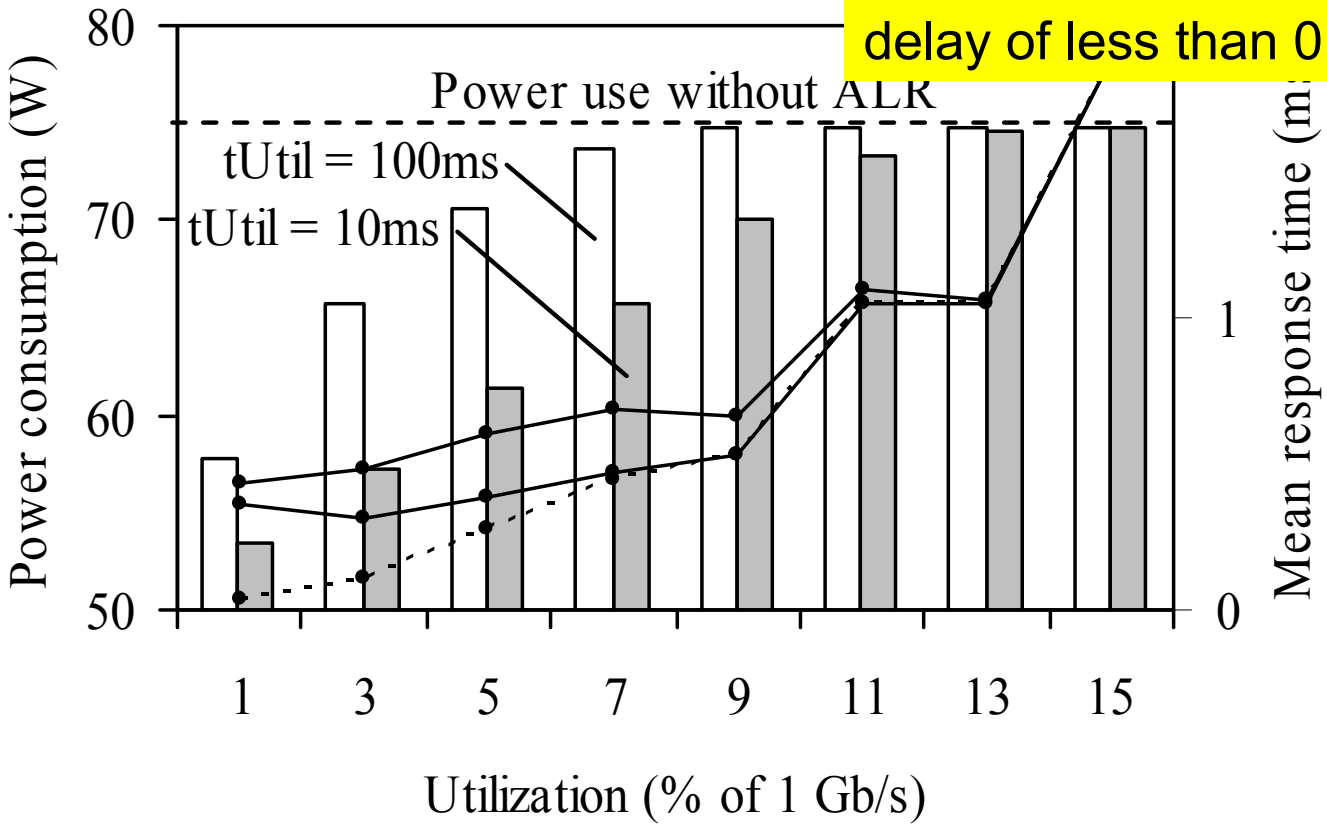

Agenda

- Introduction
- ALR control policies
- Simulation performance evaluation
- Summary and future work

Summary and conclusions

- **Investigated a utilization-threshold policy for ALR**
  - Prevents link from oscillating (and thus reduces delay)
- **At 5% utilization it is possible to…**
  - Be in low data rate for 80% of time
  - With mean packet delay of less than 0.5 milliseconds
- **Assume 100 million ALR-capable 1 Gb/s NICs**
  - Operates for 8 hours per day with 80% of time in 100 Mb/s
  - Save 4 W at 100 Mb/s (compared to 1 Gb/s)
  - Energy savings of 0.93 TWh/yr possible in the US
    - $75 million per year

Future work

- **Need to investigate deeper savings**
  - ALR may enable deeper saving in router line cards
    - For example F/V scaling of buffers, logic, etc.
- **ALR policies that use higher layer information**
  - Need to study impact on higher layer performance
- **ALR is moving towards standardization in 802.3**
  - Vote for study group at the 802.3 meeting in Dallas *last night*
    - 55 Yes, 2 No, 7 Abstain

Thank you!

- **I'll be happy to answer any questions**

Chamara Gunaratne
Ken Christensen
Department of Computer Science and Engineering
University of South Florida
Tampa, FL 33620
{pgunarat, christen}@cse.usf.edu