AMES
A Framework for Fair Bandwidth Sharing

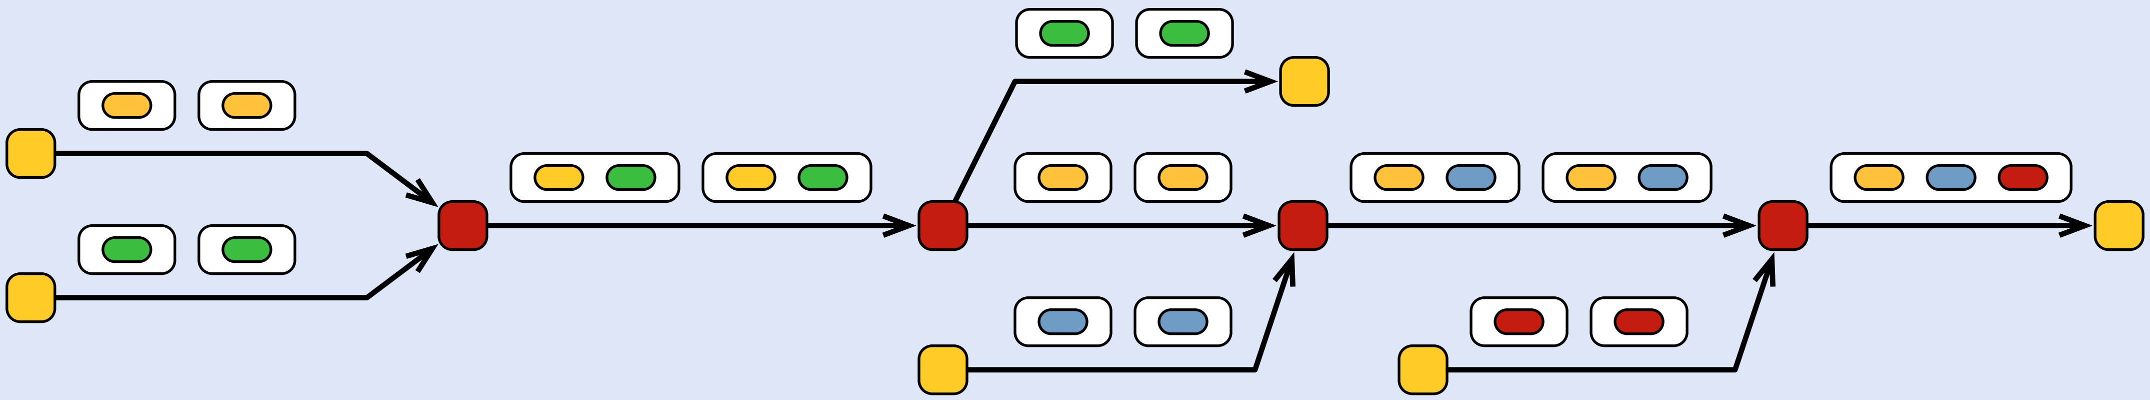

Srinivasan Ramasubramanian, Sathya Parthasarathy
University of Arizona

Arun K. Somani
Iowa State University

Motivation

Fair sharing is well-researched, why another framework?

Congestion control

Packet dropping, bandwidth estimation, rate adaptation

Packet dropping is not favorable

At high speeds as cost of retransmission increases with bandwidth-delay product

Needs to be carefully engineered to not affect traffic that adapt rates (TCP vs. UDP)

Packets and flows

Packet scheduling layer is not aware of flows

Fair sharing is defined at the level of flows

What we seek

A framework that achieves fair sharing

- Low scheduling complexity
- Less buffering requirements
- Does not perform flow-specific operations
- Does not employ packet dropping as primary method for reacting to congestion

A cross-layer approach?

- A co-ordinated approach
- A method to identify the number of flows crossing a link
- A method to identify the available bandwidth on a path
- Rate adaptation at source based on estimated bandwidth

AMES architecture

All routers employ

- Link-specific queuing
- Round-robin scheduling
- No packet dropping

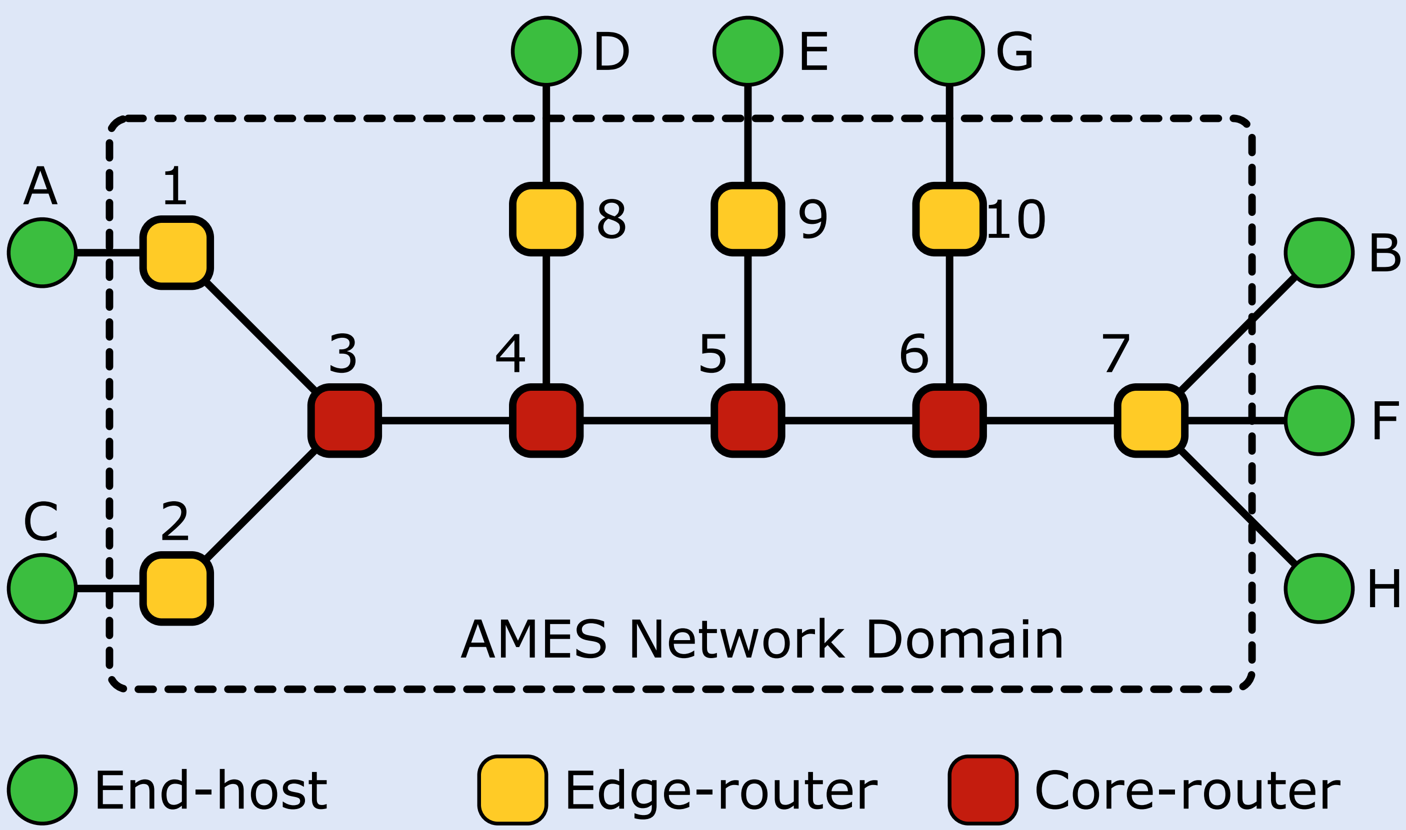

Transmission rounds

Packets are transmitted in rounds

We employ ODD and EVEN rounds - require one bit in packet header

Scheduling at a node

Select a round of packets from every eligible queue

Merge them together and transmit

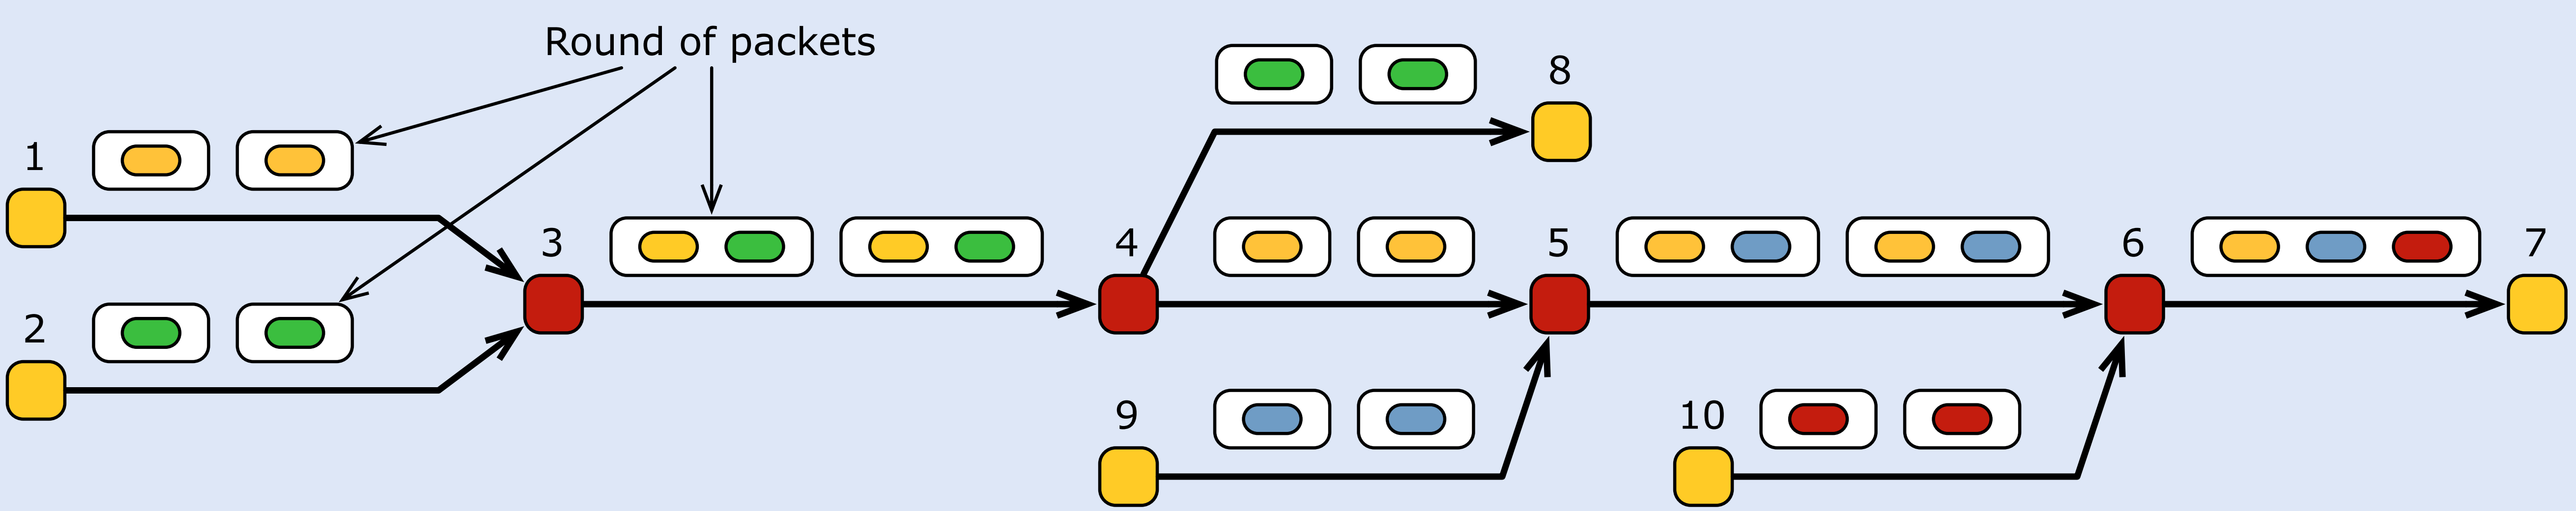

Buffering rounds

Packets are buffered in rounds

We employ ODD and EVEN rounds

Buffering rounds are independent from transmission rounds

Two packets that are transmitted in different rounds must be separated

AMES scheduling

Service provided to a queue

- Minimum packet size selected for transmission in that from from the queue
- A queue is credited with this amount when selected for transmission

Minimum service provided in a round

- Minimum packet size among all queues selected for transmission in that round
- Once a round of transmission is considered, this quantity is decremented from all queues

Similar to Deficit Round Robin (DRR), but

- The next queue selection in AMES is not performed on a per-packet basis
- The next packet for transmission comes either from the same queue or the next queue
- AMES scheduling complexity is O(1)

An example with different packet sizes

All flows are assumed to have the same priority

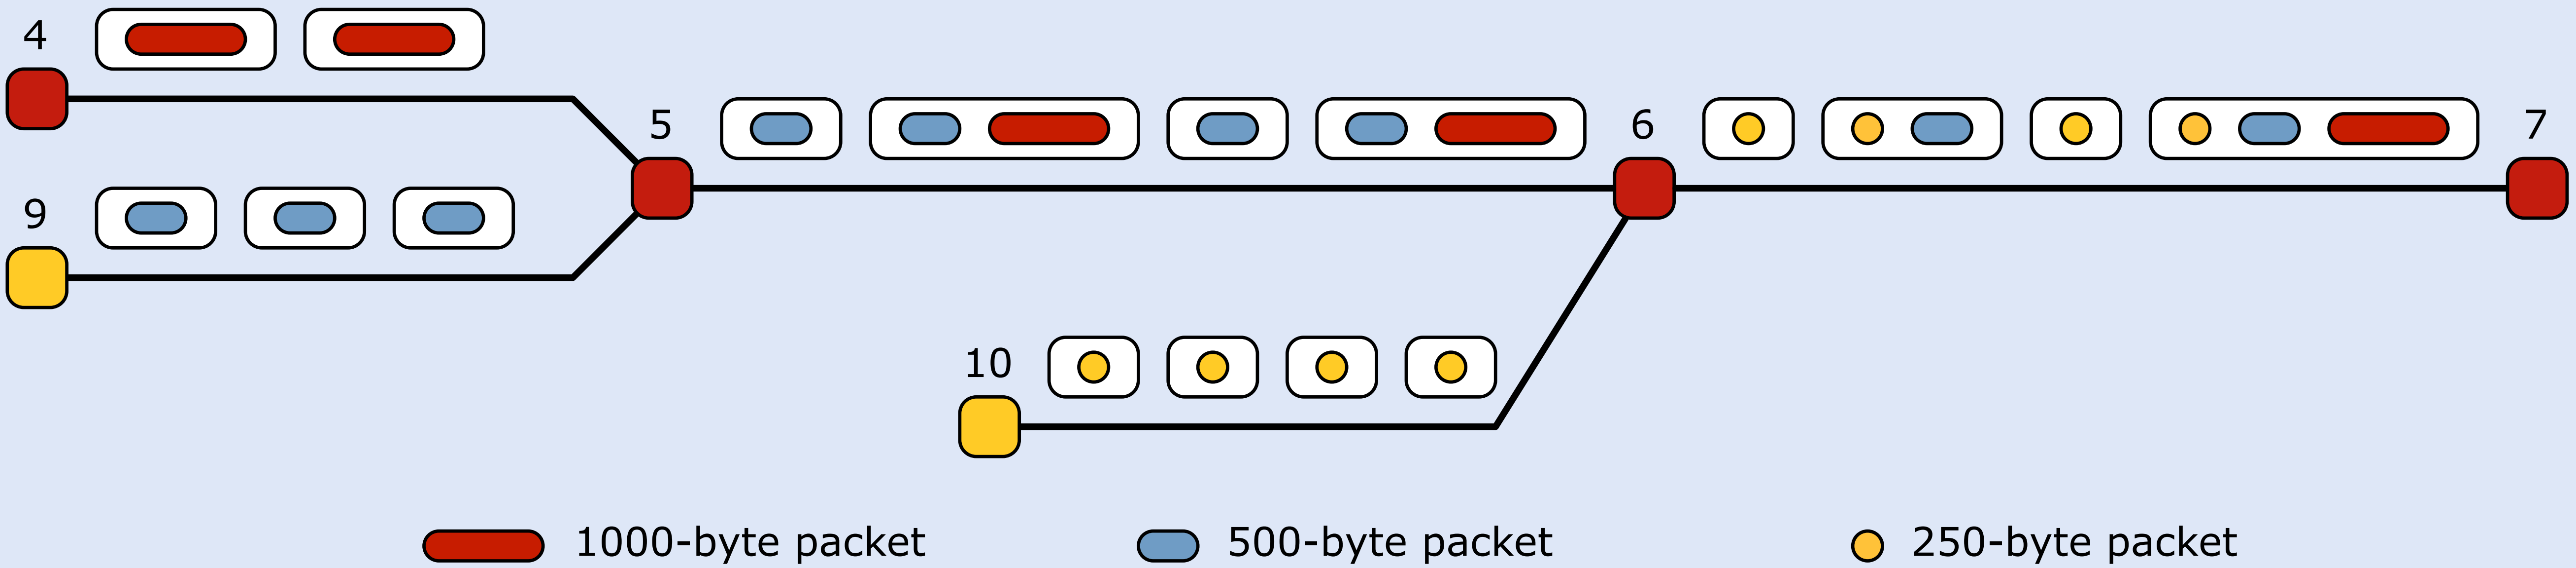

An example with different priorities

Different priorities for flows

4 - 7 flow has priority 4

9 - 7 flow has priority 2

10 - 7 flow has priority 1

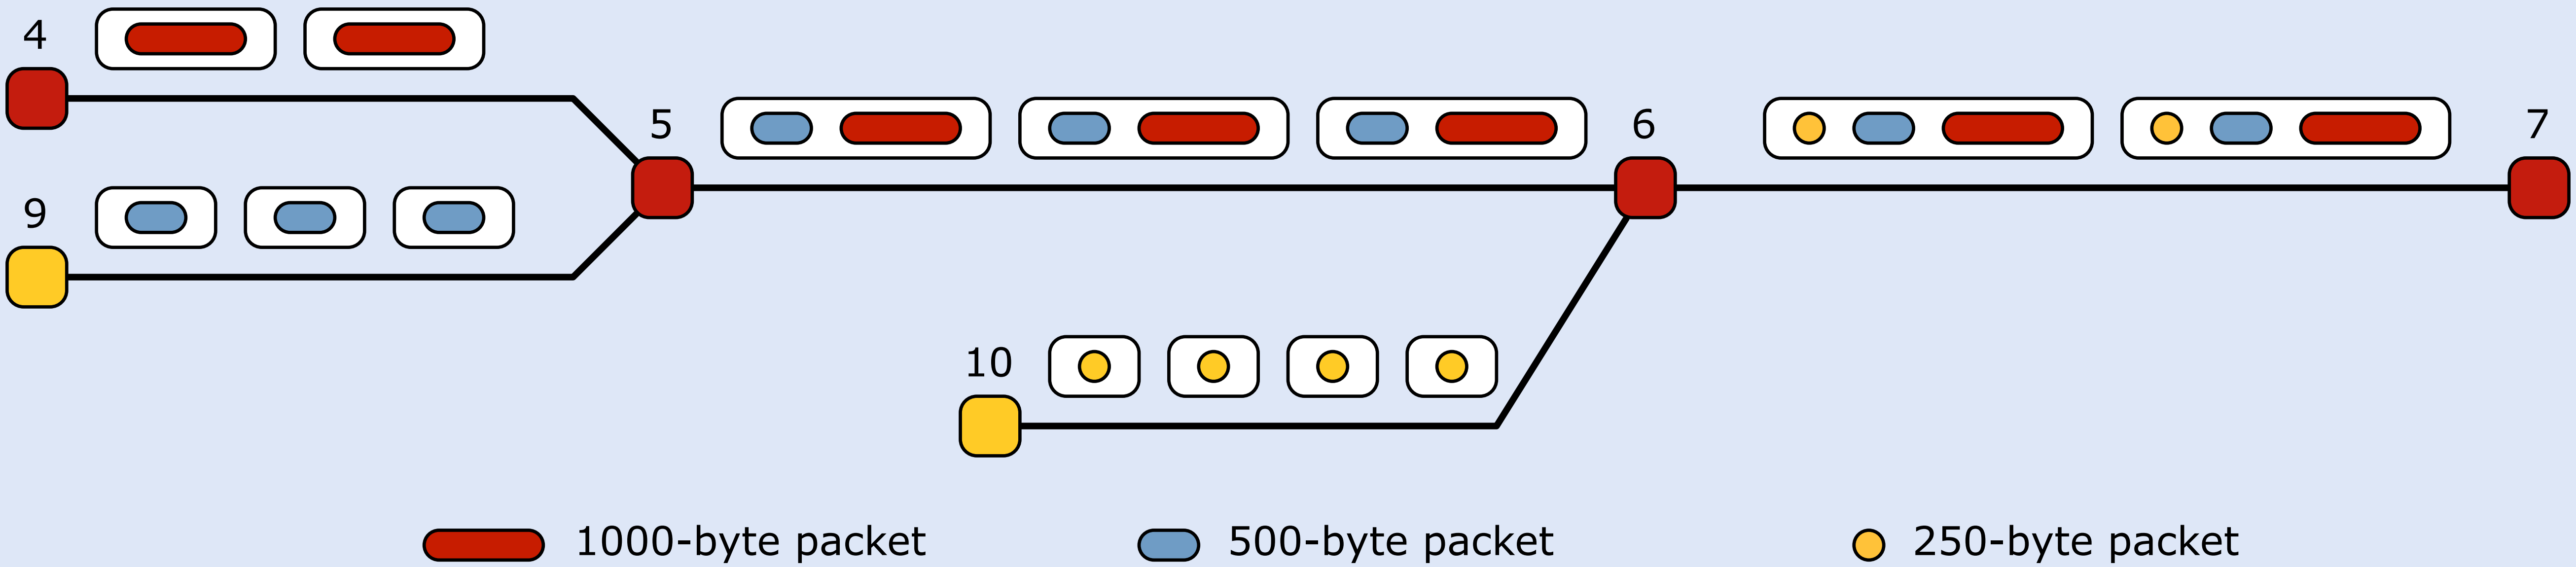

No packet dropping

Routers have buffer thresholds

When threshold is met, a 1-bit notification is sent to neighboring router

- Neighboring router stops transmission, until the notification is relased

Enough buffer space over the threshold is provided at routers

- Size depends on propagation delay and transmission speed
- No packets are dropped

Performance evaluation

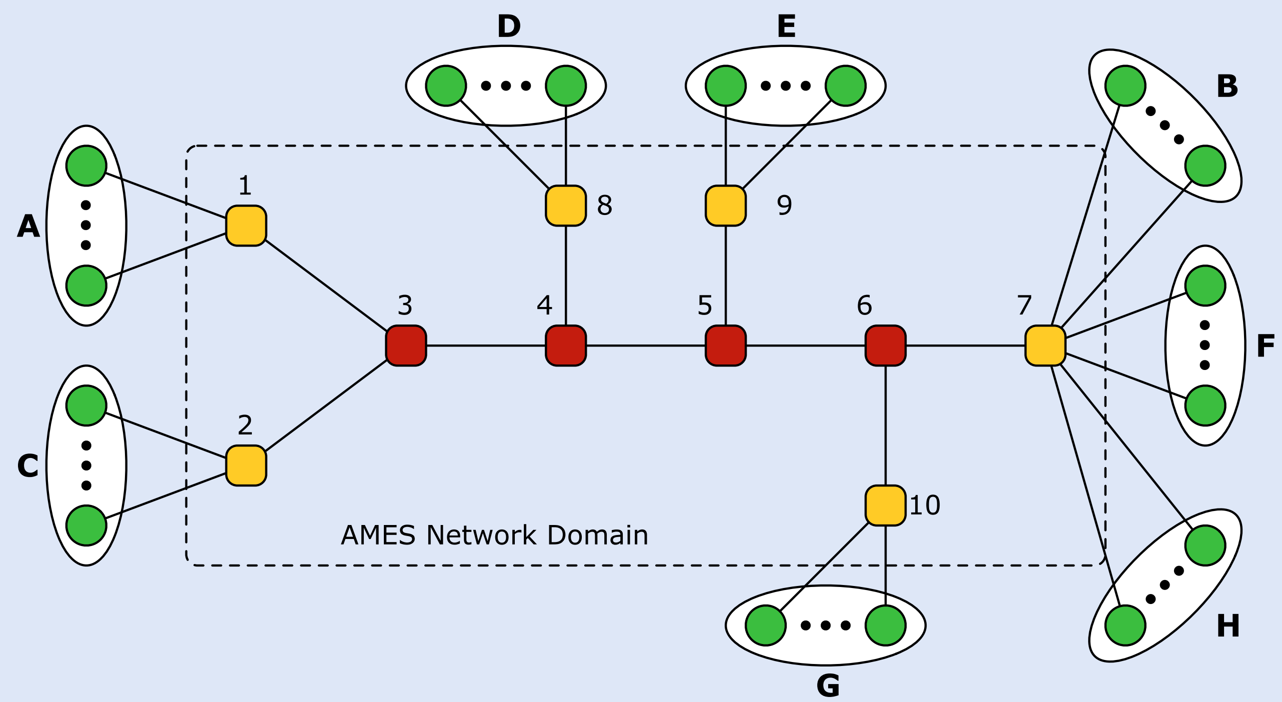

Traffic flows from A-B, C-D, E-F, G-H

Each group has multiple traffic streams (TCP/UDP)

We measure the average bandwidth received by a flow in a group

10 Mbps link, with buffer threshold of 15 KBytes

Compared to CSFQ (Core Stateless Fair Queuing)

S. Ramasubramanian, S. Parthasarathy, A. K. Somani., "AMES: A framework for fair bandwidth sharing"

Performance evaluation: Scenario 1a

TRAFFIC FOR SCENARIO 1.

Flow group	Number of flows	Flow type	Packet length distribution	Mean packet length (in bytes)
$A - B$	4	UDP/TCP-Reno*	Fixed	1000
$C - D$	8	TCP-Reno	Fixed	2000
$E - F$	4	UDP	Fixed	500
$G - H$	4	UDP	Fixed	250

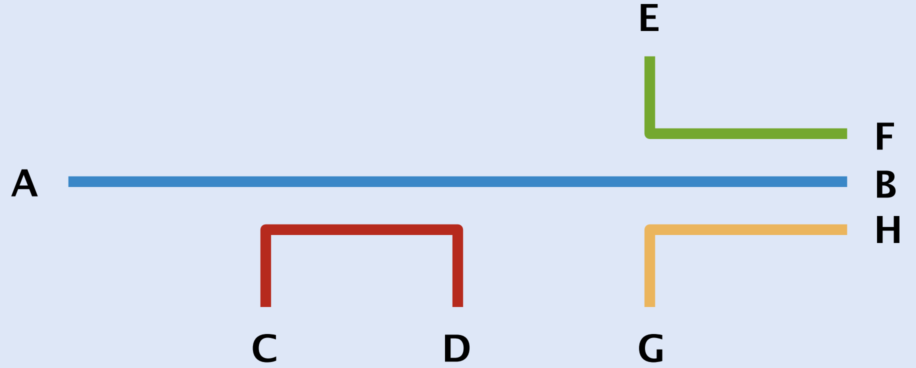

A-B traffic set to UDP

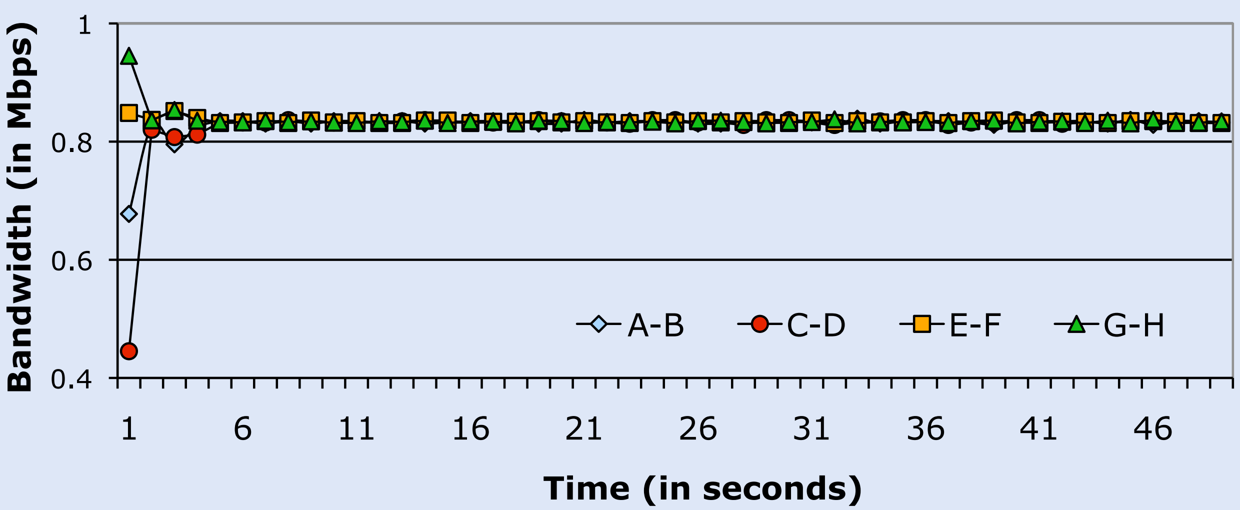

AMES

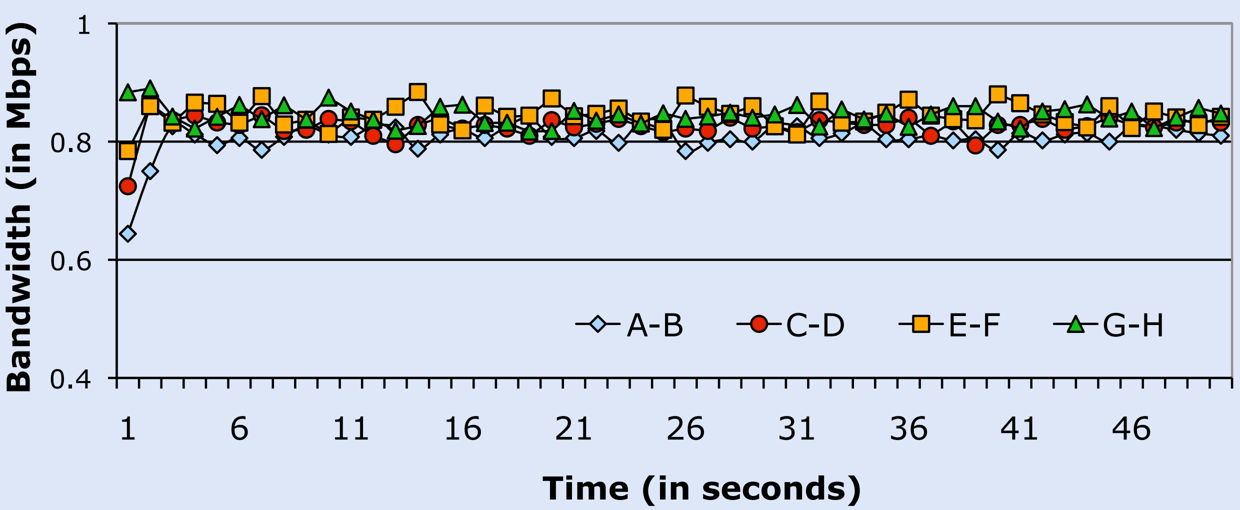

CSFQ

Performance evaluation: Scenario 1b

TRAFFIC FOR SCENARIO 1.

Flow group	Number of flows	Flow type	Packet length distribution	Mean packet length (in bytes)
$A - B$	4	UDP/TCP-Reno*	Fixed	1000
$C - D$	8	TCP-Reno	Fixed	2000
$E - F$	4	UDP	Fixed	500
$G - H$	4	UDP	Fixed	250

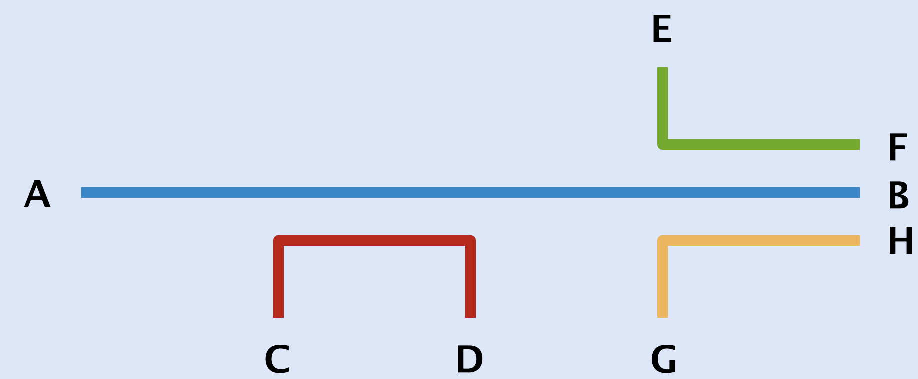

A-B traffic set to TCP

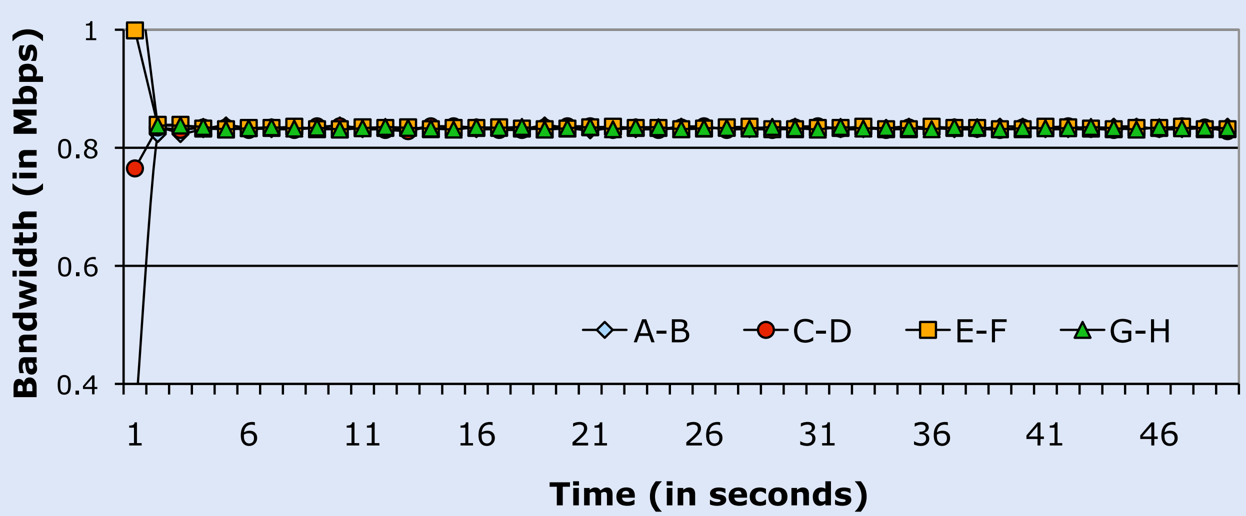

AMES

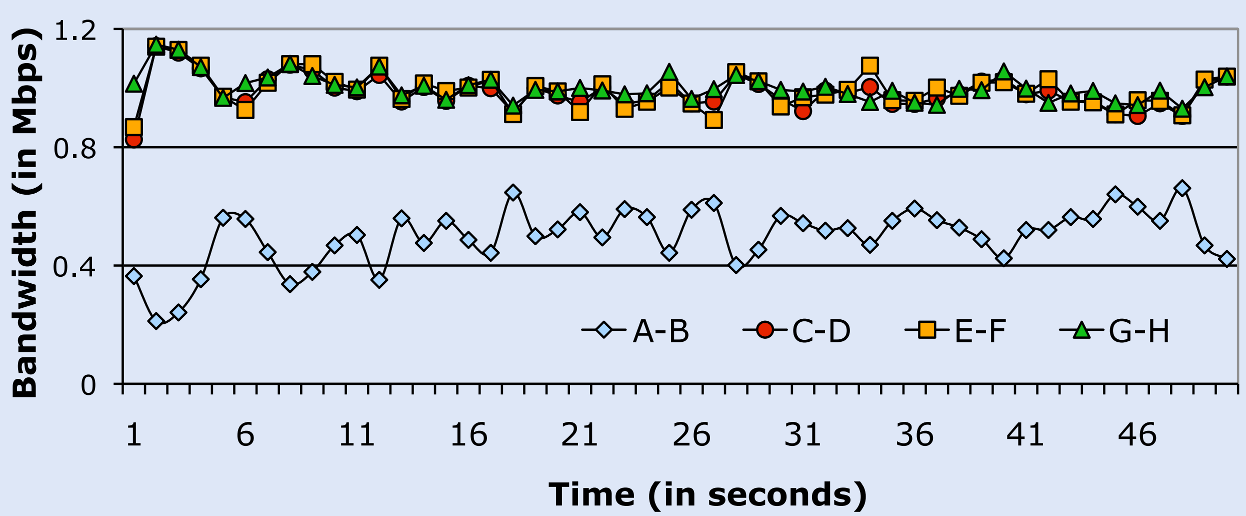

CSFQ

Performance evaluation: Scenario 2a

TRAFFIC FOR SCENARIO 2.

Flow group	Number of flows	Flow type	Packet length distribution	Mean packet length (in bytes)	Weight
$A-B$	4	TCP-Reno	Fixed	1000	2
$C-D$	8	TCP-Reno	Fixed	2000	1
$E-F$	4	UDP/TCP-Reno*	Fixed	500	1
$G-H$	4	UDP/TCP-Reno*	Fixed	250	1

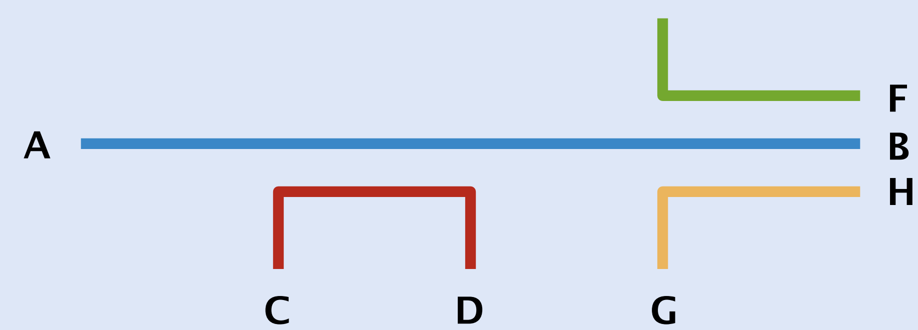

E-F and G-H traffic set to UDP

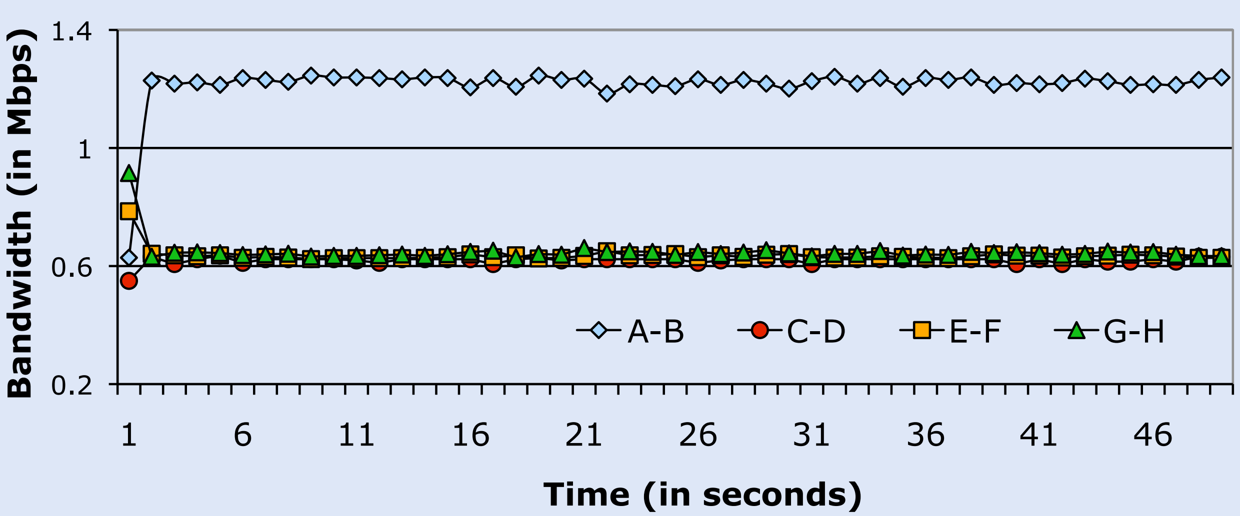

AMES

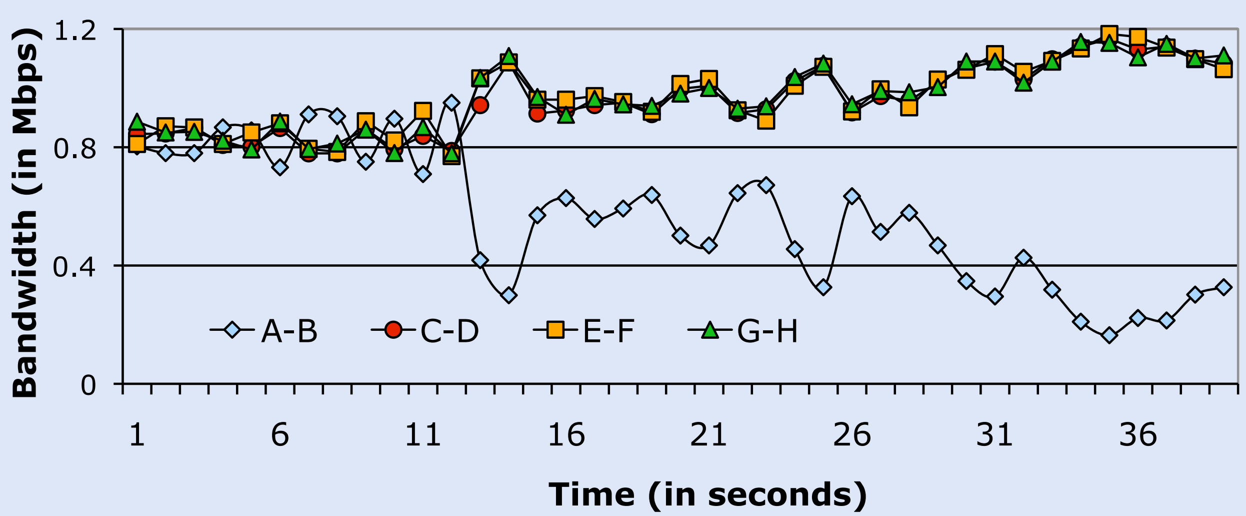

CSFQ

Performance evaluation: Scenario 2b

TRAFFIC FOR SCENARIO 2.

Flow group	Number of flows	Flow type	Packet length distribution	Mean packet length (in bytes)	Weight
$A - B$	4	TCP-Reno	Fixed	1000	2
$C - D$	8	TCP-Reno	Fixed	2000	1
$E - F$	4	UDP/TCP-Reno*	Fixed	500	1
$G - H$	4	UDP/TCP-Reno*	Fixed	250	1

E-F and G-H traffic set to UDP

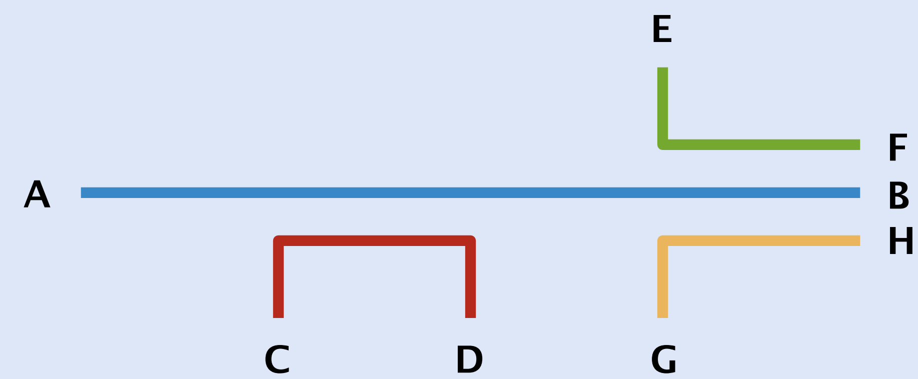

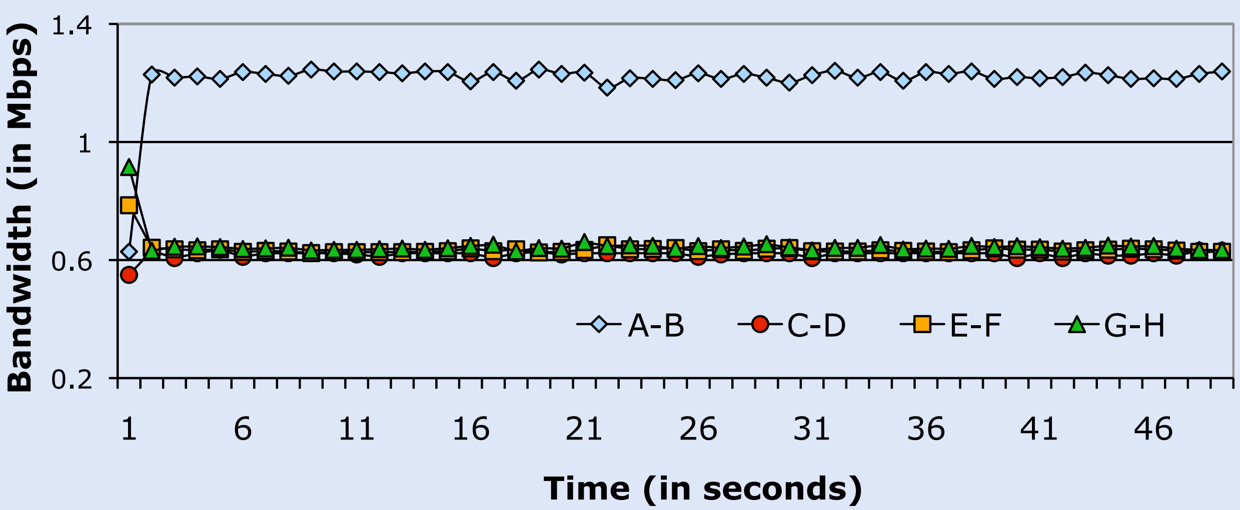

AMES

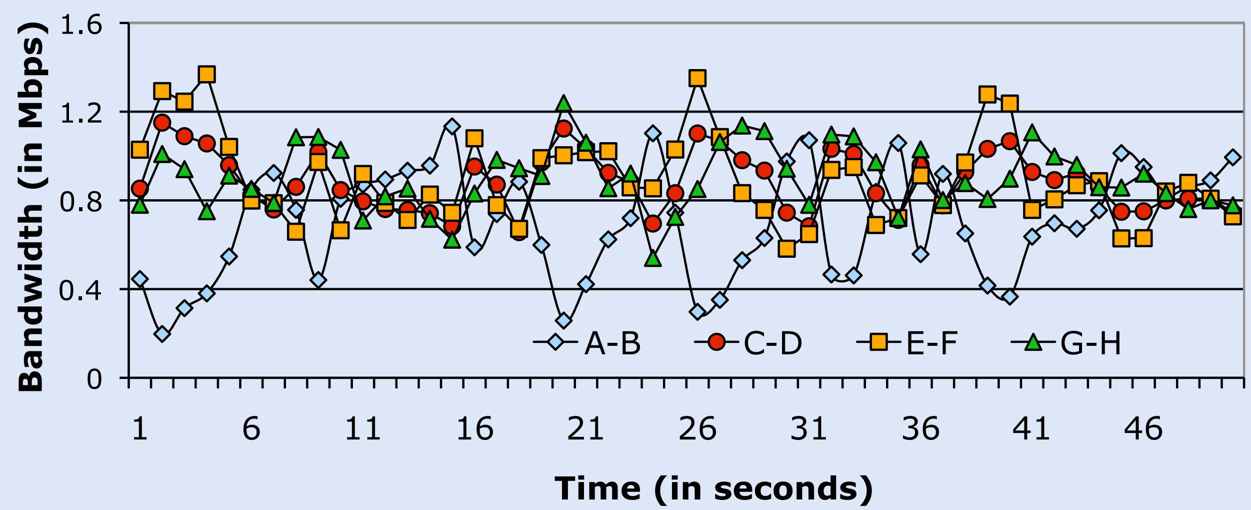

CSFQ

Performance evaluation: Scenario 3

TRAFFIC FOR SCENARIO 3.

Flow group	Number of flows	Flow type	Packet length distribution	Mean packet length (in bytes)	Weight
$A-B$	4	TCP-Reno	Fixed	1000	1
$C-D$	4	TCP-Reno	Fixed	2000	1
$E-F$	4	UDP	Fixed	500	1
$G-H$	4	UDP	Fixed	250	1

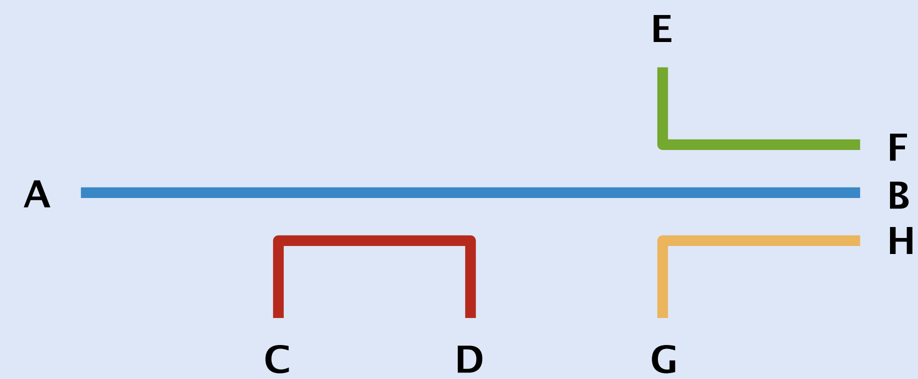

AMES: C-D traffic cannot gain more bandwidth

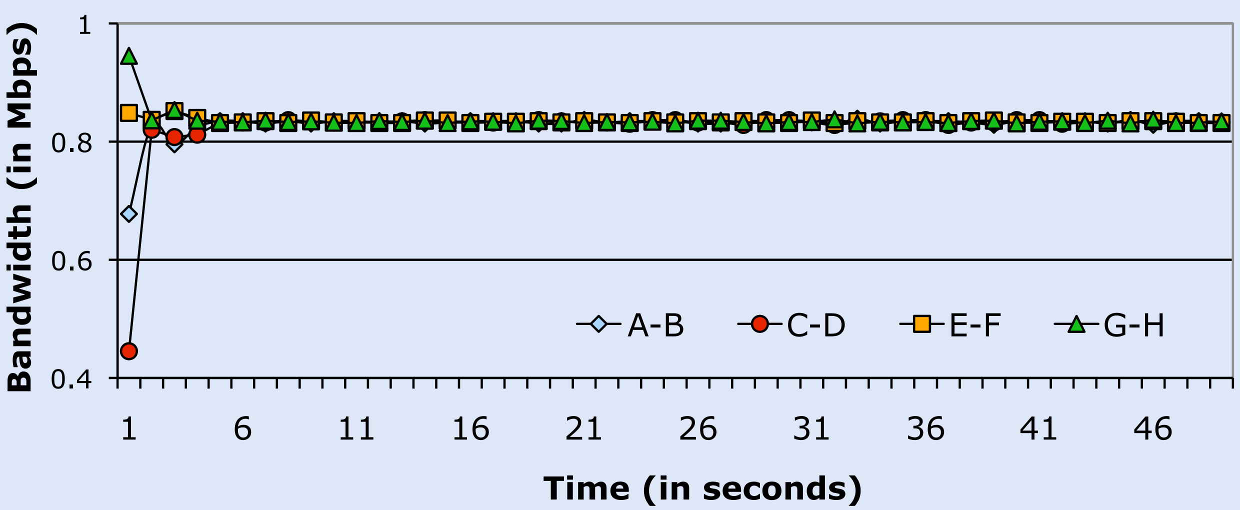

AMES

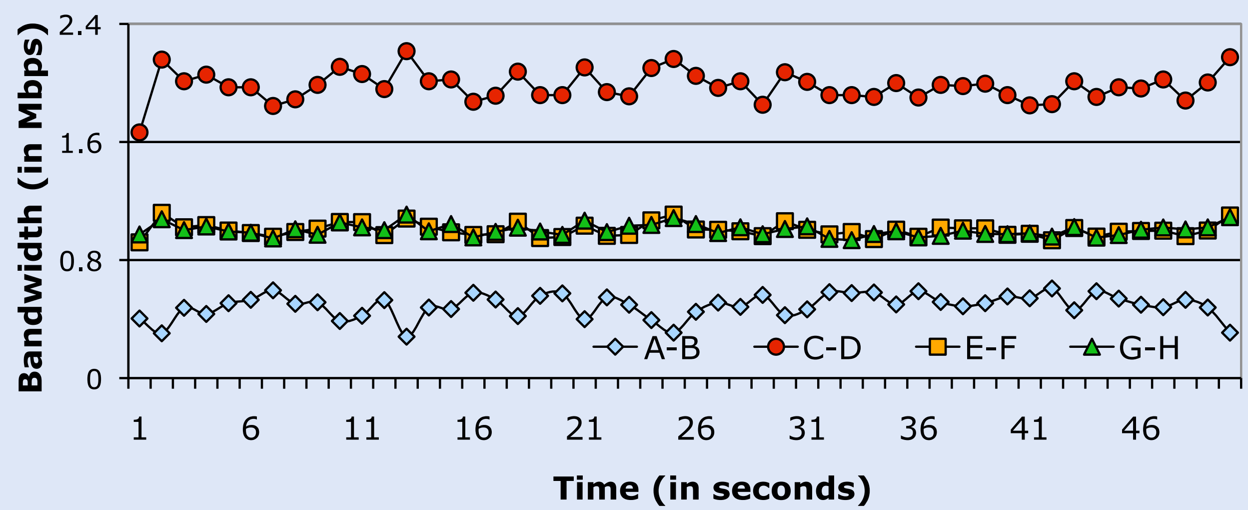

CSFQ

Performance evaluation: Scenario 4

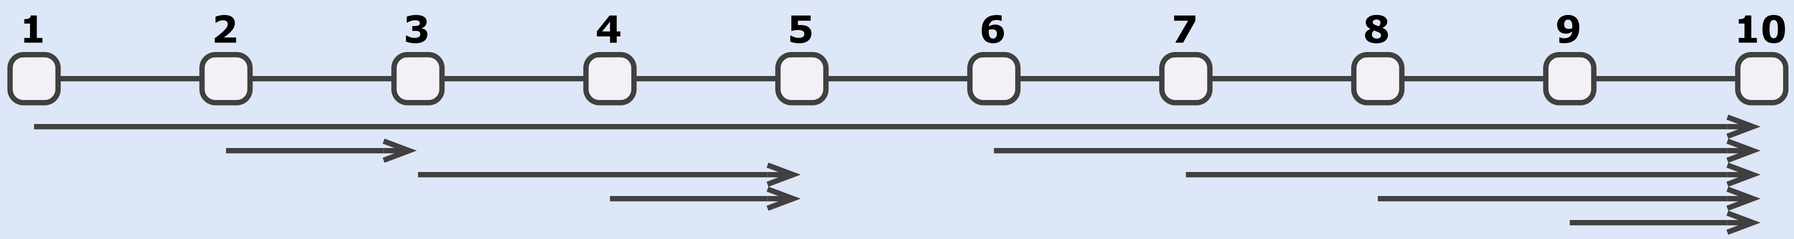

Traffic (s–d)	(1–10)	(2–3)	(3–5)	(4–5)	(6–10)	(7–10)	(8–10)	(9–10)
Start time (in intervals)	1	11	21	31	41	51	61	71

Source	Time (in intervals)							
	1	11	21	31	41	51	61	71
1	1.00	0.50	0.50	0.33	0.33	0.33	0.25	0.20
2	0	0.50	0.50	0.67	0.67	0.67	0.75	0.80
3	0	0	0.50	0.33	0.33	0.33	0.375	0.40
6	0	0	0	0	0.67	0.33	0.25	0.20

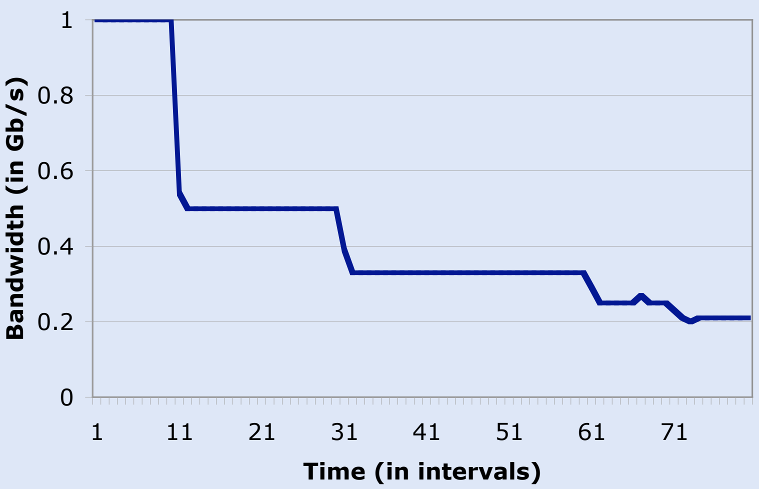

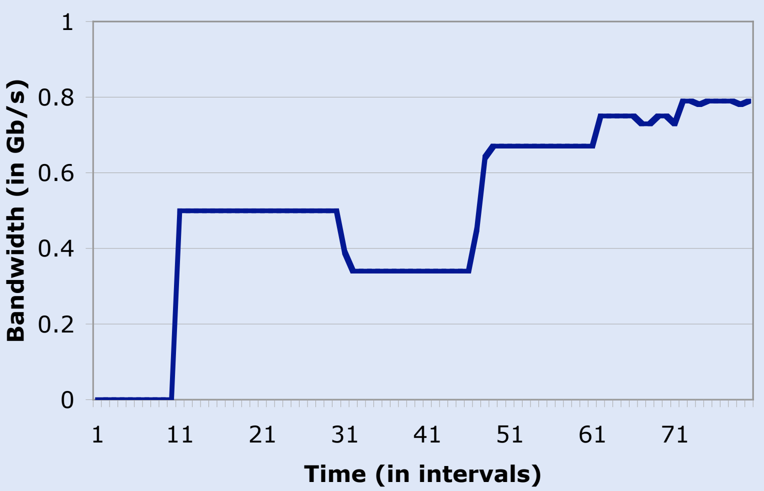

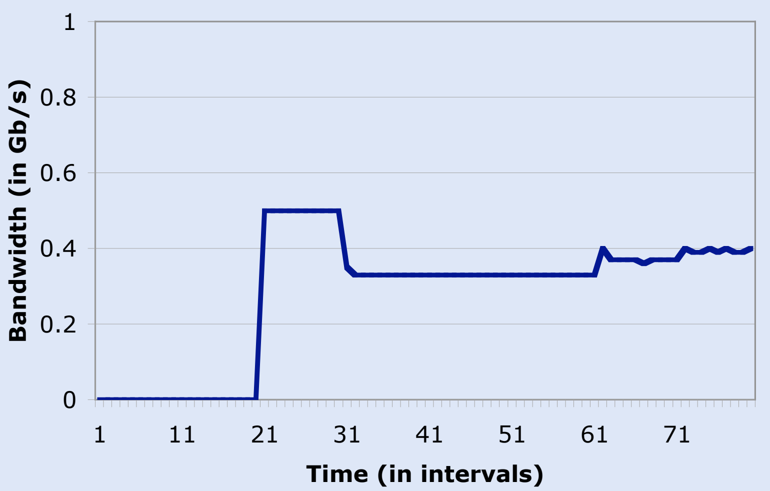

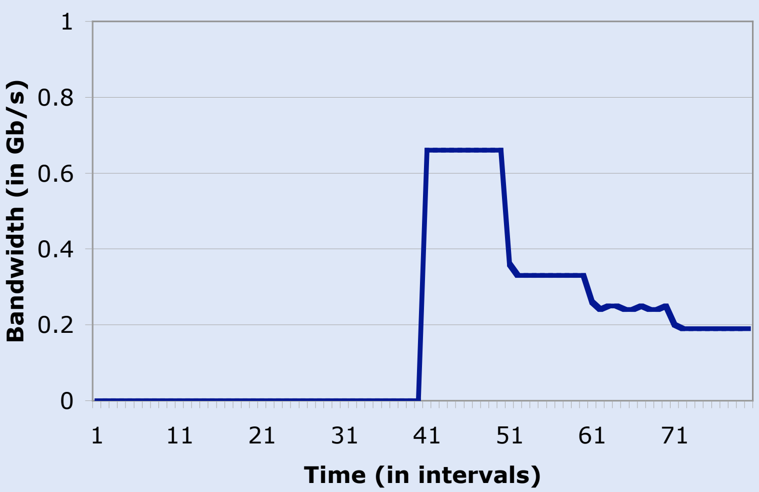

Summary

New framework for fair bandwidth sharing

- No packet dropping
- Transmission and buffering in rounds
- O(1) scheduling complexity
- Provides basic level of fairness

Generalization

- Every packet carries the **rate** at which it has to be transmitted (similar to CSFQ)

Key observation

- Intermediate routers **must not** operate in a work-conserving manner
- CSFQ determines drop probabilities using the rates; but routers are still work-conserving