Performance Mapping

from cells to applications

Simon S. Lam

Department of Computer Sciences
The University of Texas at Austin

lam@cs.utexas.edu

http://www.cs.utexas.edu/users/lam/

The Problem

- **Observation**

  To an application, the end-to-end delays and loss rate of its data units are more relevant performance measures than ones specified for cells

- **Application data unit** (e.g., video picture)

  carried in an IP packet or multiple IP packets

- **IP packet**

  carried in multiple ATM cells

- **Performance mapping**

  cell-level $\rightarrow$ packet-level $\rightarrow$ application-level

Different Approaches

- **Analyzing application-level performance**
  - no additional support
- **Improving application-level performance**
  - early packet discard
  - ATM block transfer (ABT)
  - forward error correction
- **Guaranteeing application-level performance**
  - Burst Scheduling Networks
  - OS support in end systems

Modeling and Simulation

- **Derive and measure application-level performance from cell-level guarantees**

  [*Cidon et al.*]

- **Message loss probability differs greatly from cell loss probability and depends on**
  - peak rate and message size of connection
  - whether cell loss guarantee is per-connection or global

Early Packet Discard

- **Discard entire packets rather than individual cells**

  improves packet loss probability
  for the same cell loss probability

- **Admission of packet based upon buffer utilization**

  several algorithms [*Romanow-Floyd, Turner*]

- **Additional support**
  - per-connection state information
  - packet boundary recognition

ATM Block Transfer (Immediate Transmission)

- **Flow is a sequence of blocks**

  each block bracketed by two RM cells

- **Initial connection setup with peak cell rate but no bandwidth allocation**

- **Entire block admitted or discarded**

  improves block loss probability

- **Block admission based upon bandwidth requested in leading RM cell**

Forward Error Correction

- **Example**
  - redundant parity cell per block
  - destination can reconstruct block from n out of $n+1$ cells

- **Both block delay and loss probability of FEC flow *may* be improved**

  worse for non-FEC flows !

- **Additional support**
  - in end systems
  - in switches to selectively discard cells

Burst Scheduling Networks [*Lam & Xie*]

- **Similar to ABT but with** *QoS guarantees* **for blocks (bursts)**

- **Additional support**
  - support for ABT
  - rate-based packet scheduling
  - jitter control

- **Real-time VBR service**
  - worst-case delay of blocks is bounded
  - loss rate can be specified (≥ 0)

    determines extent of overbooking

Channel Utilization vs. Loss Rate

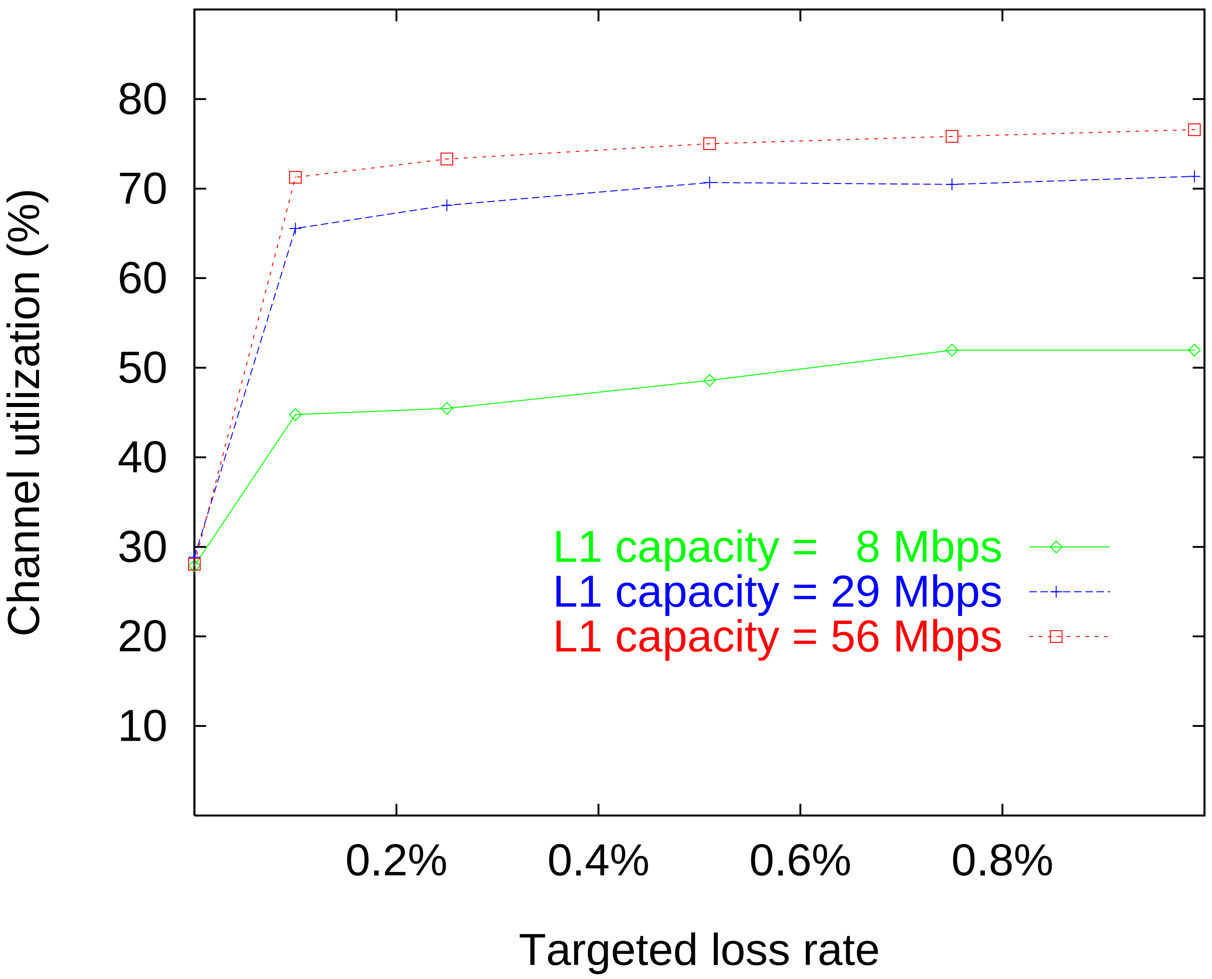

Central Limit Theorem at Work

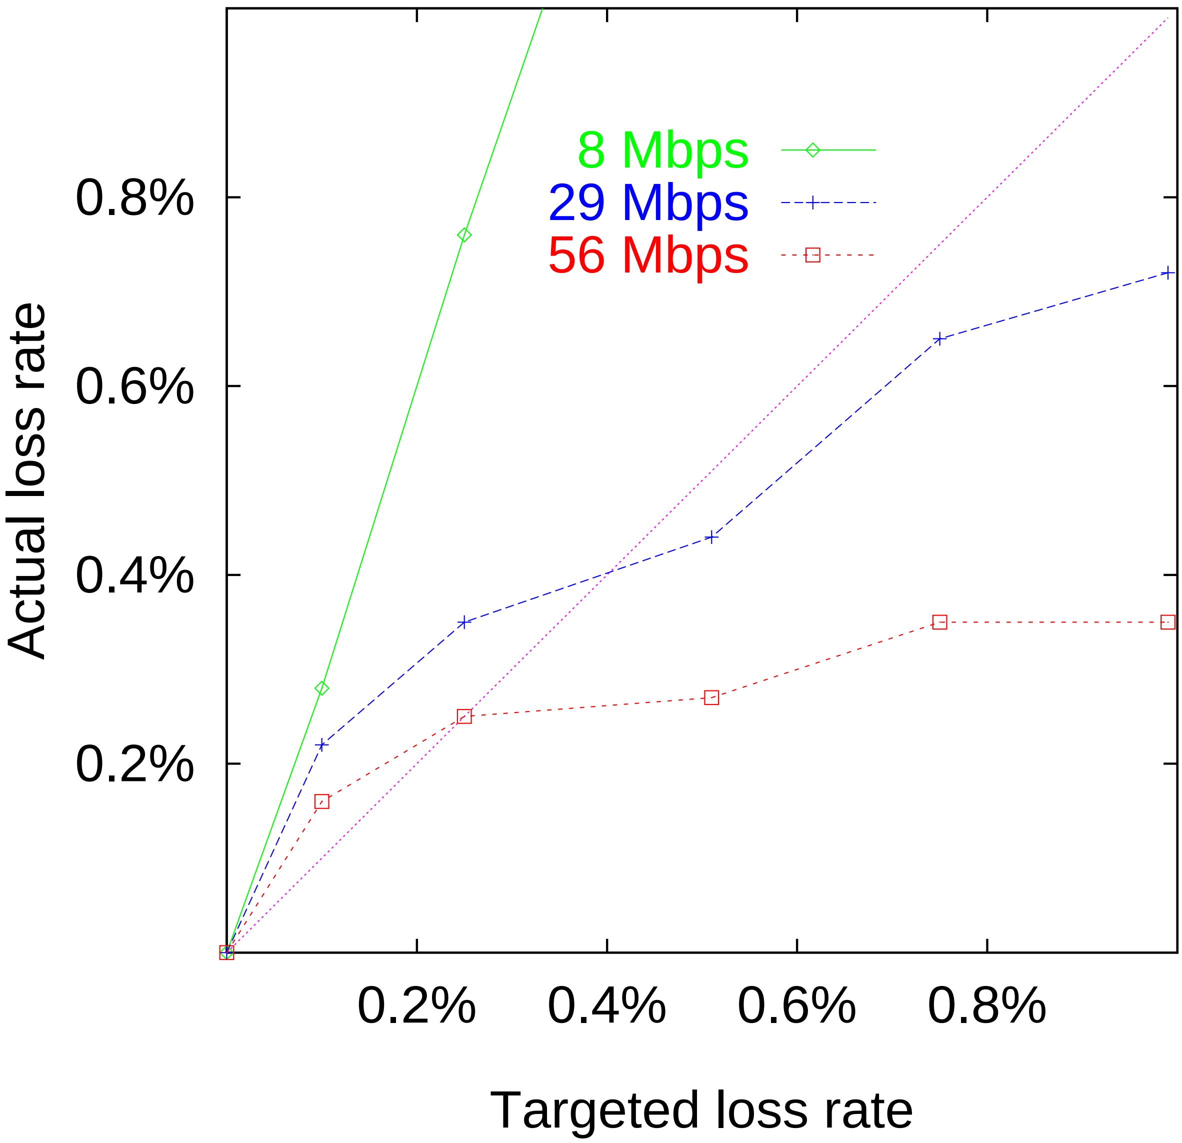

OS Support in End Systems

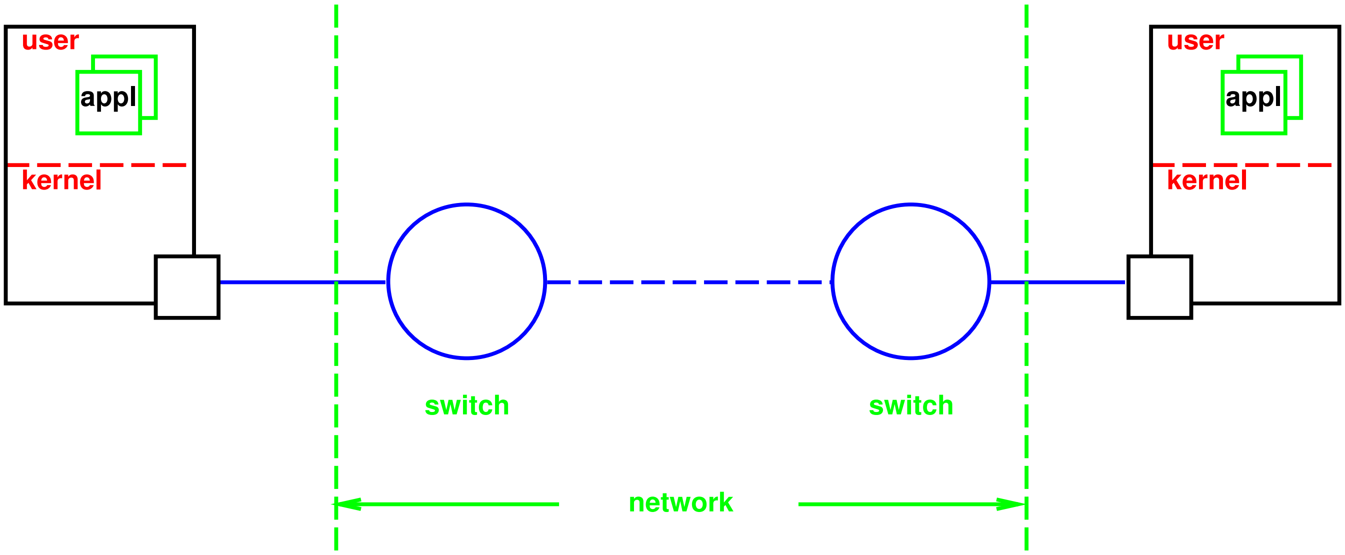

- **Adaptive rate-controlled CPU scheduler**

  [*Yau & Lam*]

- **Efficient data transfer from user space to network interface**

  e.g., zero-copy [*Dittia, Parulkar, & Cox*]