

MaxNet NetLab Presentation

Hailey Lam 08.08.08

A decorative graphic consisting of several horizontal lines of varying lengths and colors (teal, white, and light blue) extending from the right side of the slide towards the center.

Outline

- MaxNet as an alternative to TCP
- Linux implementation of MaxNet
- Demonstration of fairness, quick convergence, etc.
- Incorporating ADPM with MaxNet

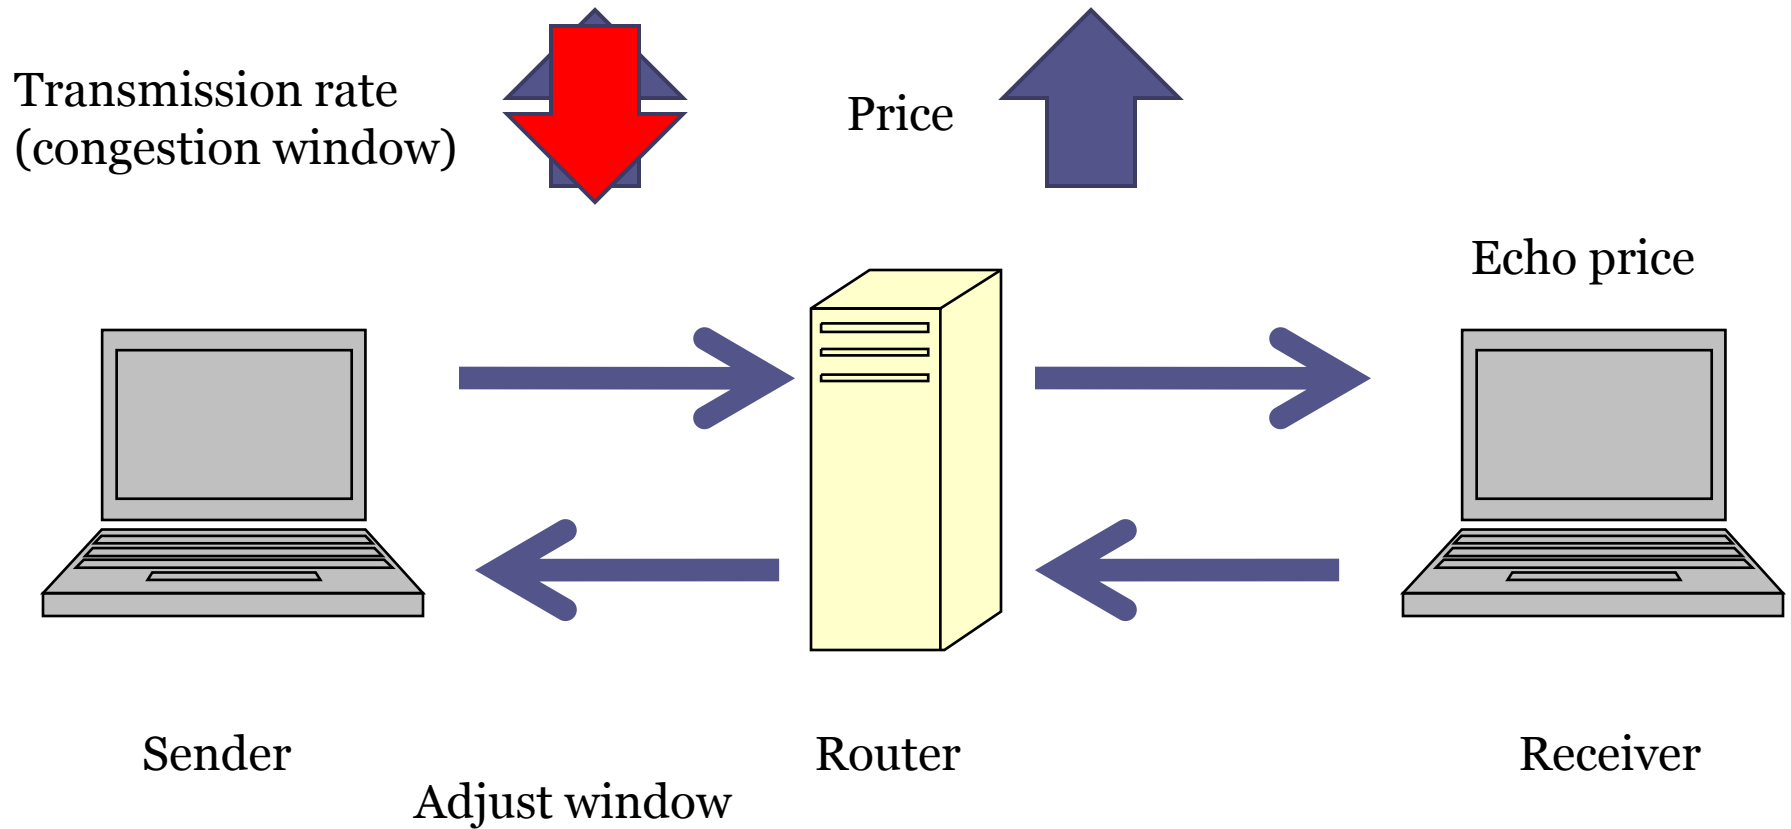
Shortcomings of TCP

Sensitivity to losses

Round-trip time (RTT) unfairness $\frac{rate_2}{rate_1} = \left(\frac{RTT_1}{RTT_2}\right)^2$

Slow start-up times for small flows

MaxNet Overview



MaxNet Benefits from Explicit Feedback

- Avoid queueing, loss instead of reacting to it
 - Avoid hindering flows
- Flows achieve max-min fairness
 - Maximize the minimum transfer rate
 - Without affecting minimum, maximize second minimum rate, repeat.

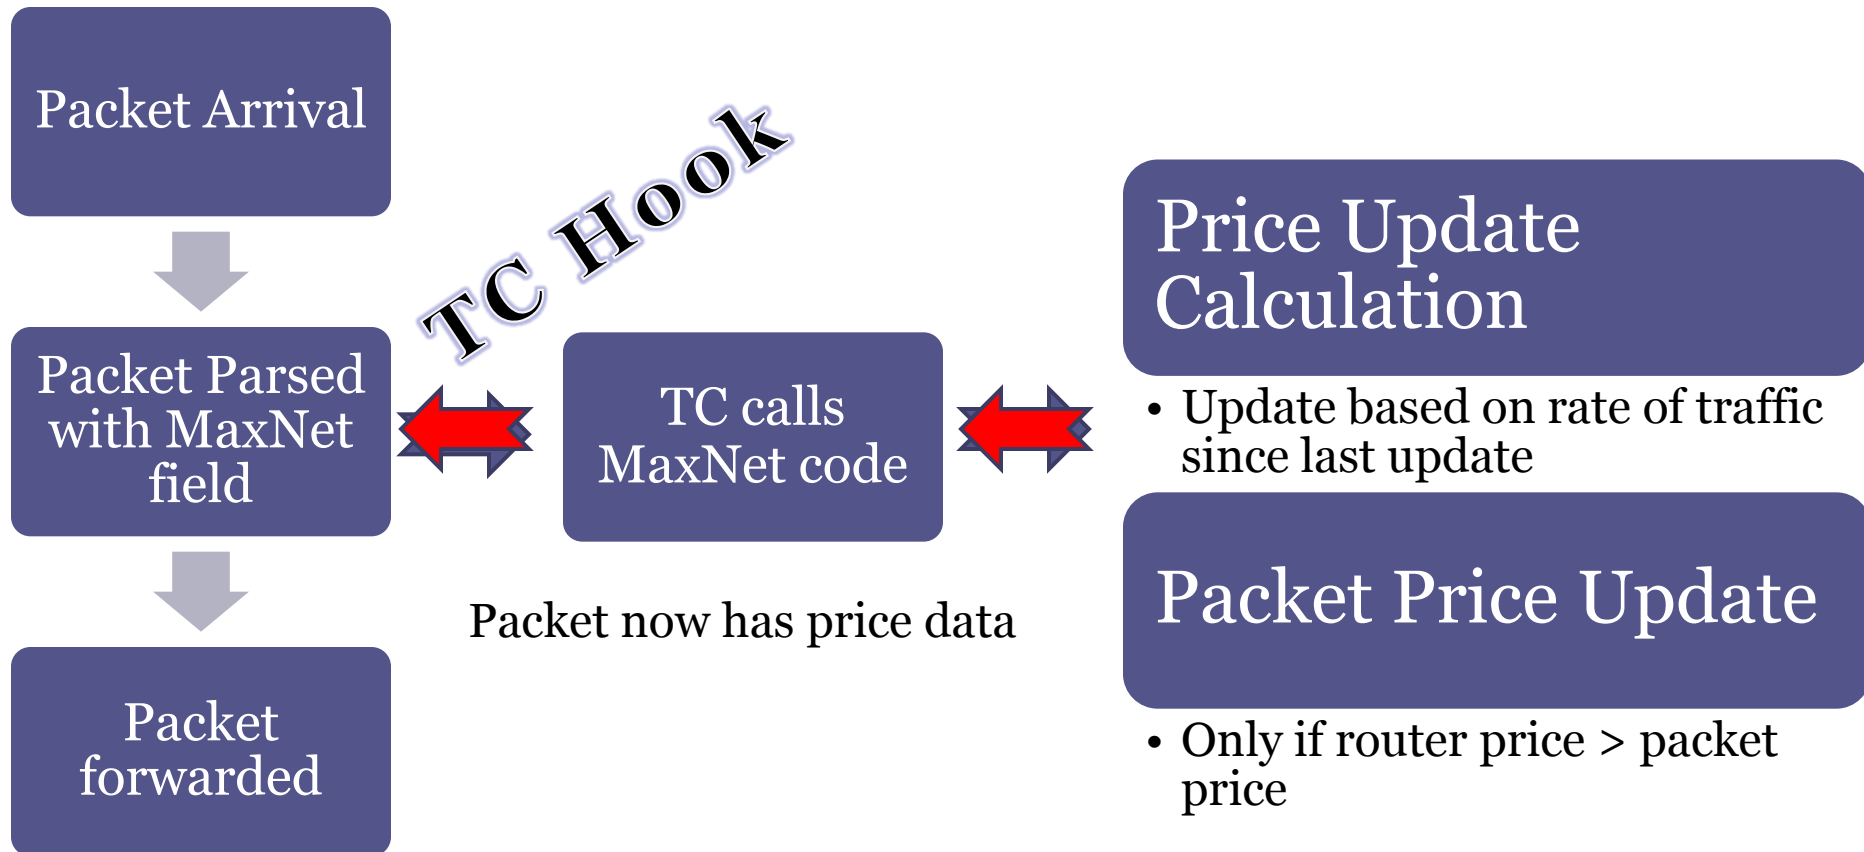
Max-Min Fairness



Other Protocols using Explicit Feedback

- XCP (Katabi, Handley, Rohrs 2002)
 - At capacity, new flow bandwidth allocation too slow
- RCP (Dukkipatti, Kobayashi, Zhan-Shen, Mckeown 2005)
 - Signals RTT for more responsiveness
- REM, JetMax, RED/ECN, many others

TC on Routers to Implement MaxNet



Choice of Host Kernel

- Linux kernel version 2.6.23.1
 - Benefits from improvements in SACK (selective acknowledgements)
- Modular Framework
 - Consolidates much of MaxNet code

Linux Modular Framework

**Linux Kernel Modular
Networking Framework**



Hooks from Modular Framework

Cwnd_event – Extracts MaxNet Option

Cong_avoid – Changes congestion window per algorithm

Ssthresh – Sets threshold for slowstart

Many Others...

TCP Header

TCP
Header

SACK
Header

SACK 1

SACK 2

MaxNet
Option

MaxNet Option Format

MaxNet Option
(42)
1 byte

Option Size (6)
1 byte

Echo Price
(3 bytes)

Price
(3 bytes)

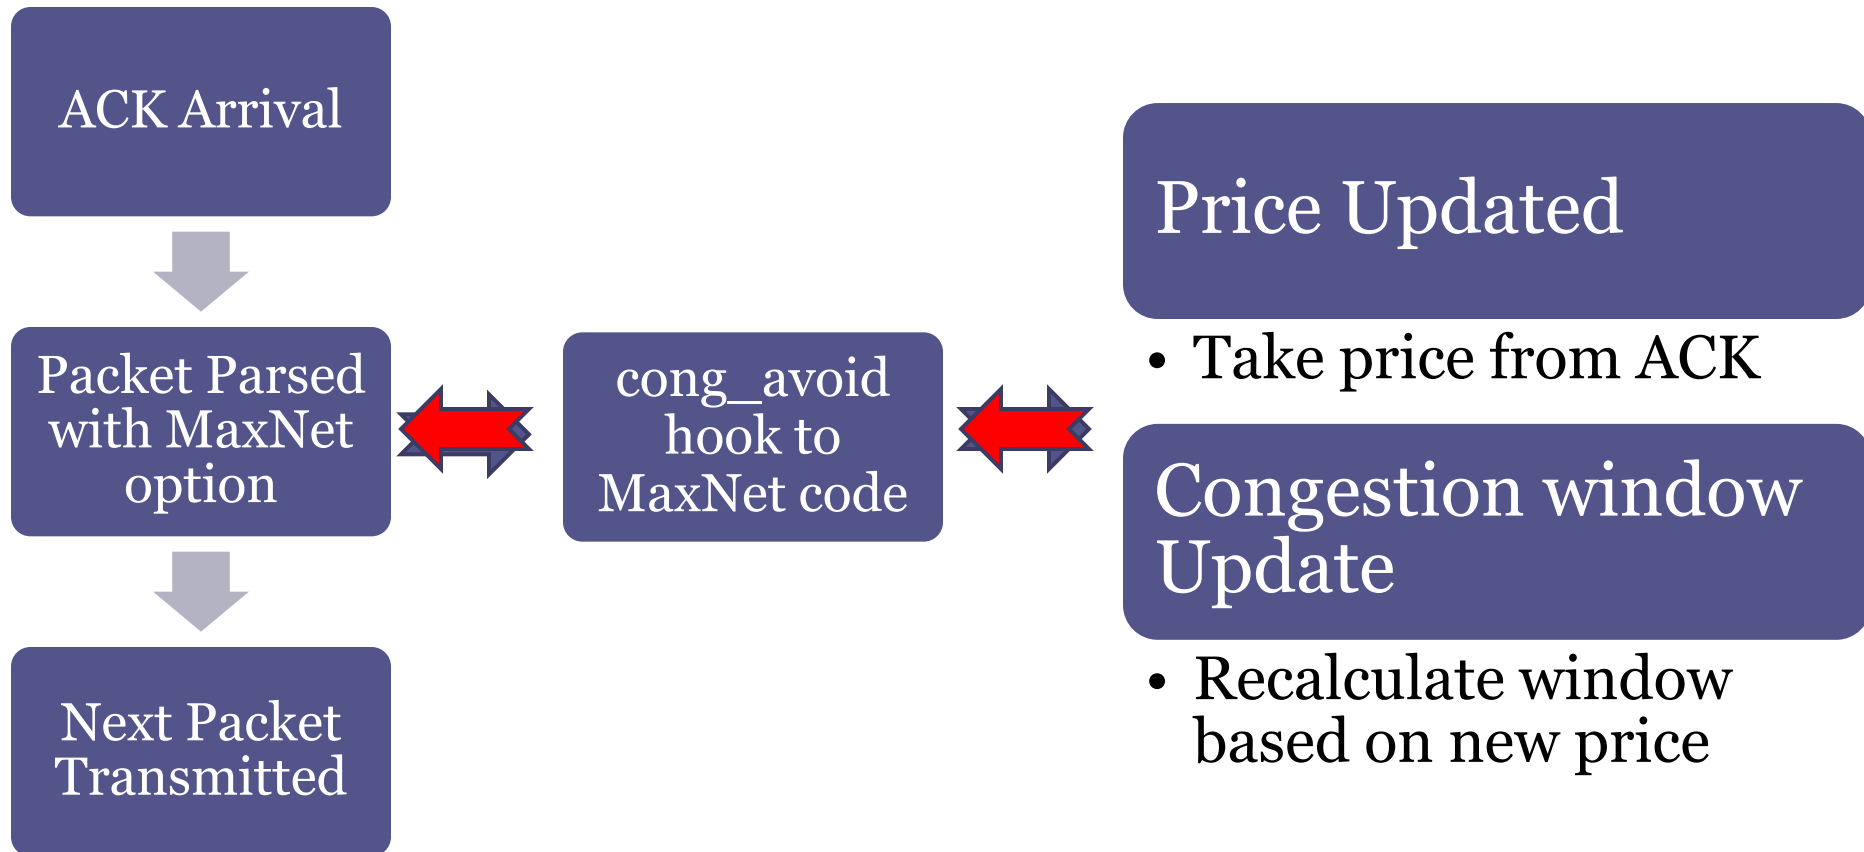
Sender Targets for Equilibrium

- Equilibrium Rate (log scale for constant relative precision)

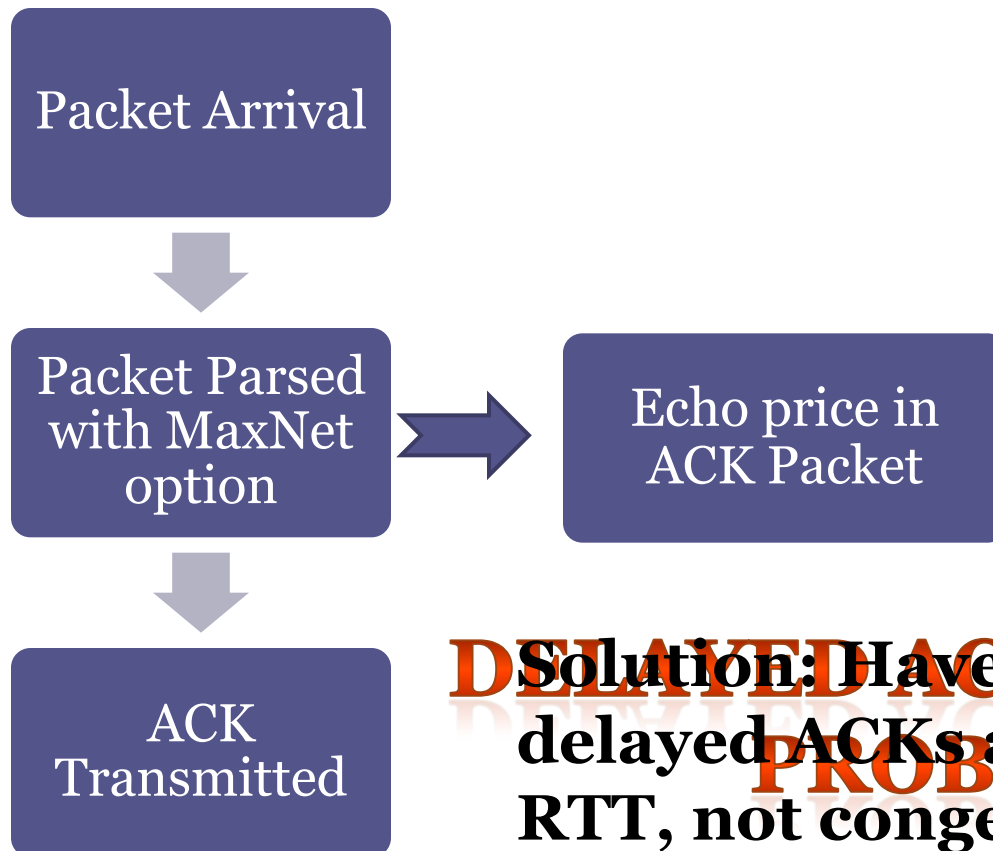
$$x_{\max} e^{-price/T}$$

- Variable ξ for stability
- Host keeps track of variables: ξ , price
 - Values in fixed point format

Sender Acknowledgement Arrival




Receiver Data Arrival




DELATED ACKS CAUSE A PROBLEM!
Solution: Have MaxNet treat delayed ACKs as increase in RTT, not congestion

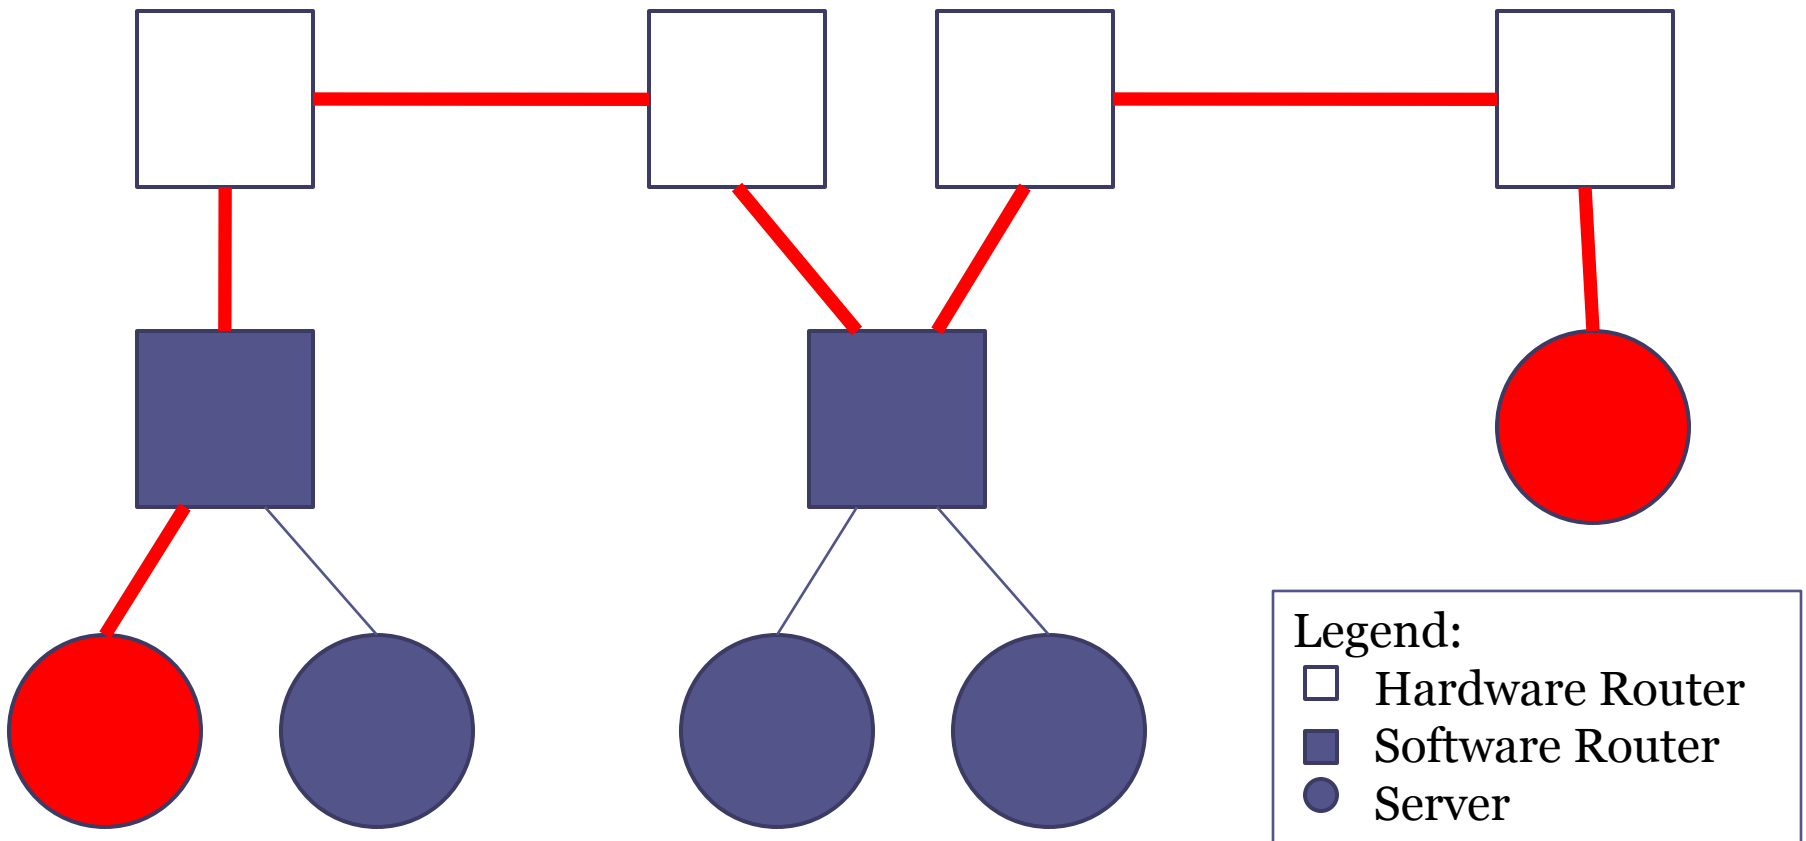
Experiment Topology




Experiment Topology




Experiment Topology



Experiment Topology



Example Data - Fairness and Quick Convergence



Bursty Data Flows

- Data transmitted/received in bursts
- Router price and CPU load increase temporarily
- Results in slower transfer rates
- TCP Pacing as a solution

Future enhancements

- “Incompatibility” with TCP
 - Greedy TCP will hog bandwidth
 - MaxNet backs off and stops sending
- Encrypted packets
 - MaxNet option in TCP header – routers cannot use

Advanced Deterministic Packet Marking (ADPM)

- Single bit (ECN) in IP header marks packets
- Determine bit value with IPid field hash
- More packets passed yield closer estimate to actual price

ADPM

- **Benefits**
 - Enables MaxNet to work with encrypted packets
 - Routers should not look inside TCP layer
- **Drawbacks**
 - Decreases accuracy of congestion level estimation

Summary

- MaxNet: Router assisted scheme
 - Avoids some pitfalls of TCP
- MaxNet ported to 2.6.23.1
 - Modular framework
 - Improved SACK handling
- Address packet encryption with ADPM