Reducing network energy consumption via sleeping and rate-adaptation

Sergiu Nedevschi

UC Berkeley

Joint work with:

Sylvia Ratnasamy (Intel Research), Lucian Popa (UC Berkeley),
Gianluca Iannaccone (Intel Research),
David Wetherall (Intel Research and U. Washington)

Overview

Last time: presented higher-layer algorithms that exploit hardware support for low-power operation [ratnasamy_1_1107.pdf]. Results were based on:

- real-world measurements of network traffic (Abilene Internet2 backbone, Intel corporate network)
- simulated models of equipment power profiles

This talk: updated results based on measured power profiles of real-world equipment

- Intel NIC [hays_01_1107]
- Cisco GSR [Chabarek, Infocom08]
- (would welcome additional data on equipment power consumption)

recap 11/07 meeting: overview

Goal: save energy <u>without compromising performance</u>

- achieving this will depend on:
  - appropriate hardware-level support for power management
  - higher-layer algorithms that invoke this support wisely
- our study
  - model two forms of hardware support
    - sleep states (low-power idle)
    - rate states (subset PHY (?))
  - design, evaluate higher-layer algorithms
  - explore how hardware support impacts savings/performance

recap 11/07 meeting: solutions

Presented two higher-layer algorithms that exploit hardware support for energy savings with controlled impact on performance

1. saving energy via sleeping: “buffer then burst”
   - sources buffer packets for up to B*ms*, then transmit buffered packets in a burst; switches sleep between bursts
   - buffer interval (**B *ms***) controls the tradeoff between energy savings (*i.e*., sleep time) and performance (*i.e*., added delay)

2. saving energy via rate adaptation: monitor queue lengths
   - adapt rate if doing so doesn’t add more than D*ms* delay to packets
   - delay bound (**D *ms***) controls the savings-performance tradeoff

recap 11/07 meeting: conclusions

- simple, practical higher-layer algorithms are effective in navigating the savings *vs.* performance tradeoff

- both sleep and rate adaptation are useful, but in different circumstances
  - sleeping typically better at low network utilizations; rate-adaptation at higher utilizations
  - crossover utilization depends greatly on equipment power profile

- rate-adaptation with uniform rates (e.g., R/4, R/2, 3R/4, R) enables higher savings than with exponential rates (R/100, R/10, R)

- low system-wide transition times are critical to maintaining acceptable performance

Outline

- Evaluation methodology
- Equipment power profiles
- Test results:
    1. sleep vs. rate-adaptation for Intel NIC
    2. sleep vs. rate-adaptation for Cisco GSR
    3. impact of system transition times
    4. impact of asymmetric operation

Evaluation methodology

- packet-level simulation (ns2)

- using real network topologies and traffic workloads
  - Abilene backbone
  - Intel enterprise network
    (scale measured traffic to explore effect of network utilization)

- based on measured power profiles for real equipment
  - Intel NIC [hays_1107]
  - Cisco GSR router [Chabarek, Infocom08]

- metrics
  - % energy savings
  - performance: 98 percentile delay

Equipment Power Profiles

	Intel NIC (R=1Gbps)	Cisco GSR (R=10Gbps)
$p_{active}(R)$	1217 mW	200+ 80W/linecard
$p_{idle}(R)$	1010 mW	200+70W/LC
$p_{active}(R/10)$	483 mW	200+32W/LC
$p_{idle}(R/10)$	314 mW	200+22W/LC
$p_{active}(R/100)$	504 mW	200+33W/LC
$p_{idle}(R/100)$	194 mW	200+13W/LC
p_{sleep}	65 mW	200W

extrapolated

chassis only

Sleep *vs.* Rate-adaptation: Intel NIC

Abilene backbone; transition time=1ms; rates=1G/100M/10Mbps

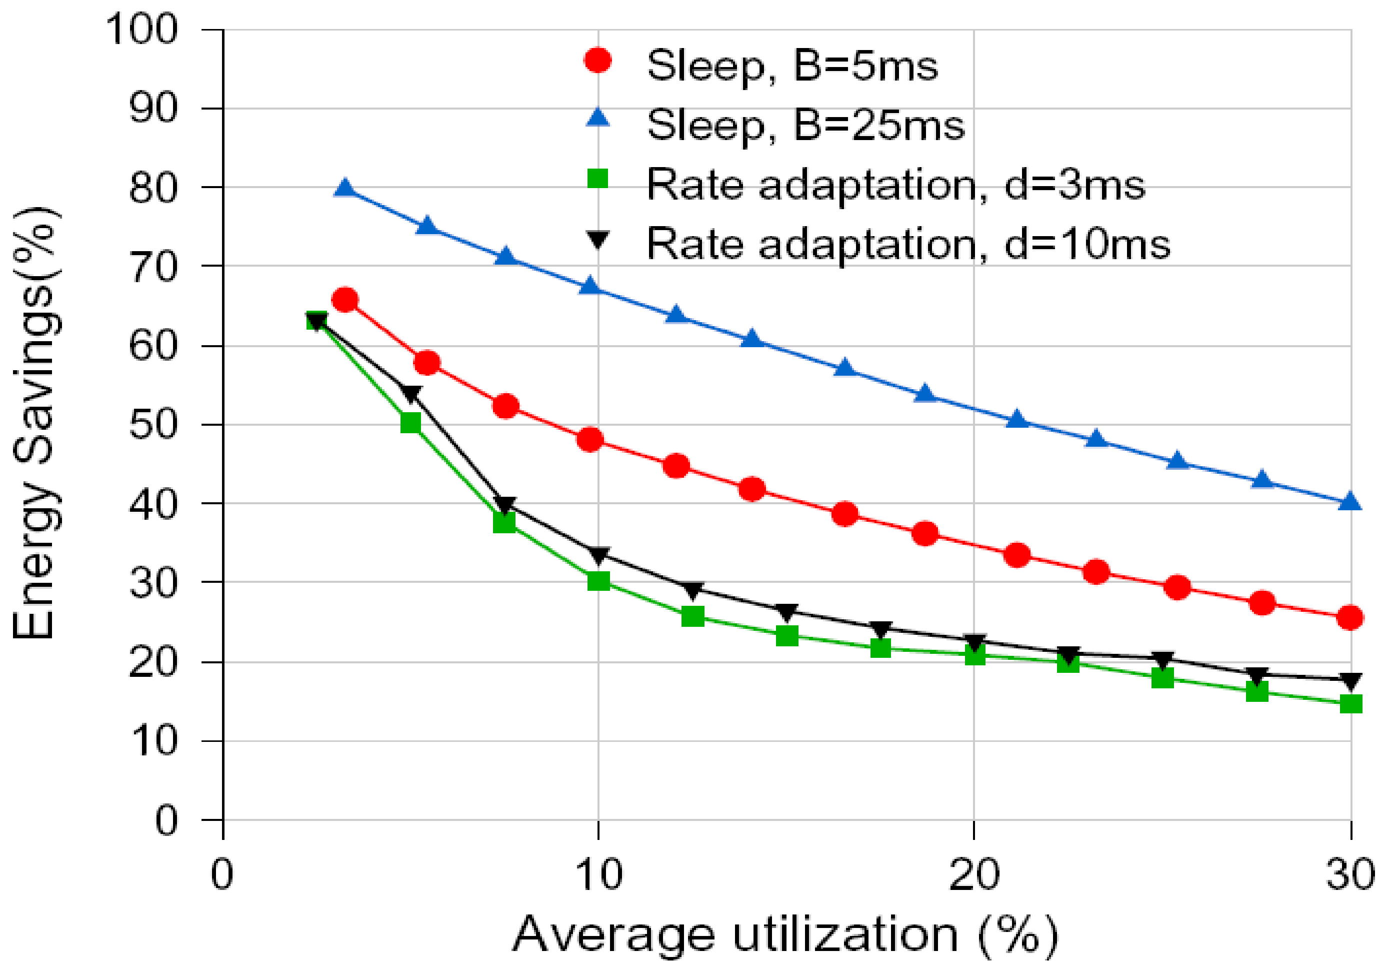

Sleep *vs.* Rate-adaptation: Cisco GSR

Abilene backbone; transition time=1ms; rates: 10G/1G/100Mbps

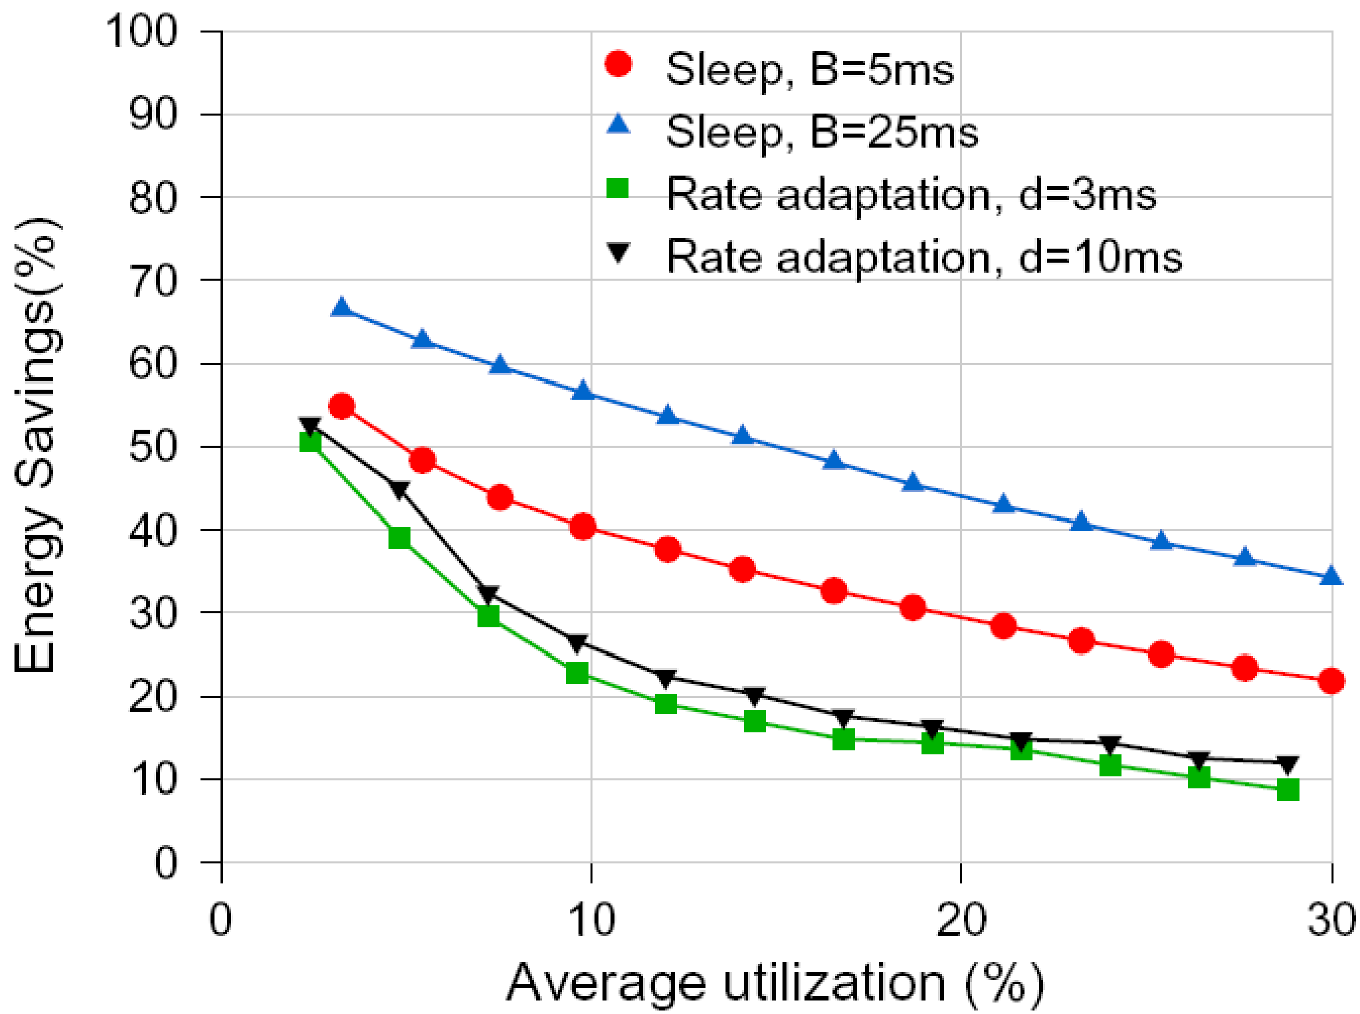

Sleep *vs.* Rate-adaptation: Cisco GSR

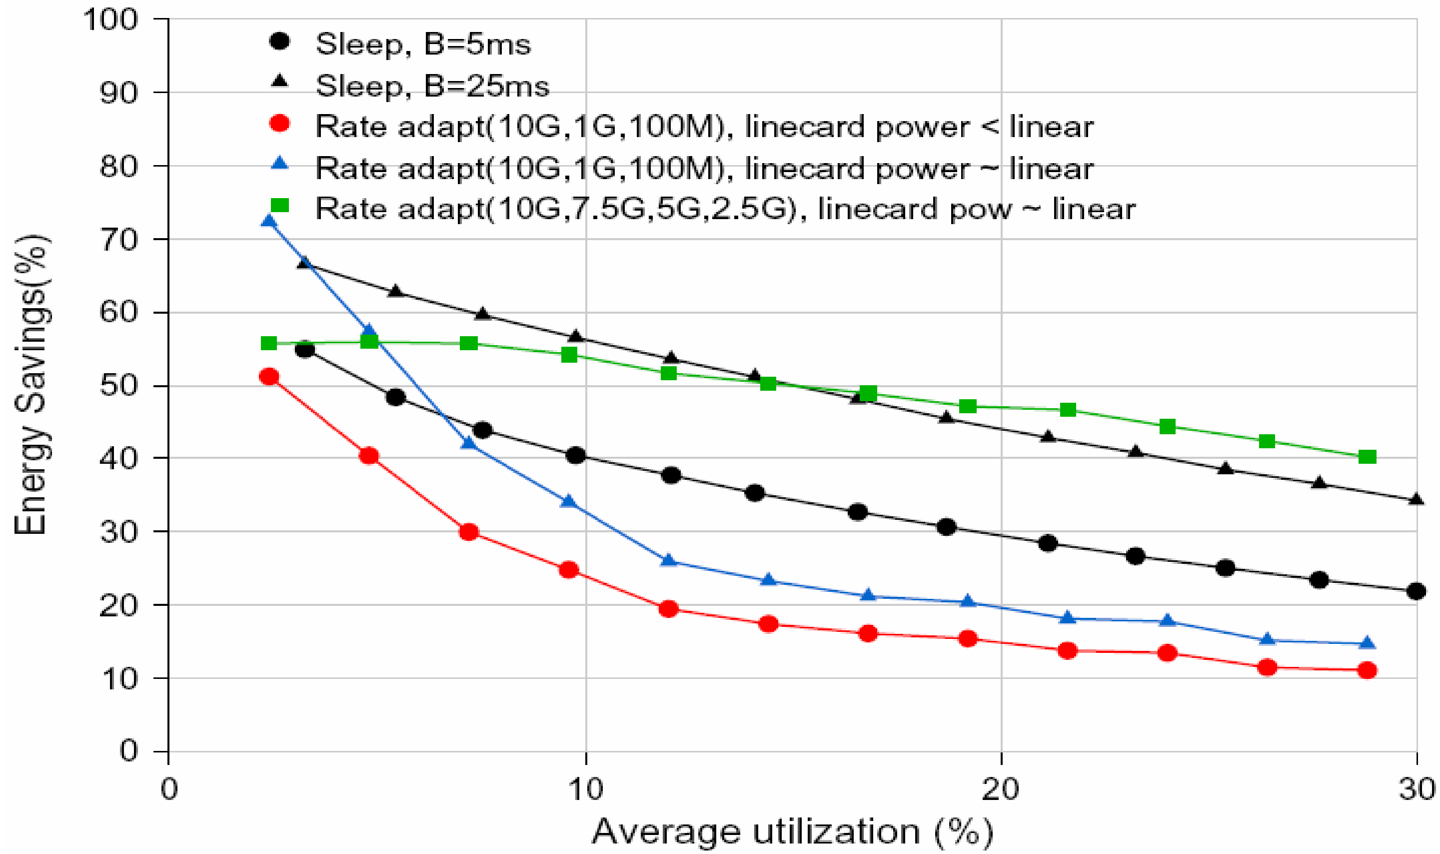

support for uniform rates (R/4, R/2, 3R/4,R) would greatly improve the savings from rate-adaptation

Impact of transition times

Using buffering

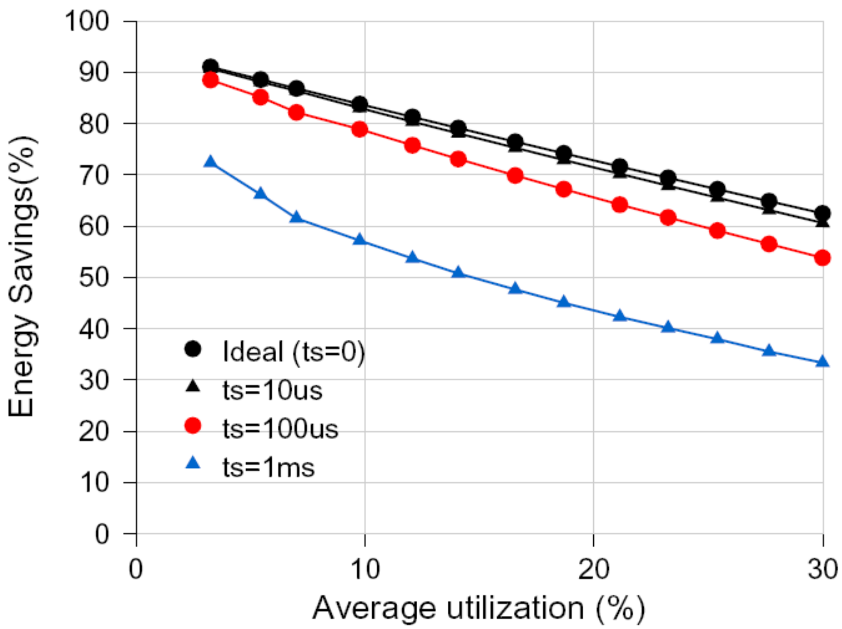

No buffering / traffic shaping

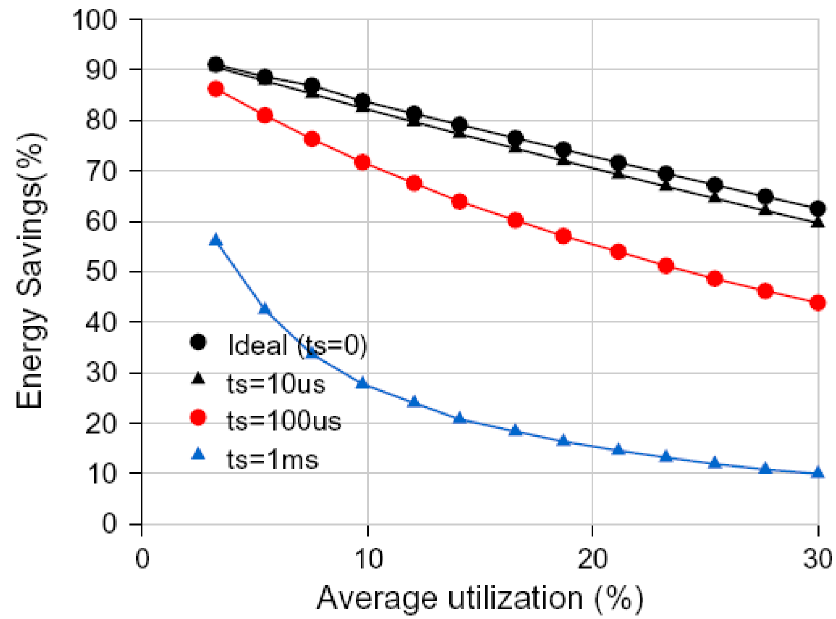

Traffic shaping is critical to achieving a good savings/performance tradeoff at higher system transition times