FCP: A Flexible Transport Framework for Accommodating Diversity

Dongsu Han (KAIST)

Robert Grandl†, Aditya Akella†,
Srinivasan Seshan*

† University of Wisconsin-Madison
*Carnegie Mellon University

Evolution of Transport Protocols

Reliability/Loss recovery
In-order delivery
Flow control
Congestion control

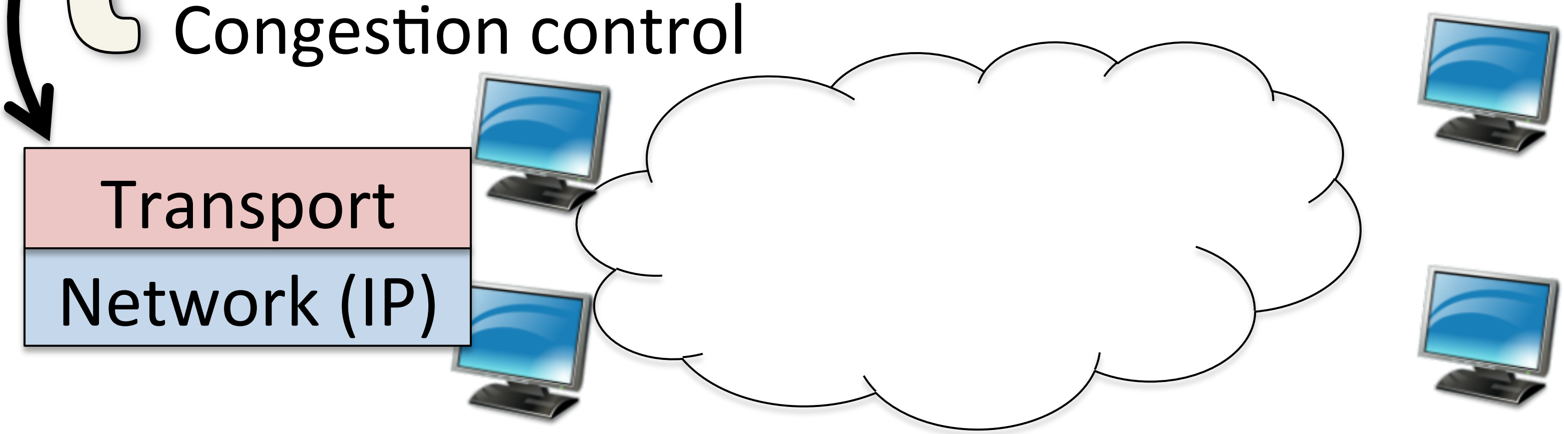

Evolution of Transport Protocols

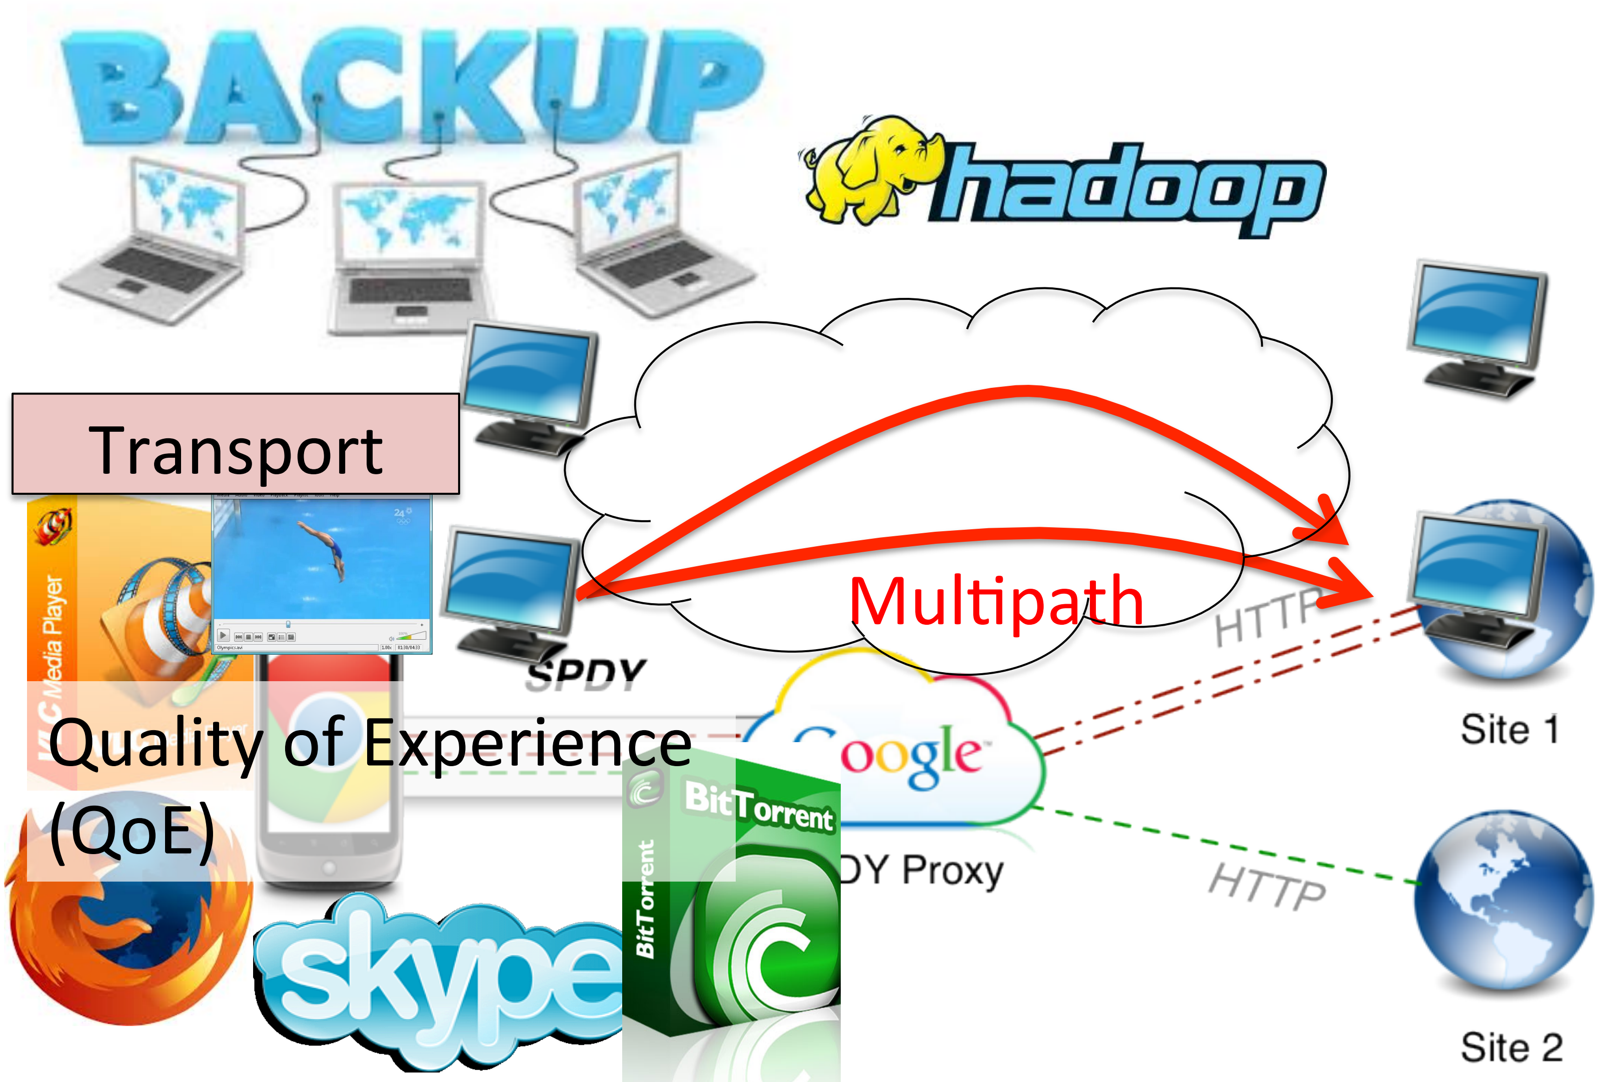

Evolution of Transport Protocols

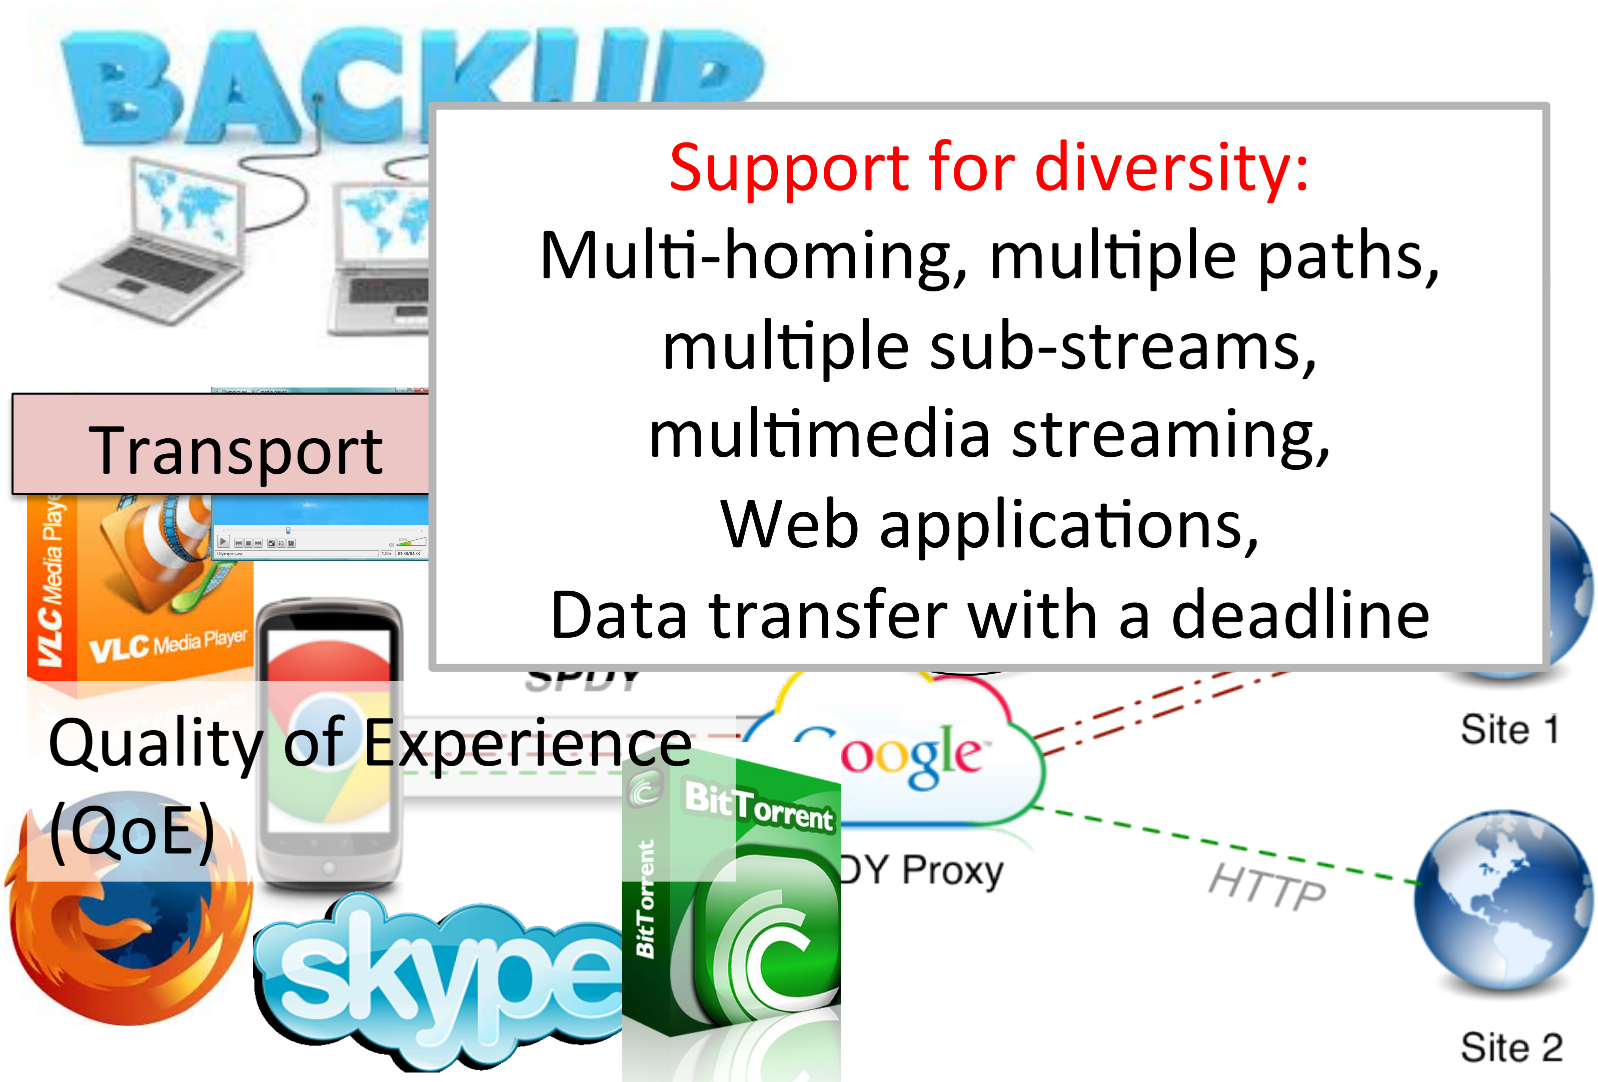

Transport

Support for diversity:
Multi-homing, multiple paths,
multiple sub-streams,
multimedia streaming,
Web applications,
Data transfer with a deadline

Quality of Experience (QoE)

Evolution of Transport Protocols

Reliability/Loss recovery

In-order delivery

Flow control

Congestion control

Transport

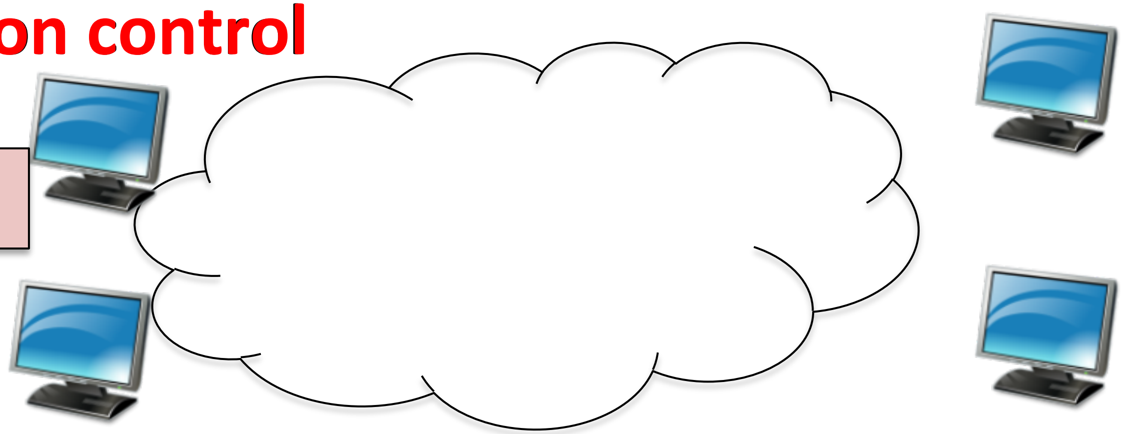

Evolution of Transport Protocols

Resource Allocation

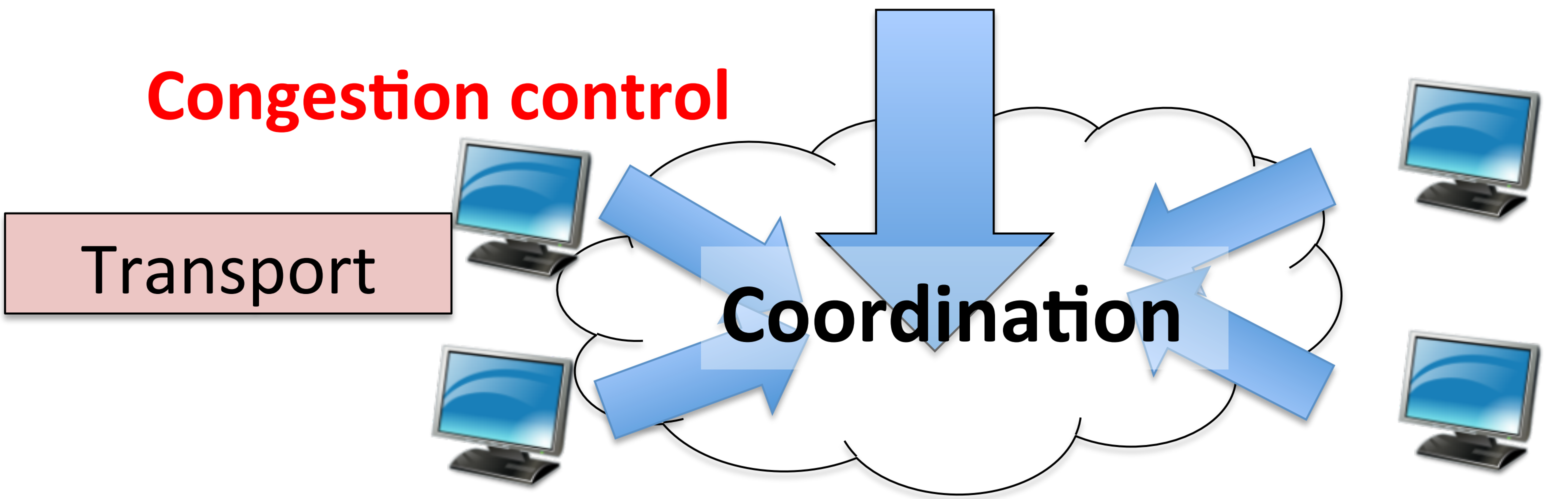

Fairness and efficiency

Problems in Supporting Diversity in Congestion Control

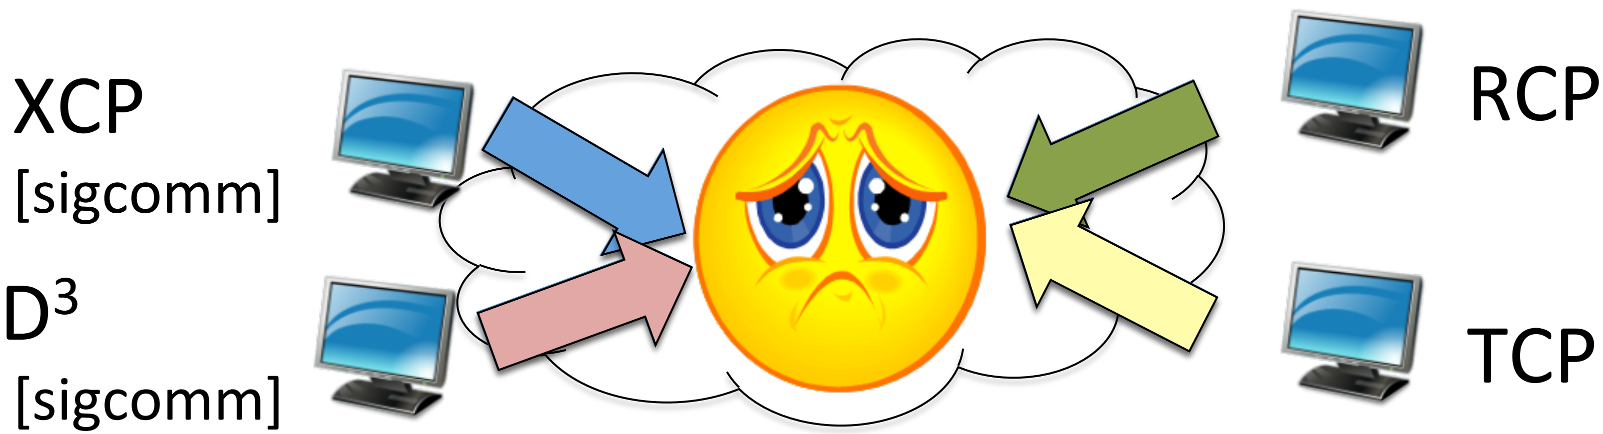

Cannot ensure coexistence.

Evolution of TCP

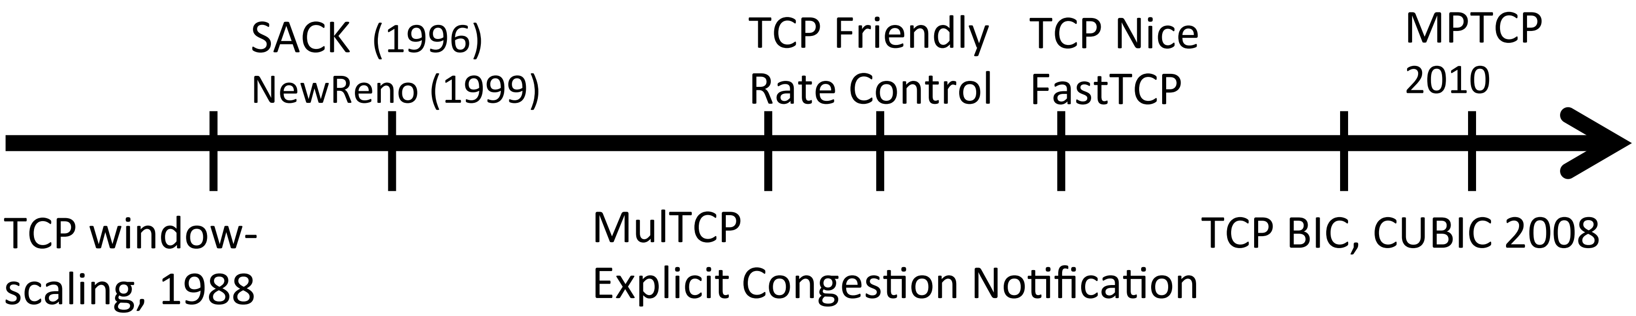

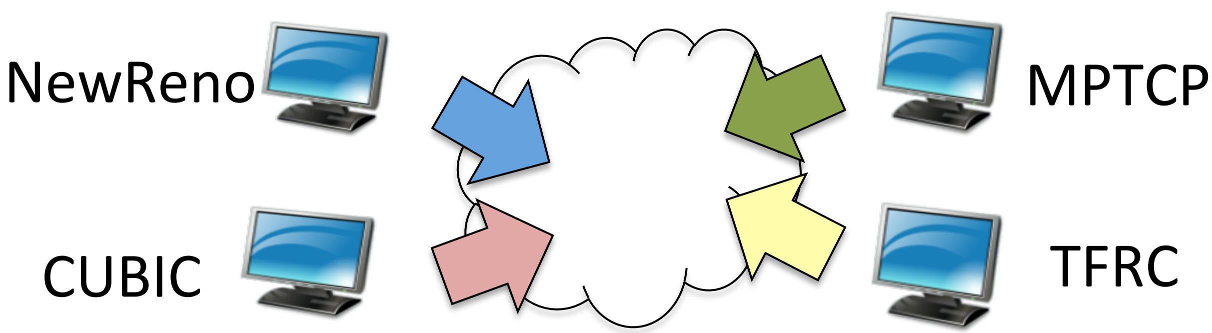

1. End-point flexibility: Purely end-point based
2. Coexistence: Invariant for fairness

(TCP friendliness)

End-point based vs. Router-Assisted

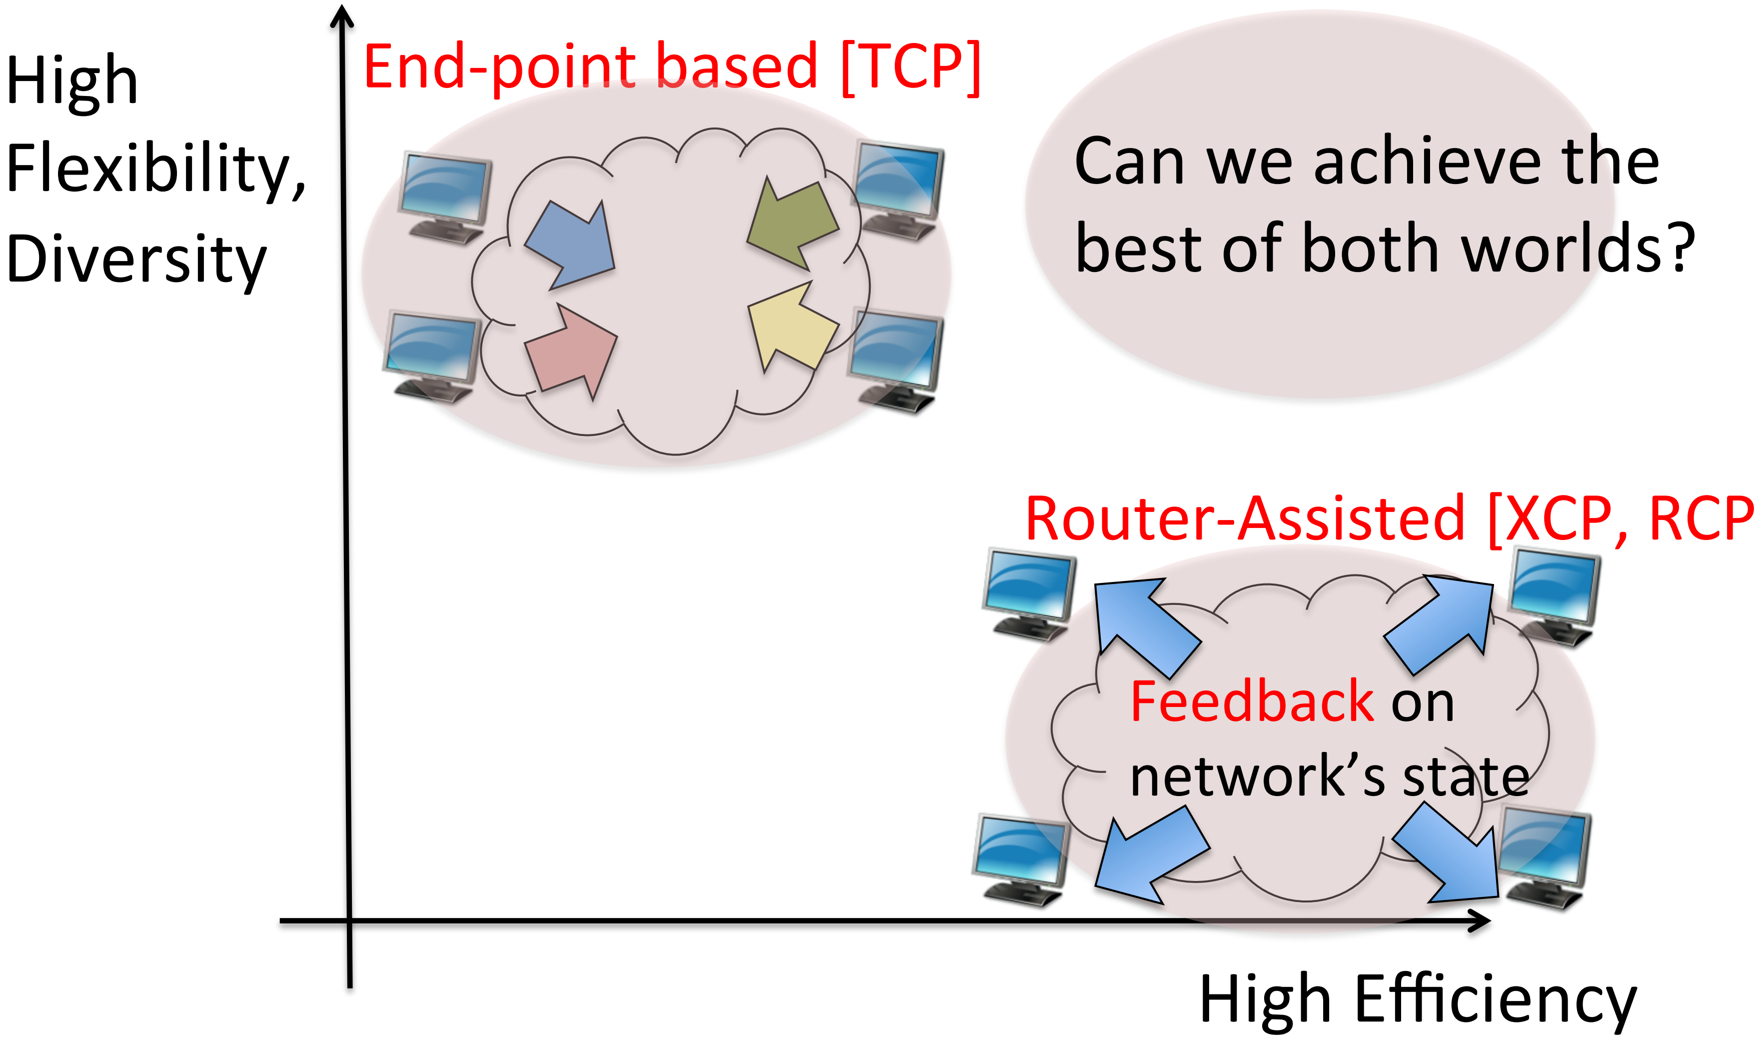

Our Approach

1. Decouple coexistence issues (fairness and efficiency) from end-point control

2. Introduce generic abstractions for resource allocation.

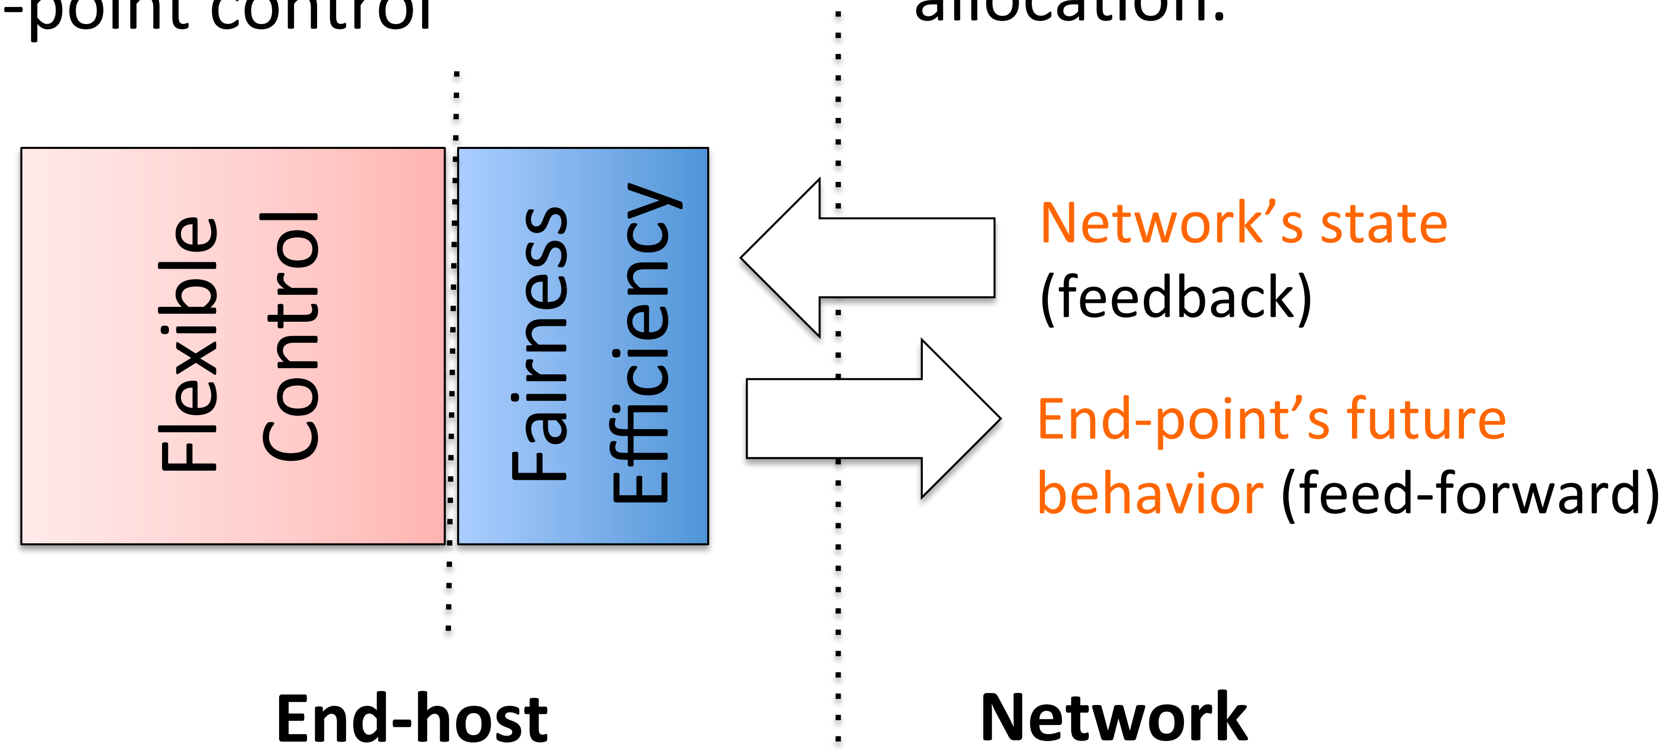

1 Decoupling for Flexibility

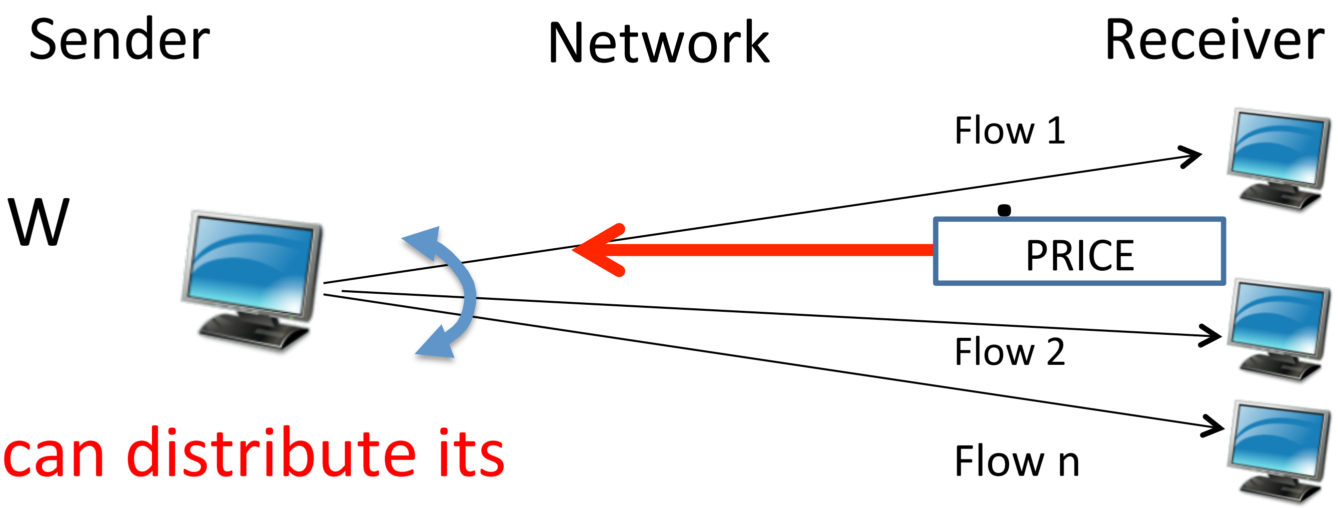

Budget W
($/sec)

Sender can distribute its budget to its flows.

Flow i's budget $w_{i,}$ subject to $W \geq \Sigma w_i$

Flow i's price: P_i ($/Byte)

Flow i's rate:

$$R_i = \text{budget/price} = w_i / P_i \text{ (Byte/sec)}$$

1 Decoupling for Flexibility

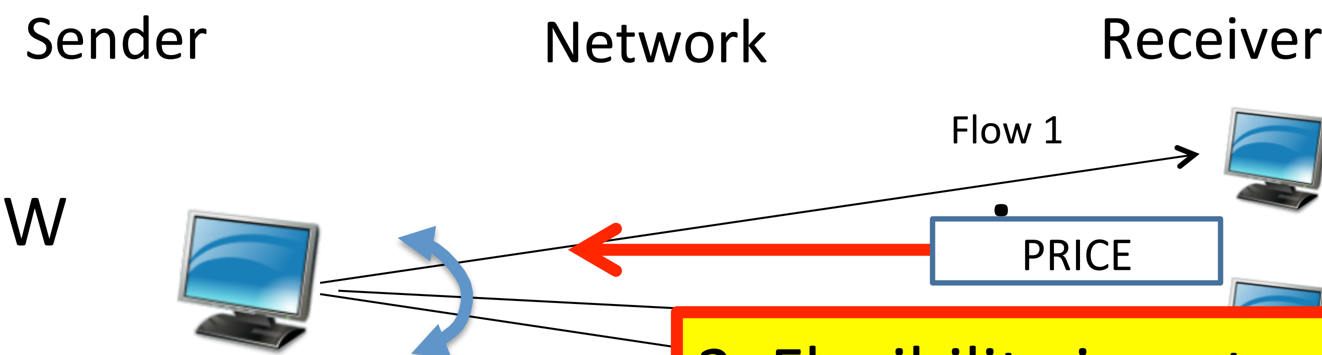

Budget W
($/sec)

2. Flexibility in network price generation.

Invariant for fairness

Flow i's budget w_i, subject to $W \geq \Sigma w_i$

1. Flexibility at the end-points in how its budget is used.

2a Feedback: Pricing

- Feedback: "congestion price" reflecting the "cost" of sending data across the link [Kelly]

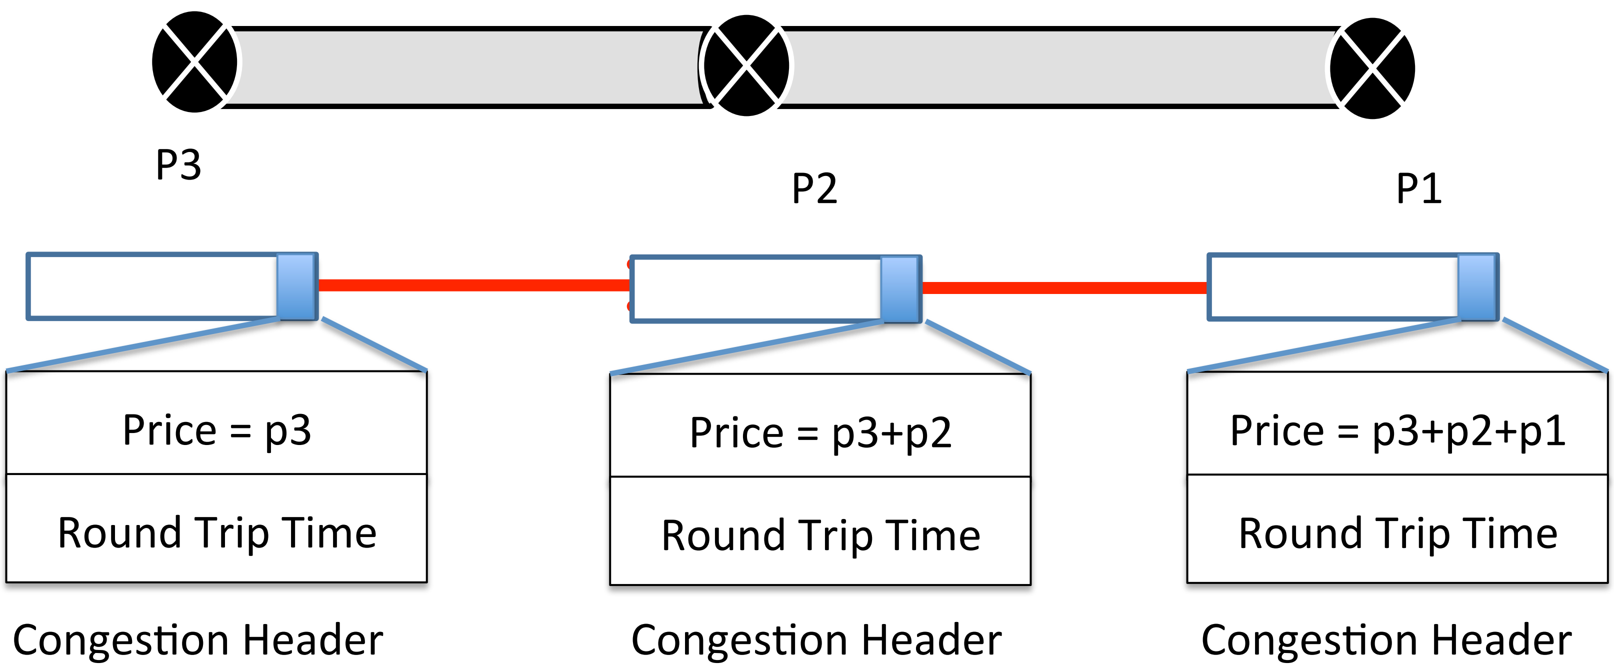

2a Feedback: Pricing

- Feedback: “congestion price” reflecting the “cost” of sending data across the link [Kelly]

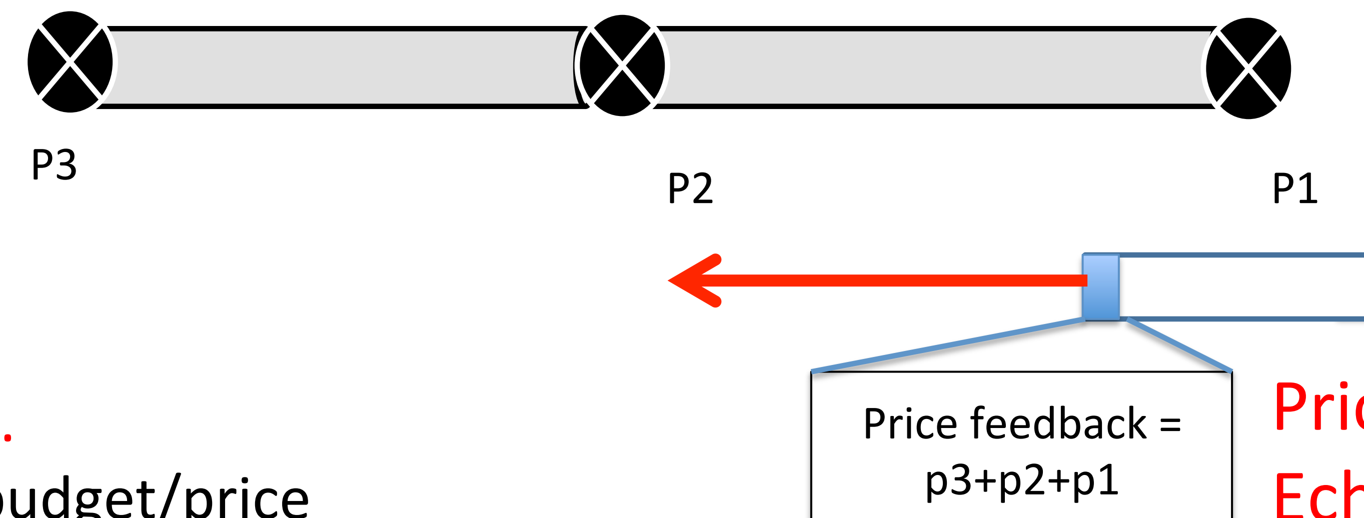

Sender updates the rate.

Rate = budget/price

Price Echoback

This implements proportional fairness [Kelly].

2a Feedback: Pricing

Router updates the price, upon packet reception.

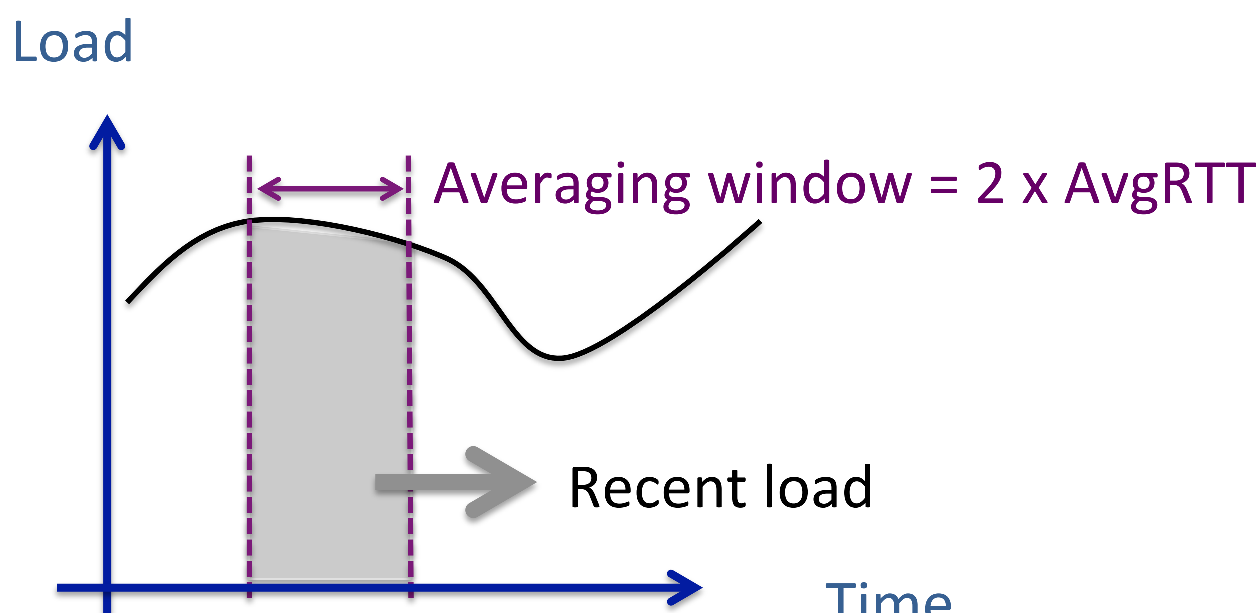

Price ($/byte) = f(Average recent load)

2a Feedback: Pricing

Router updates the price, upon packet reception.

Instantaneous Incoming budget ($)
= Price ($/byte) x Bytes Received (bytes)

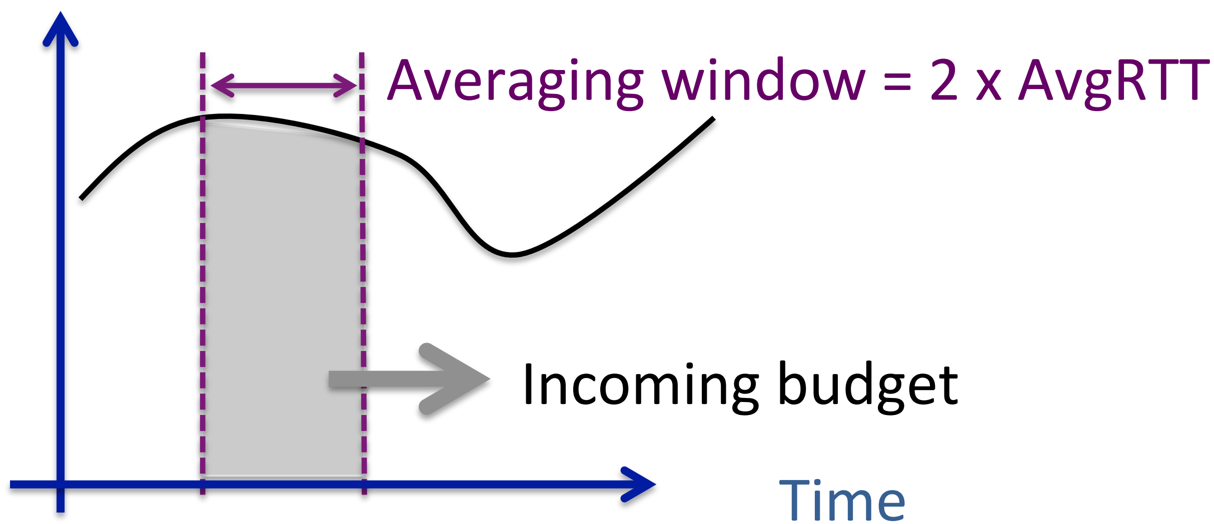

$$\text{Price (\$/byte)} = \frac{\text{Average Incoming Budget}}{\text{Remaining Link Capacity}}$$

2a Feedback: Pricing

Router stores recent history of price, $price(t)$

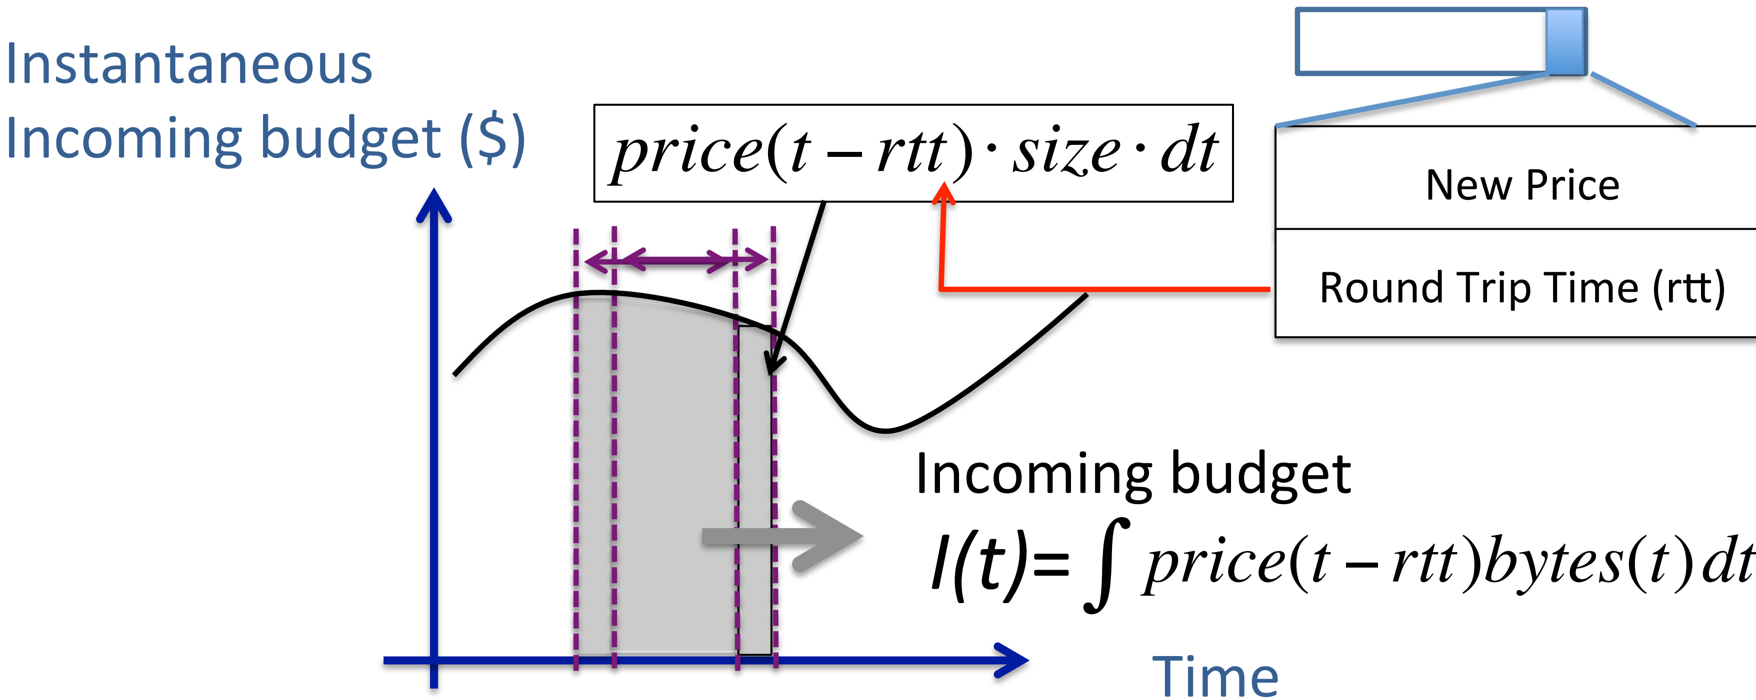

$$\text{Price (\$/byte)} = \frac{\text{Average Incoming Budget}}{\text{Remaining Link Capacity}}$$

2b Feed-forward: Preloading

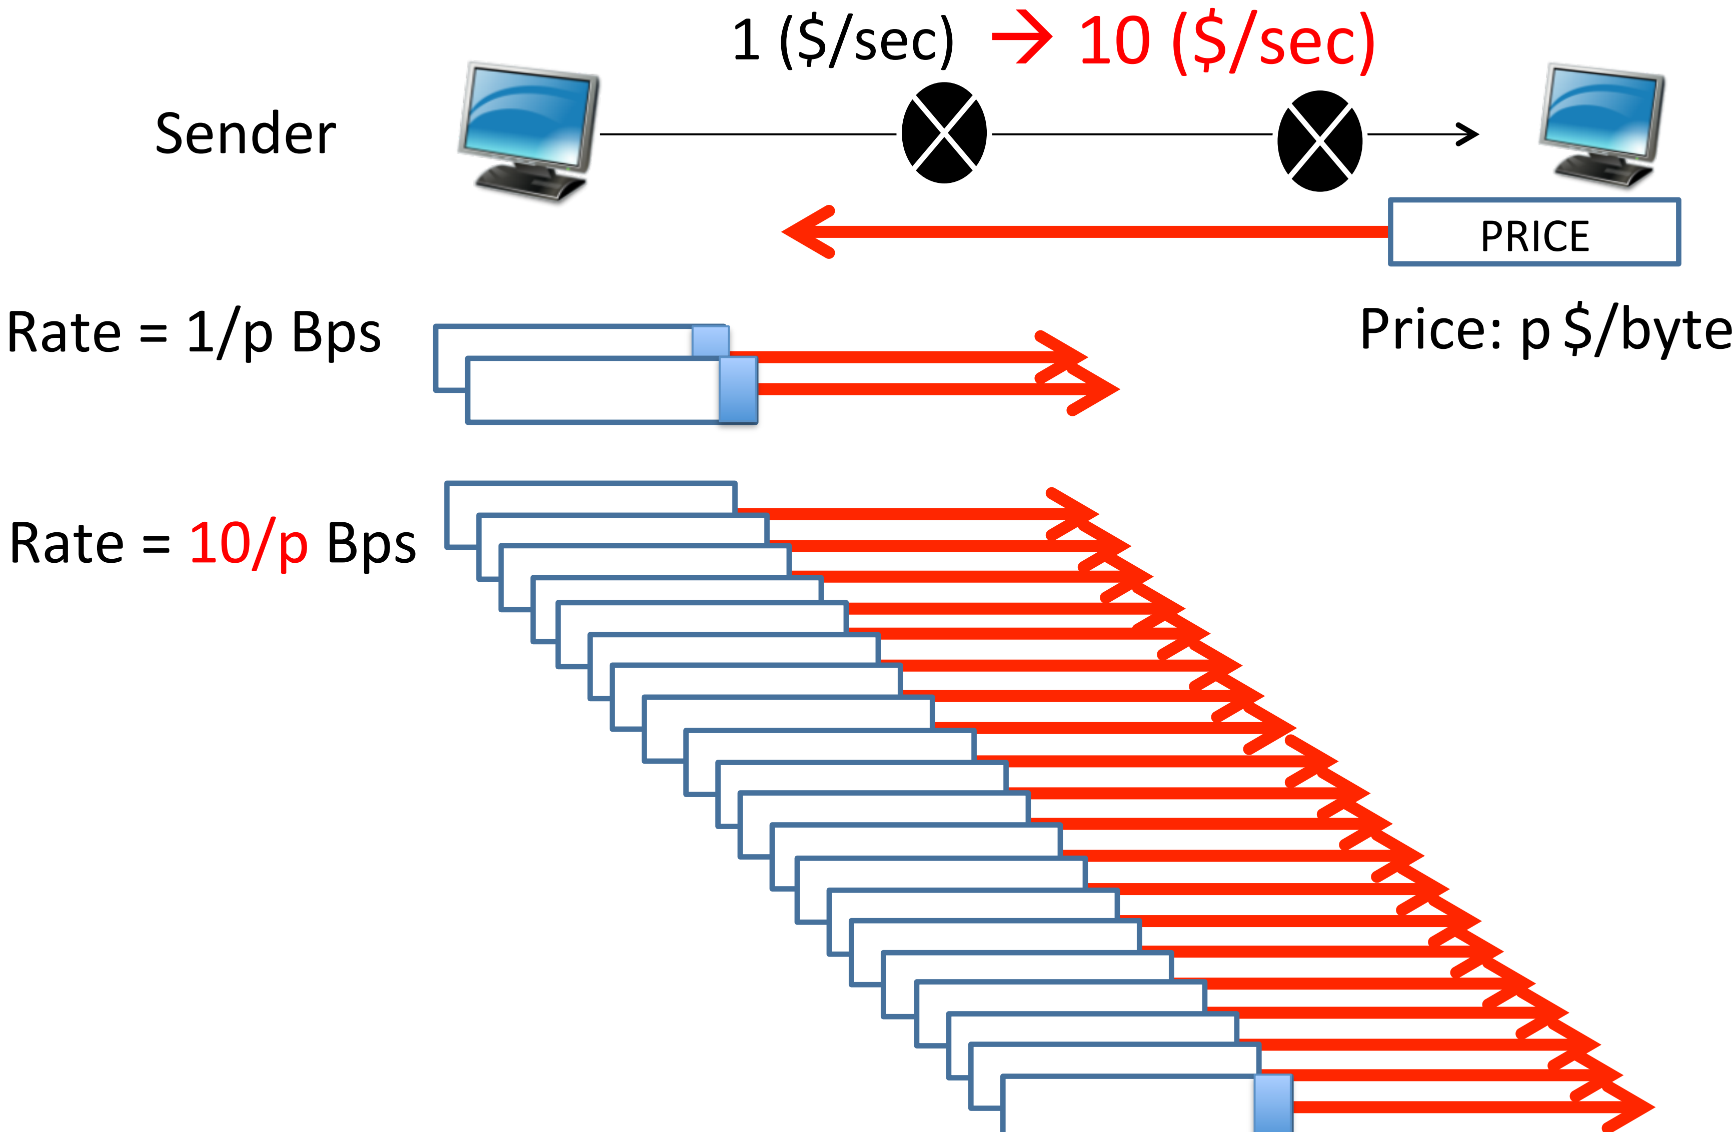

2b Feed-forward: Preloading

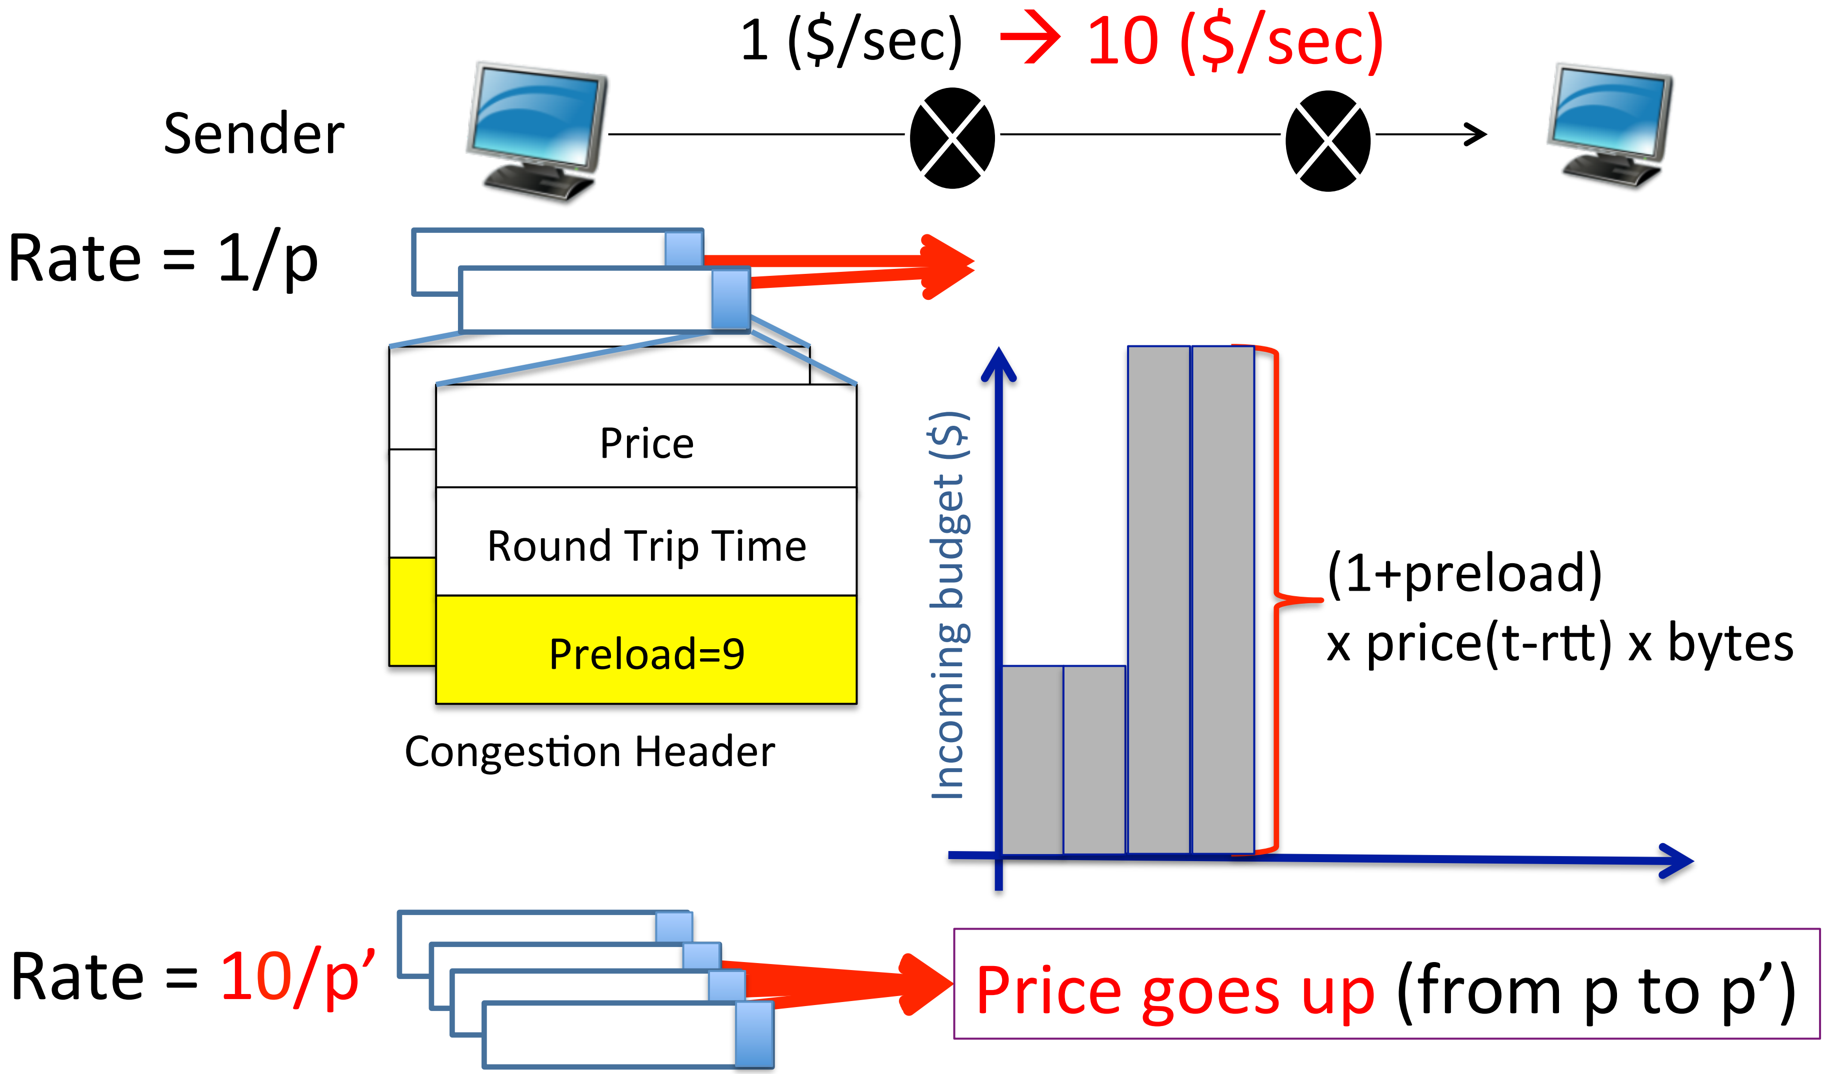

Evaluation

1. Basic performance of FCP

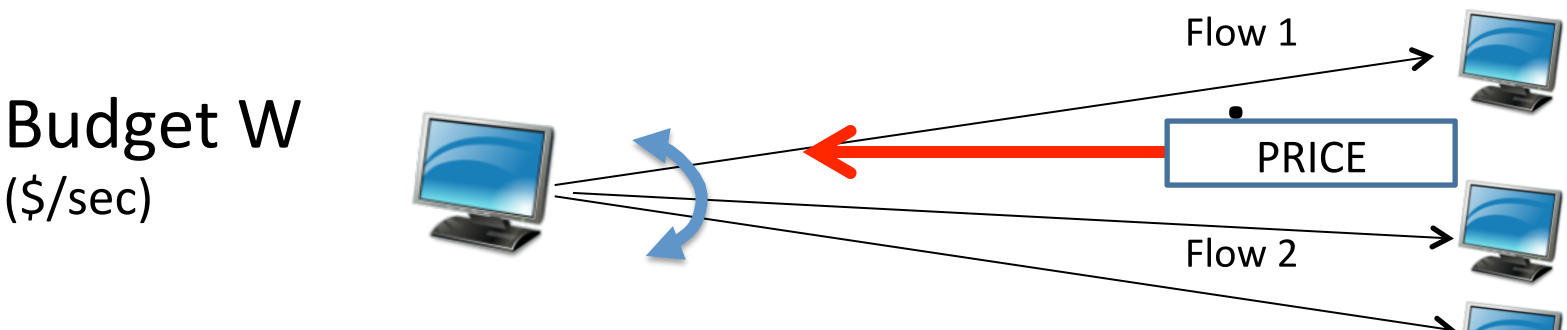

2. End-point budget allocation.

3. Flexibility in network price generation.

Fast Convergence/Accurate Feedback

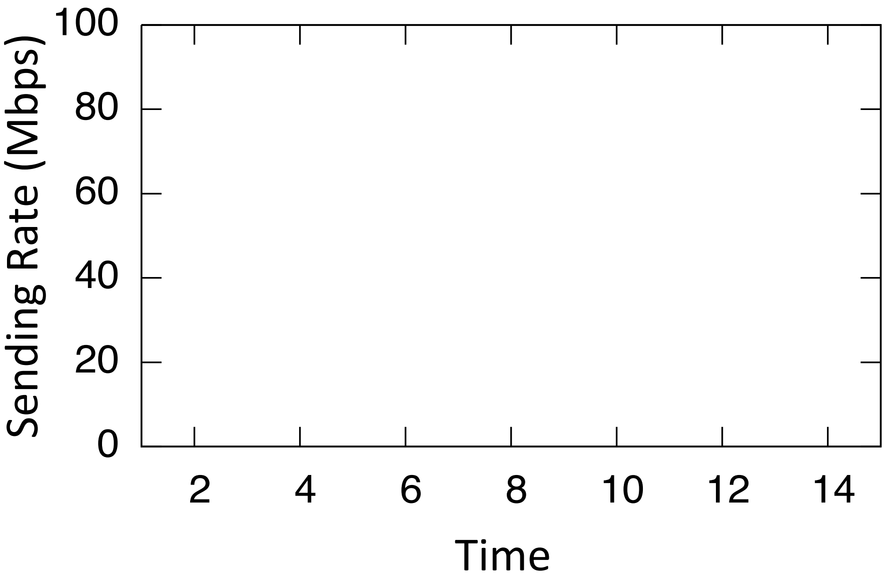

RTT = {25ms, 50ms, 125ms, 250ms, 625ms }

Fast Convergence/Accurate Feedback

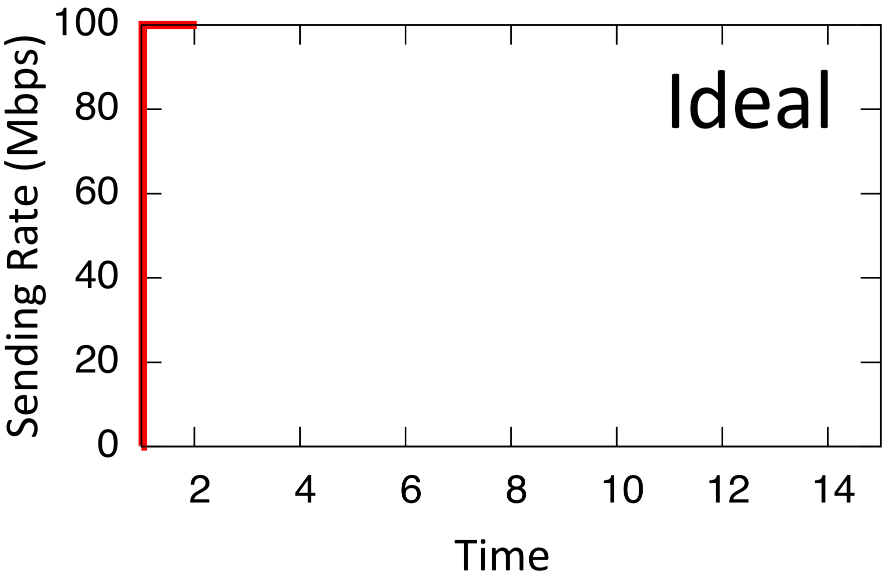

RTT = {25ms, 50ms, 125ms, 250ms, 625ms }

Fast Convergence/Accurate Feedback

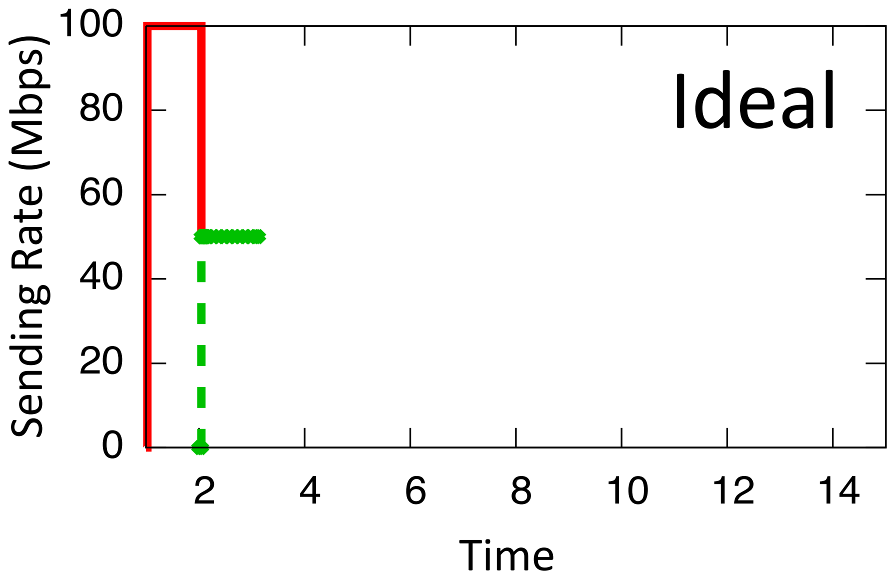

RTT = {25ms, 50ms, 125ms, 250ms, 625ms }

Fast Convergence/Accurate Feedback

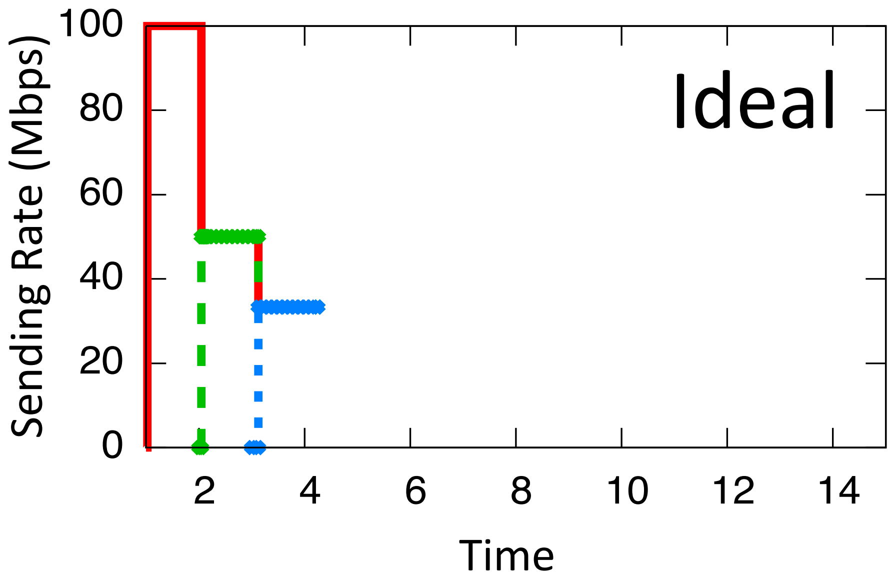

RTT = {25ms, 50ms, 125ms, 250ms, 625ms }

Fast Convergence/Accurate Feedback

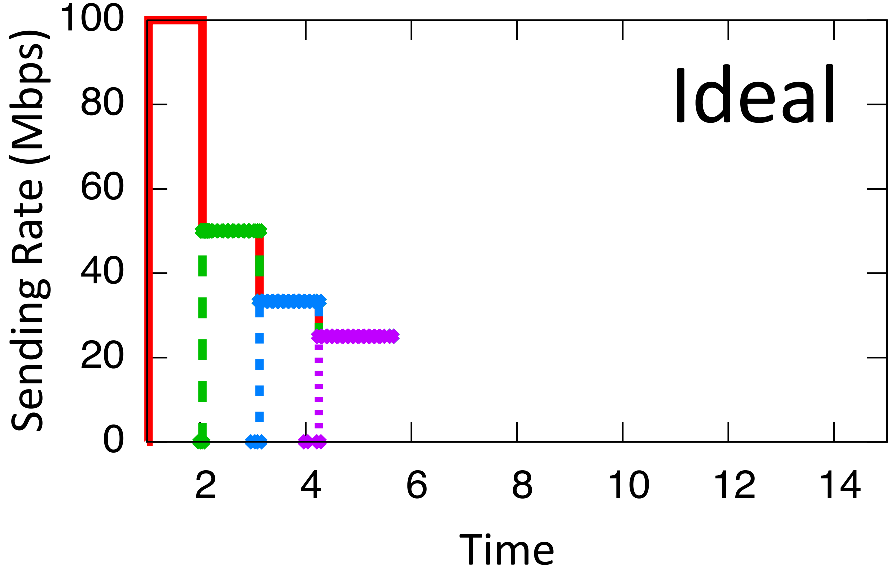

RTT = {25ms, 50ms, 125ms, 250ms, 625ms }

Fast Convergence/Accurate Feedback

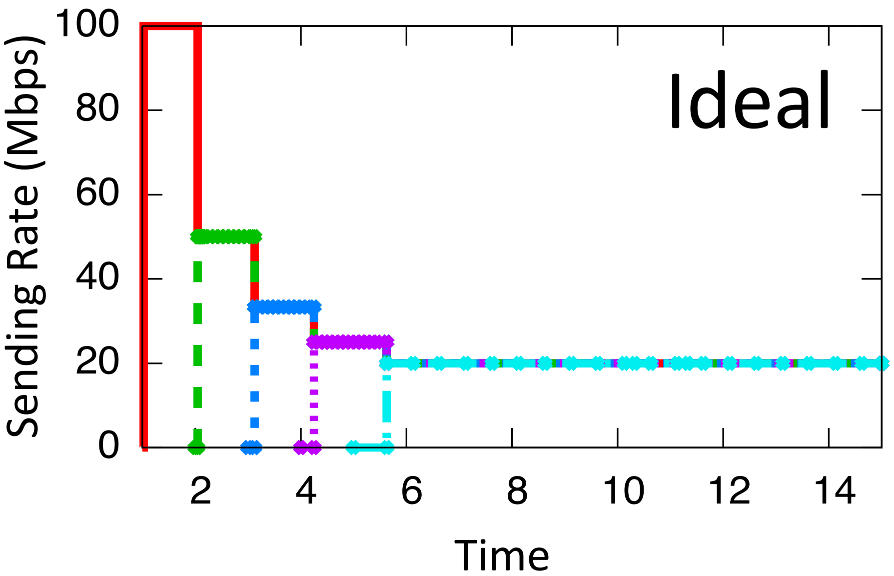

RTT = {25ms, 50ms, 125ms, 250ms, 625ms }

Fast Convergence/Accurate Feedback

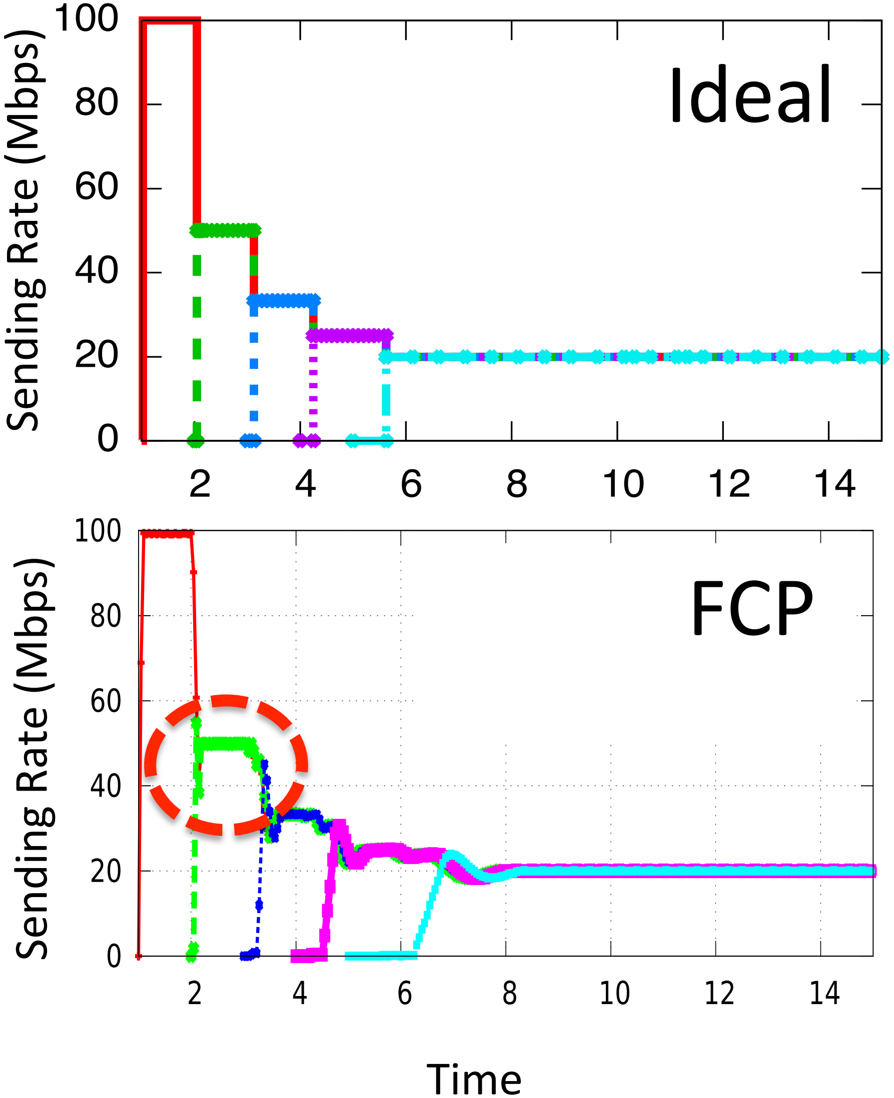

Fast Convergence/Accurate Feedback

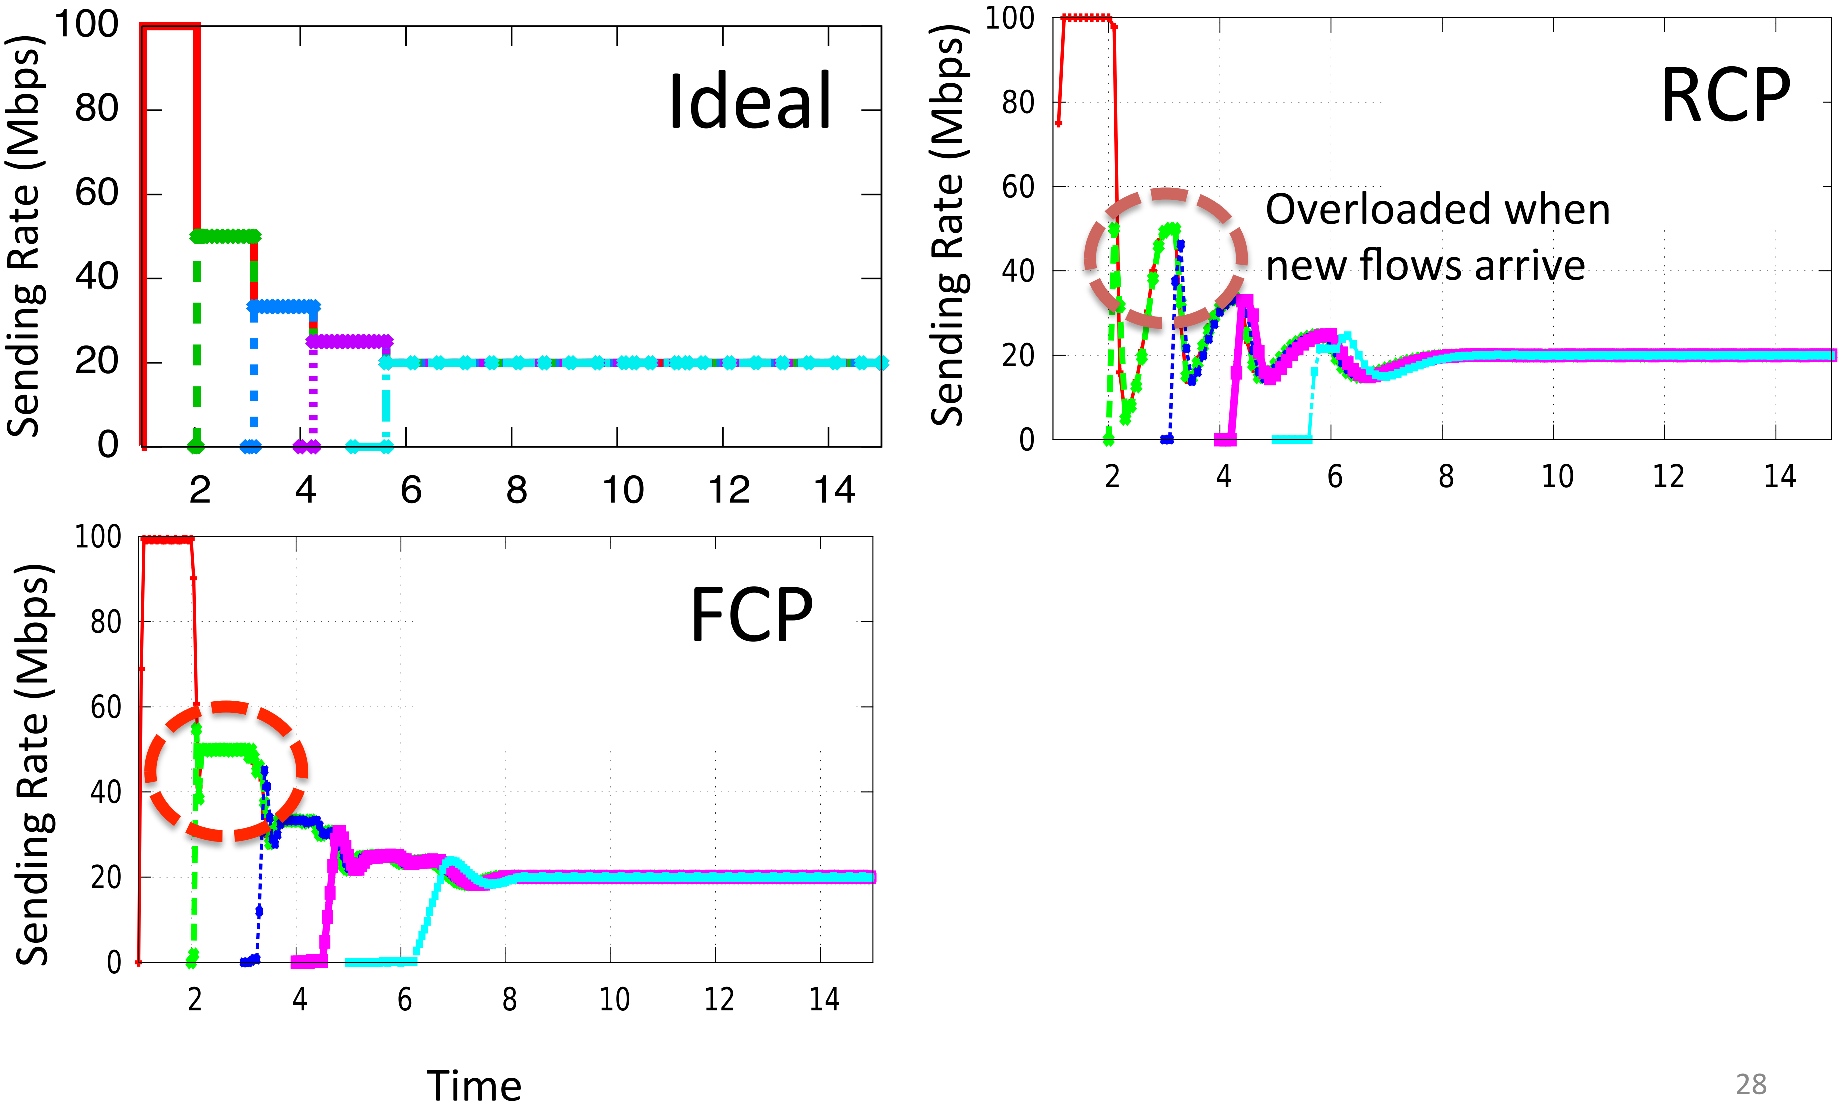

Fast Convergence/Accurate Feedback

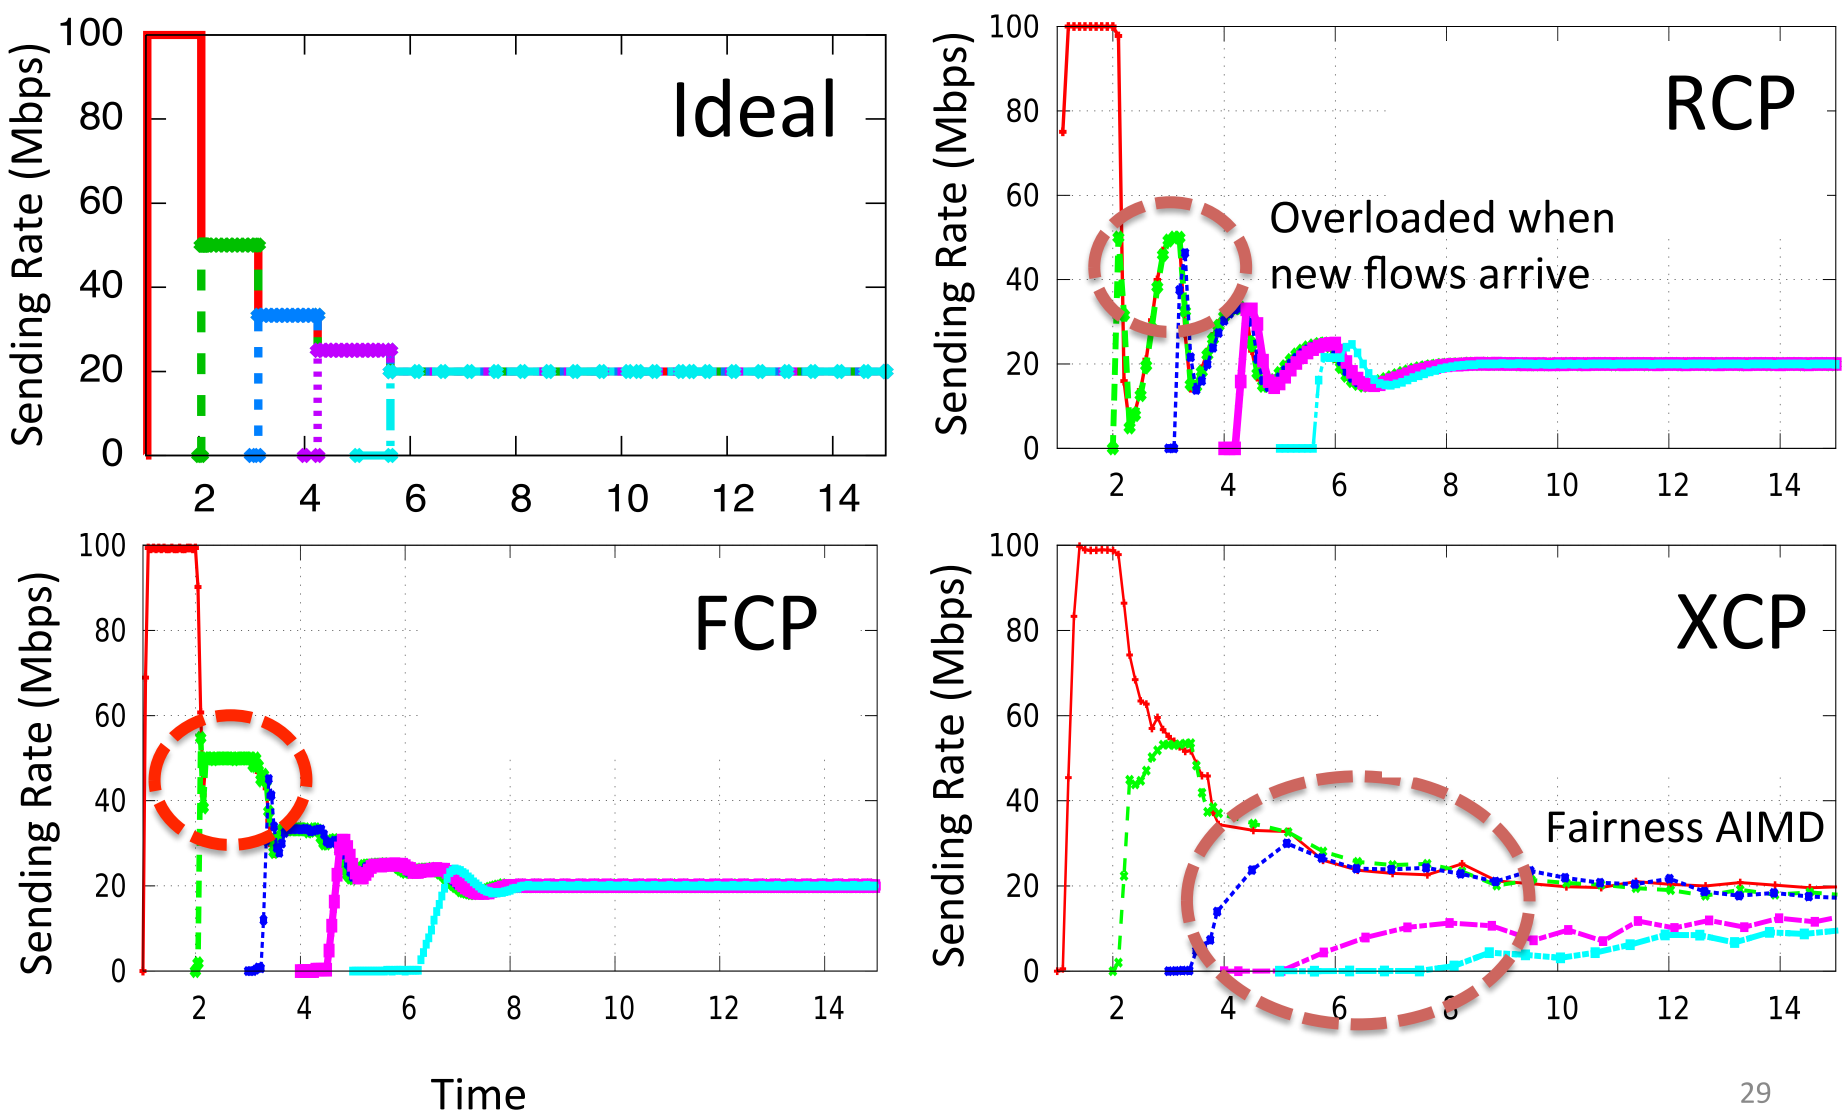

Faster Convergence/Accurate Feedback

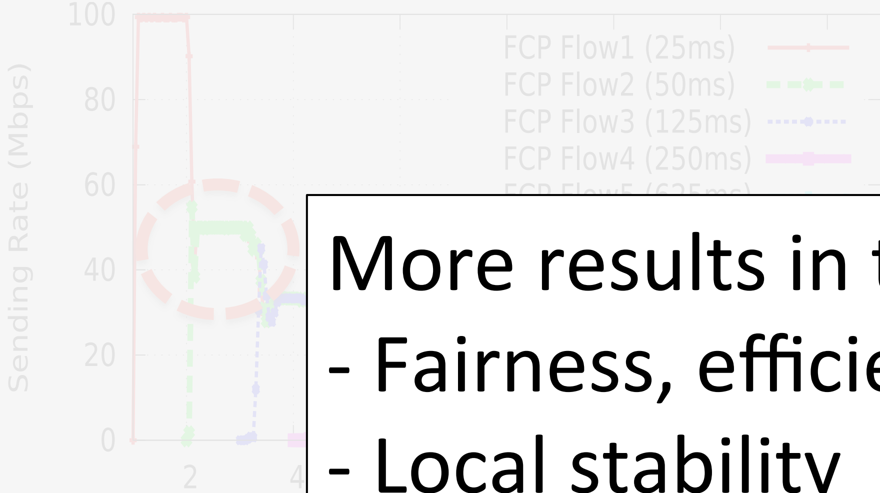

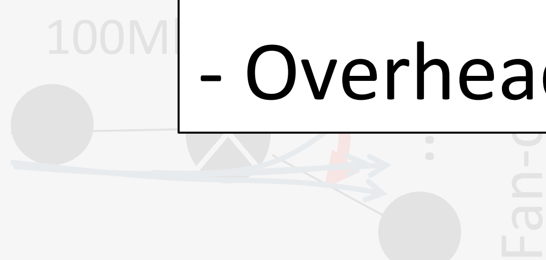

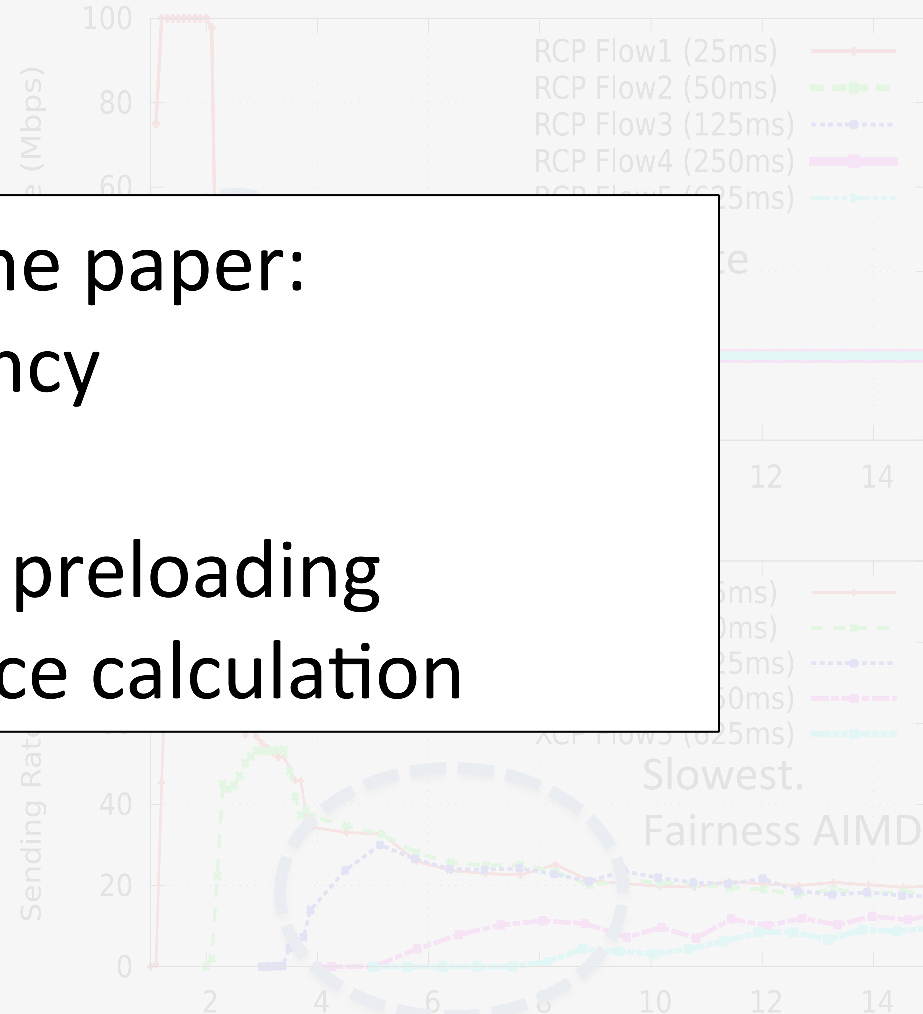

More results in the paper:
- Fairness, efficiency
- Local stability
- Effectiveness of preloading
- Overhead of price calculation

FCP's converges faster and the feedback is more accurate.

End-point Flexibility: Topology

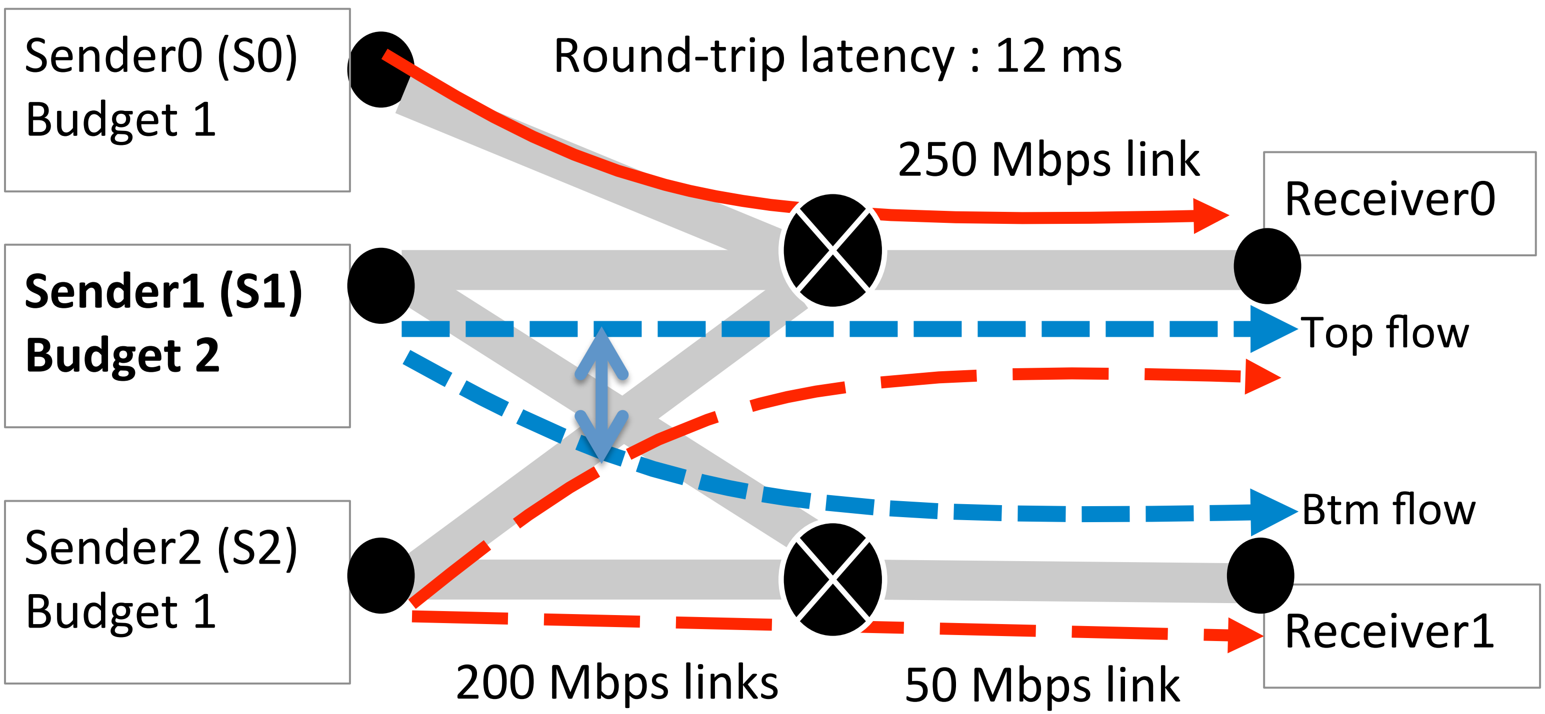

Budget allocation is up to end-points

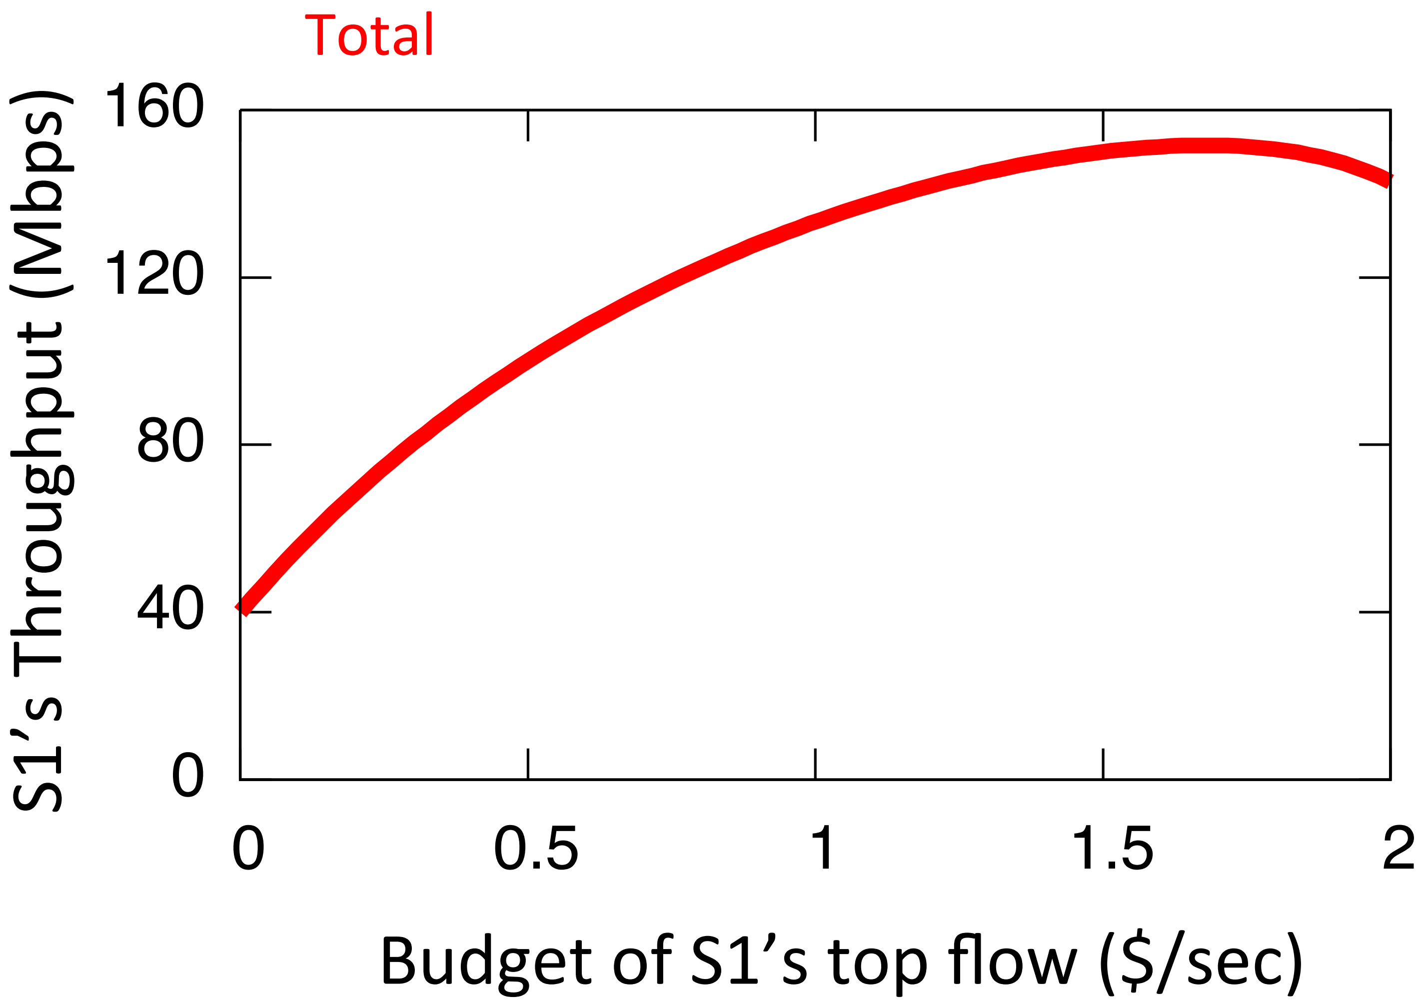

Budget allocation is up to end-points

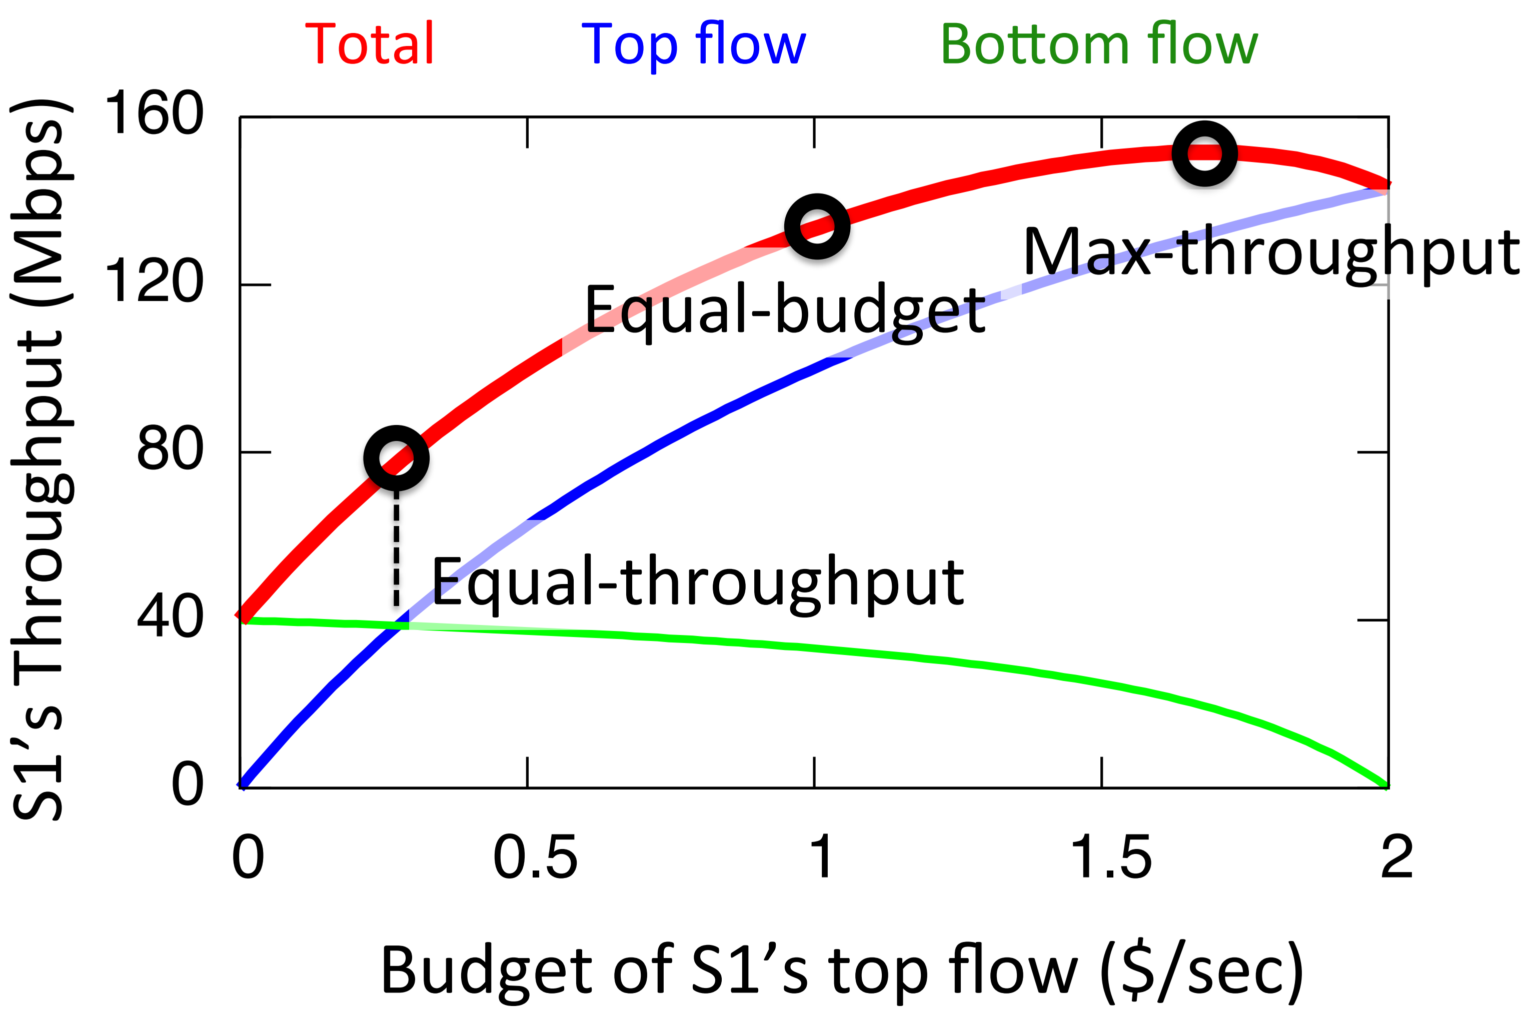

End-point Diversity: Background flows

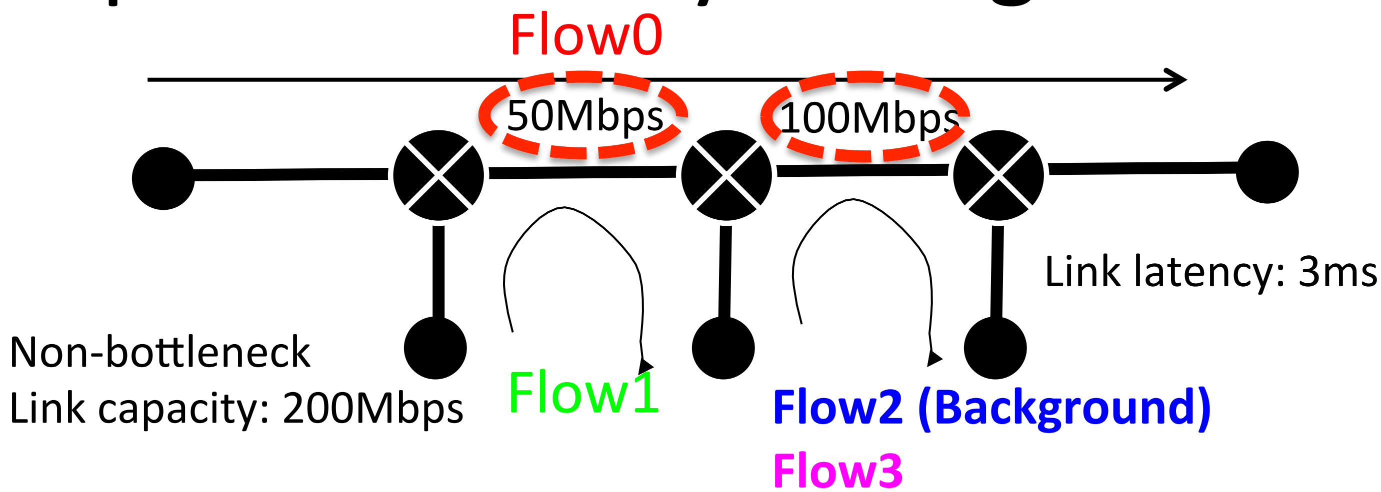

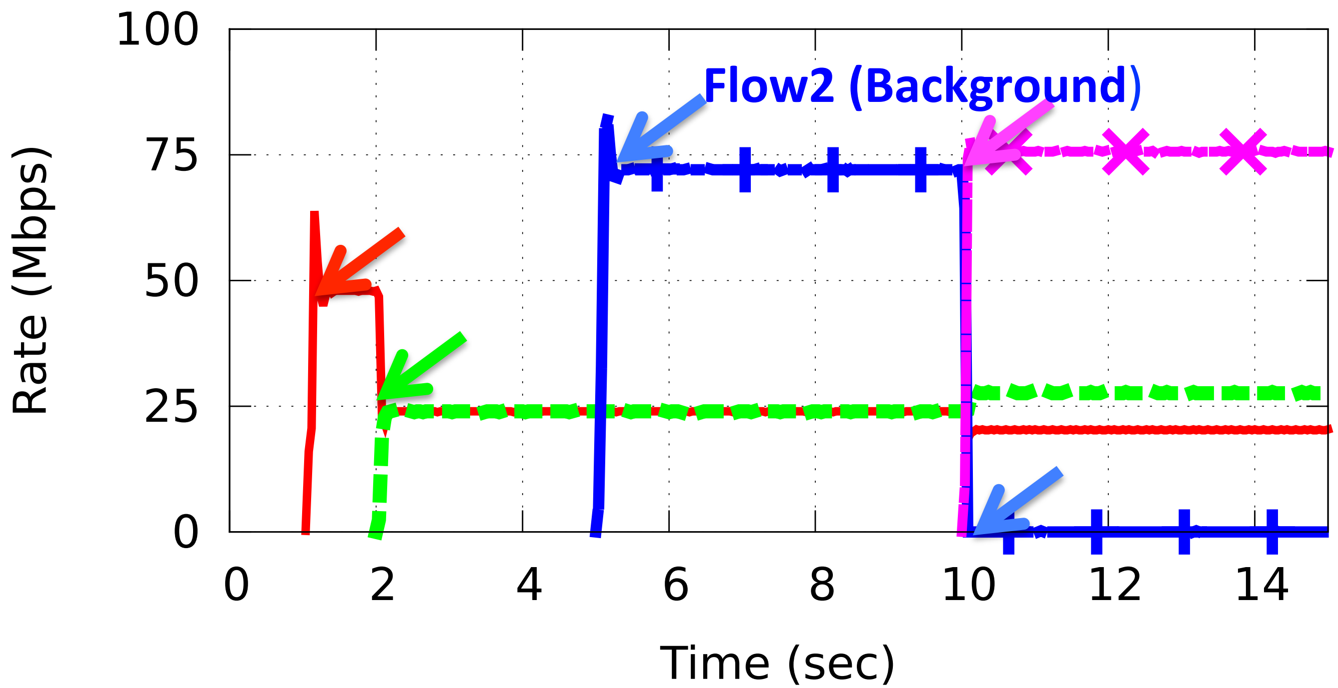

Diversity in Network Pricing

- FCP can also support diverse behaviors in network price generation.
- Examples (in the paper)
  - Deadline support [D^3 SIGCOMM 11]
  - Aggregate resource allocation in a multi-tenant data-center
  - Stable bandwidth allocation for streaming
  - Multicast congestion control

Deadline Support [D^3 SIGCOMM 11]

Deadline flows preload to specify the desired rate

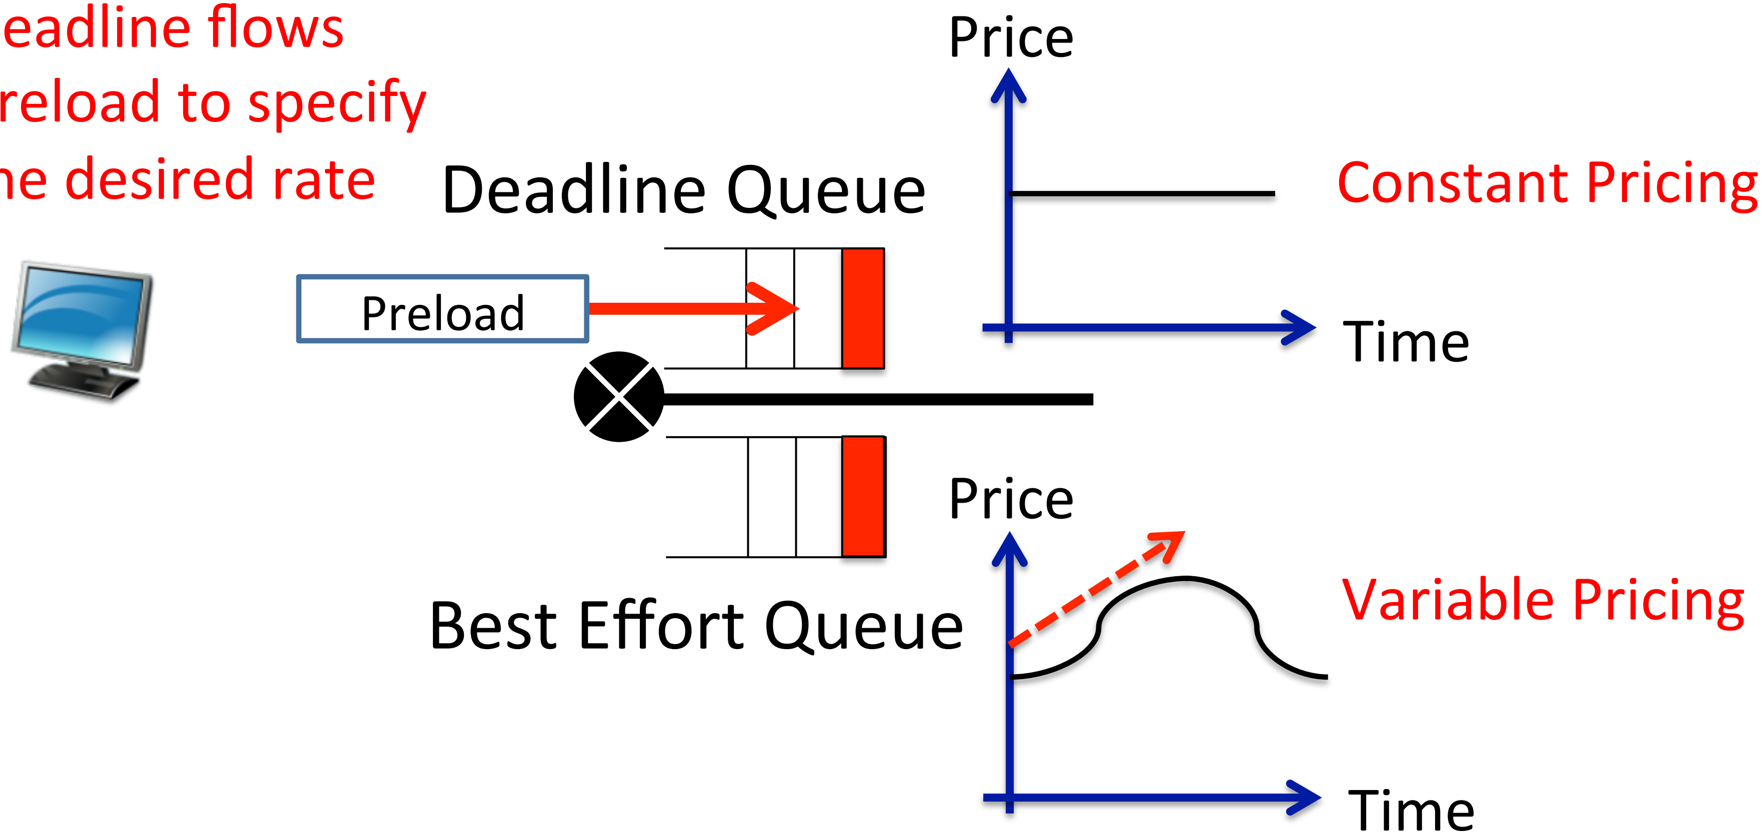

$$\text{Price (\$/byte)} = \frac{\text{Average Incoming Budget}}{\text{Remaining Link Capacity}}$$

Differential Pricing for Deadline Support

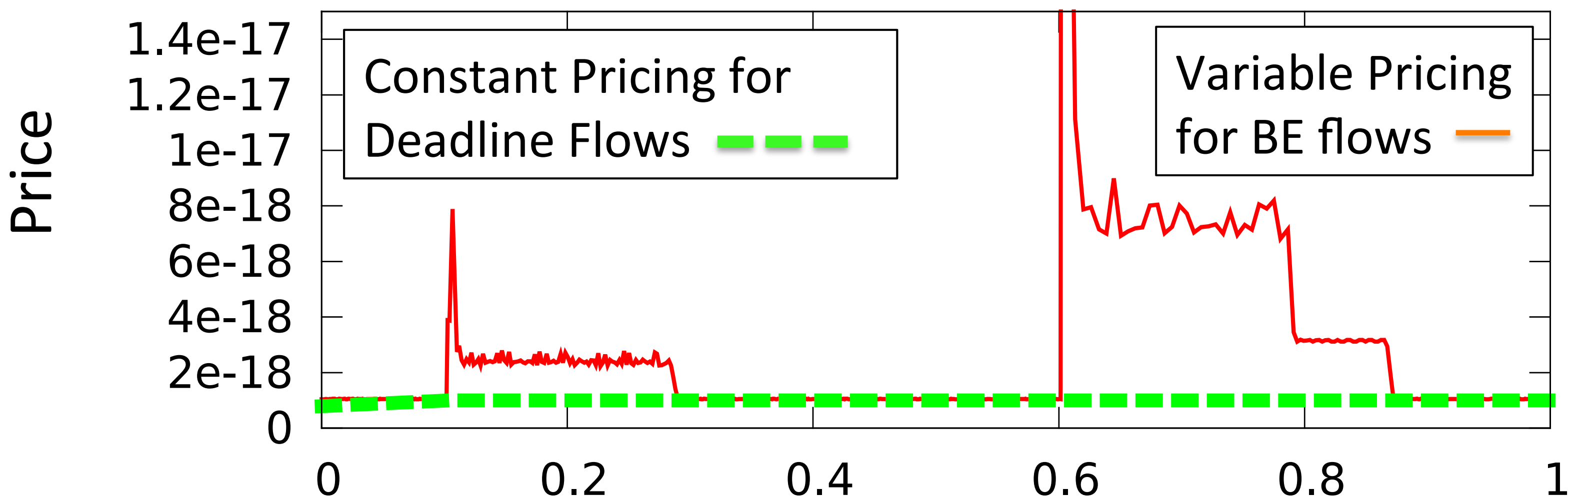

Differential Pricing for Deadline Support

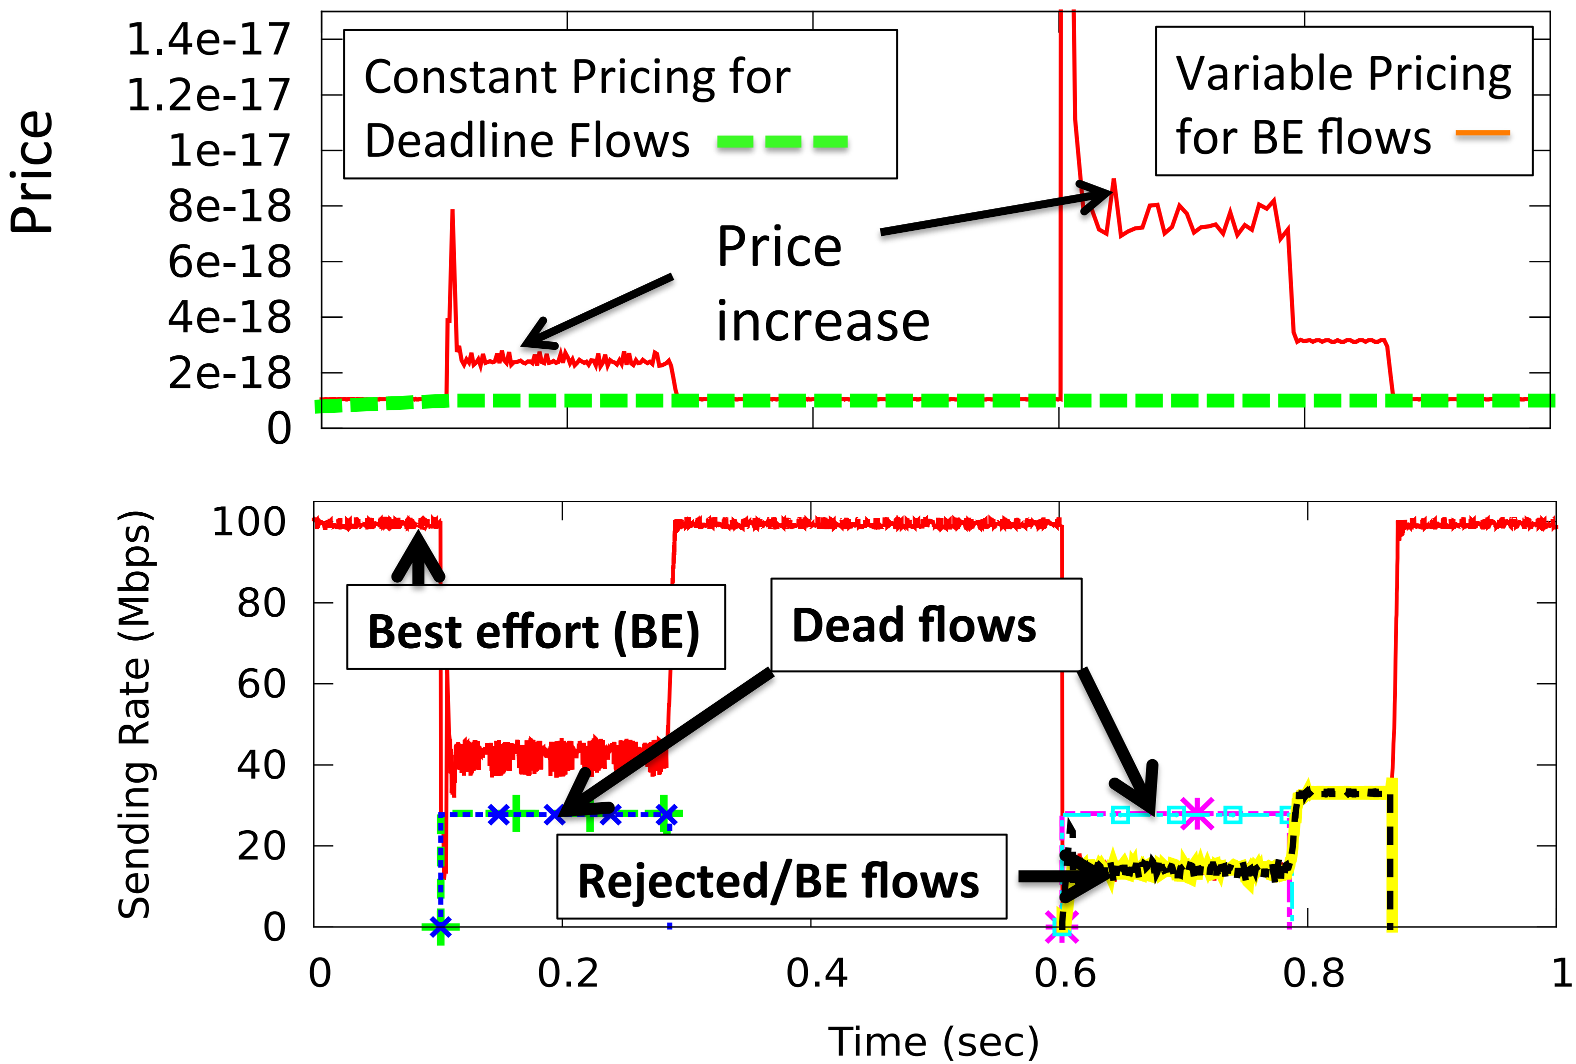

Conclusion

- FCP accommodates diverse behaviors in resource allocation while utilizing explicit feedback.
- FCP maximizes end-point's flexibility by simplifying the mechanism of coexistence.
- FCP's explicit feed-back and feed-forward provides a generic interface for efficient resource allocation.