Bandwidth Allocation

H. T. Kung
Division of Engineering and Applied Sciences
Harvard University

November 30, 1998

Copyright 1998 by H. T. Kung

Presentation Plan

Given a link (or virtual circuit) between two nodes, consider the following two problems:

1. Allocate a certain bandwidth for the link

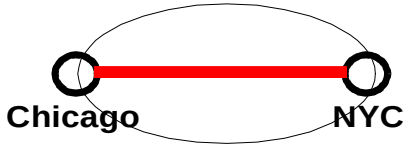

2. Allocate the bandwidths that various traffic classes over the link can consume

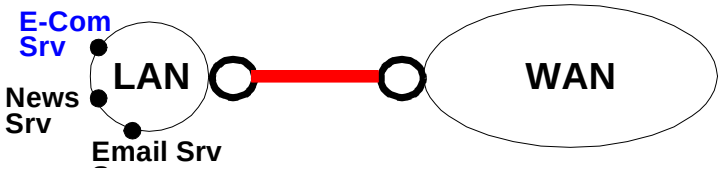

Problem 1: Allocate Link Bandwidth over a Network

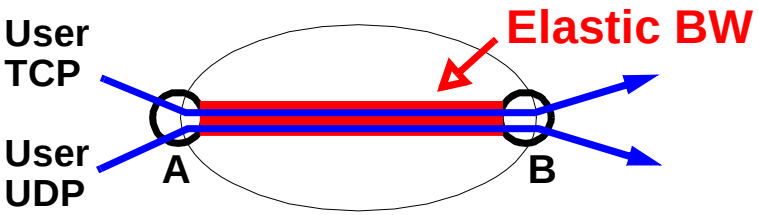

Want "Elastic" services between two communicating edge nodes A and B:

- Allow bandwidth overflow when extra network bandwidths become available
- Can guarantee minimum bandwidth and bound peak bandwidth

Illustration of Bandwidth Elasticity

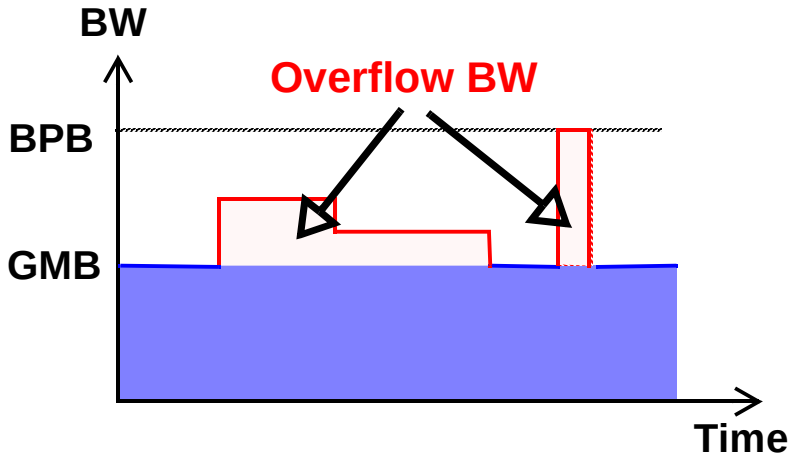

BPB: Bounded Peak Bandwidth

GMB: Guaranteed Minimum Bandwidth

Good Elastic Services

1. Must be easily understood by customers
2. Must be implementable by service providers
3. Must be able to achieve high degrees of network utilization and fairness
4. Must be able to handle both TCP and UDP traffic load

Five Dimensions of Elastic Services

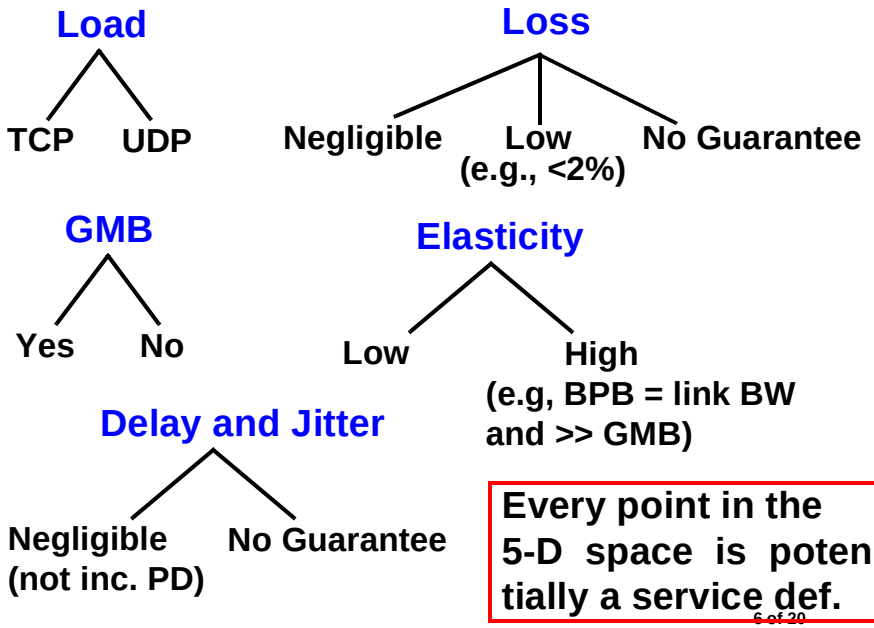

Every point in the 5-D space is potentially a service def.

Application Examples and Their Desired Services /1

Large File Transfer

- TCP, low loss
- TCP, GMB, high elasticity, low loss

Aggregate of Interactive Web Sessions

- TCP, GMB, low elasticity, low loss

Real-Time Control

- UDP, negligible loss, jitter and delay

NFS

- UDP, low loss
- UDP, low loss, GMB

Application Examples and Their Desired Services /2

Interactive Voice

- UDP, GMB, low elasticity, negligible loss and delay
- UDP, GMB, low elasticity, low loss

Video streaming

- UDP, GMB, low elasticity, low loss
- TCP, GMB, low elasticity, low loss

Some Solution Approaches

1. **Diff-Serv-PS-GMB: Diff-Serv Premium Services**
   - Achieving GMB via admission
2. **Diff-Serv-AS-GMB: DiffServ Assured Services, with GMB**
   - Add GMB to DiffServ AS via admission and some sophisticated RIO-like buffer management
3. **TCP-Trunking-GMB**
   - Add GMB to TCP trunking (see later slides)

Service-Solution Association

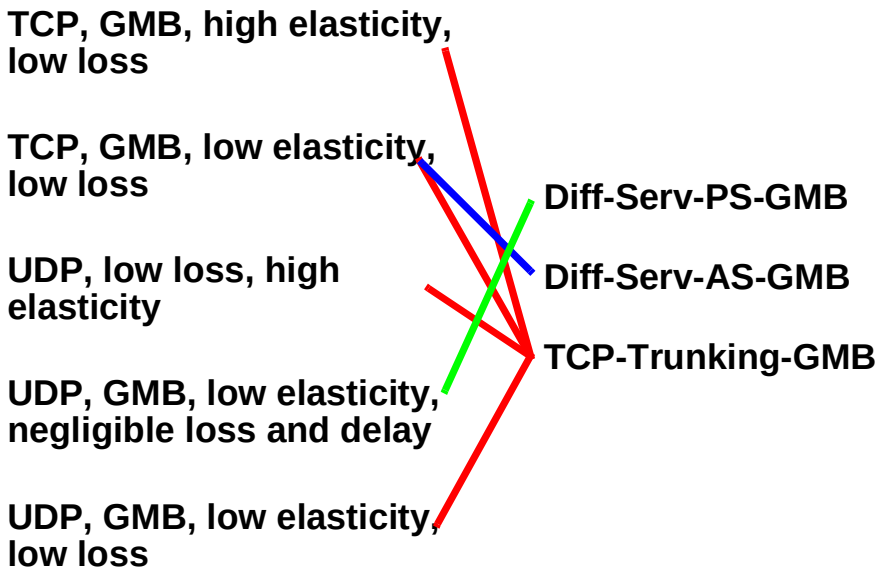

TCP-Trunking-GMB (HT Kung and ShieYuan Wang)

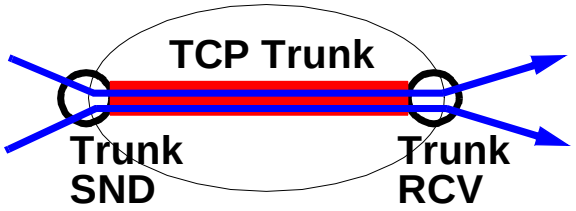

Trunk SND:

- S1: Send pkts *unconditionally* at the rate of GMB
- S2: Send additional pkts under TCP congestion control

Trunk RCV:

- Forward received pkts immediately

Illustration of TCP-Trunking-GMB

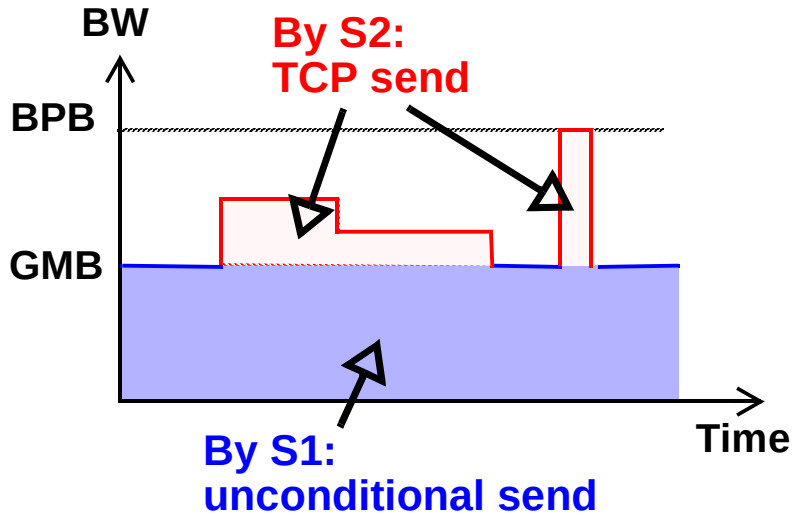

Working of TCP-Trunking-GMB (Validated on Experimental Networks)

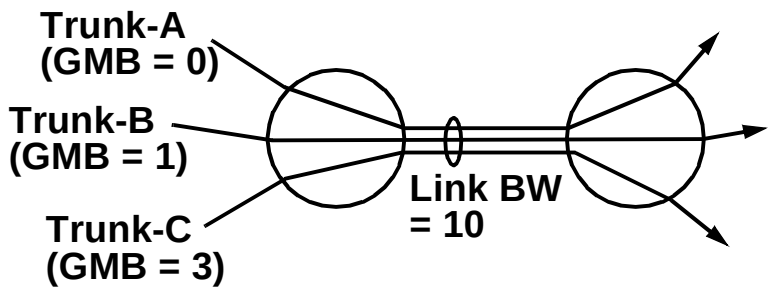

Achieved Bandwidths:

Trunk A: 0 + 2 = 2

Trunk B: 1 + 2 = 3

Trunk C: 3 + 2 = 5

Properties of TCP-Trunking-GMB

- Trunks can each achieve their GMBs and share the remaining bandwidth fairly
- A trunk will recover lost pkts sent under TCP by S2, but will not recover those sent unconditionally by S1. (This is OK.)

Problem 2: Allocate Bandwidths for Traffic Classes over a Link

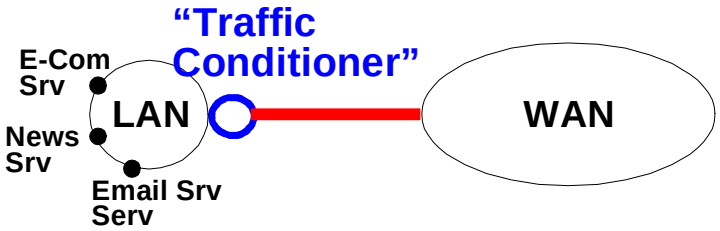

- Classify and monitor flows
- Provide differentiated treatments to packets based on their classes

Experimental system: Nortel's Traffic Conditioner (1997)

Products from companies such as Packeteer, Check Point, and Ukiah

Packeteer Approach

- Classify traffic based on IP address, port/service, URL, TOS, etc.
- Shape traffic by making modifications to TCP ACK packets

Packeteer's TCP Shaper /1

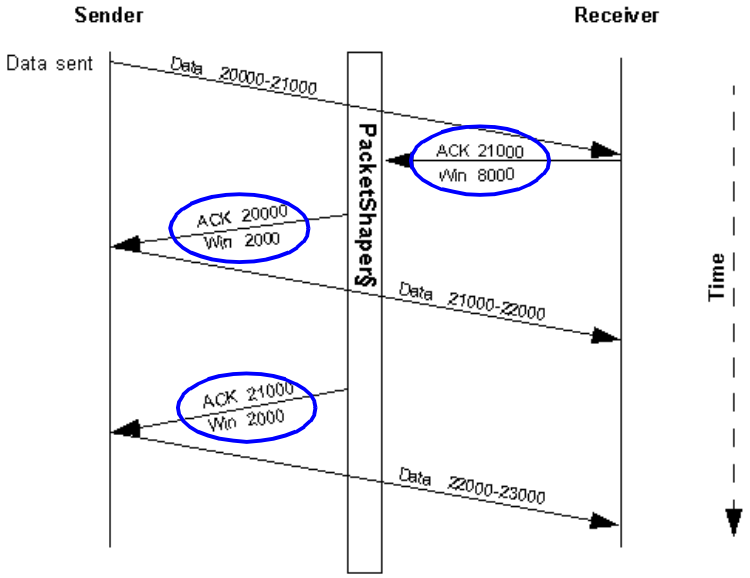

Packeteer's TCP Shape /2

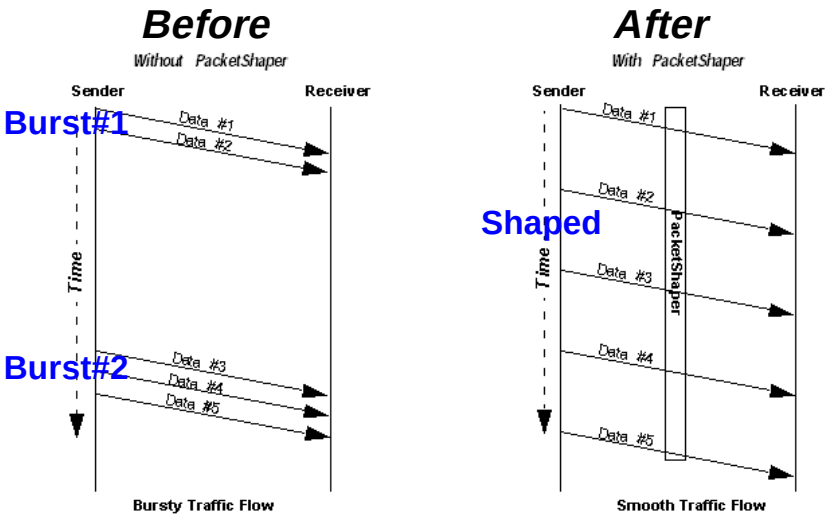

Concluding Remarks

- **Traditional bandwidth allocation is fixed-bandwidth allocation such as use of dedicated circuit services**
- **New bandwidth allocation methods will provide elastic bandwidth, and use of the bandwidth will be class-based. Research is still needed before these new services become widespread**

References

[1] Kung, H. T. and Wang, S. Y., TCP Trunking, July 1998. http://www.eecs.harvard.edu/~htk/paper/tcp-trunk-798.pdf

[2] B. Nandy, N. Seddigh, A.S.J. Chapman and J. Hadi Salim, A Connectionless Approach to Providing QoS in IP Networks, 1998

[3] Chapman, A. and Kung, H. T., Automatic Quality of Service in IP Networks, Proc. the Canadian Conference on Broadband Research, April 1997, pp. 184-189