Principles for Internet Congestion Management

Lloyd Brown, Albert Gran Alcoz, Frank Cangialosi, **Akshay Narayan**, Mohammad Alizadeh, Hari Balakrishnan, Eric Friedman, Ethan Katz-Bassett, Arvind Krishnamurthy, Michael Schapira, Scott Shenker

ETH zürich

BROWN

COLUMBIA UNIVERSITY

Congestion Control Today: A Wild West

Congestion Control Today: A Wild West

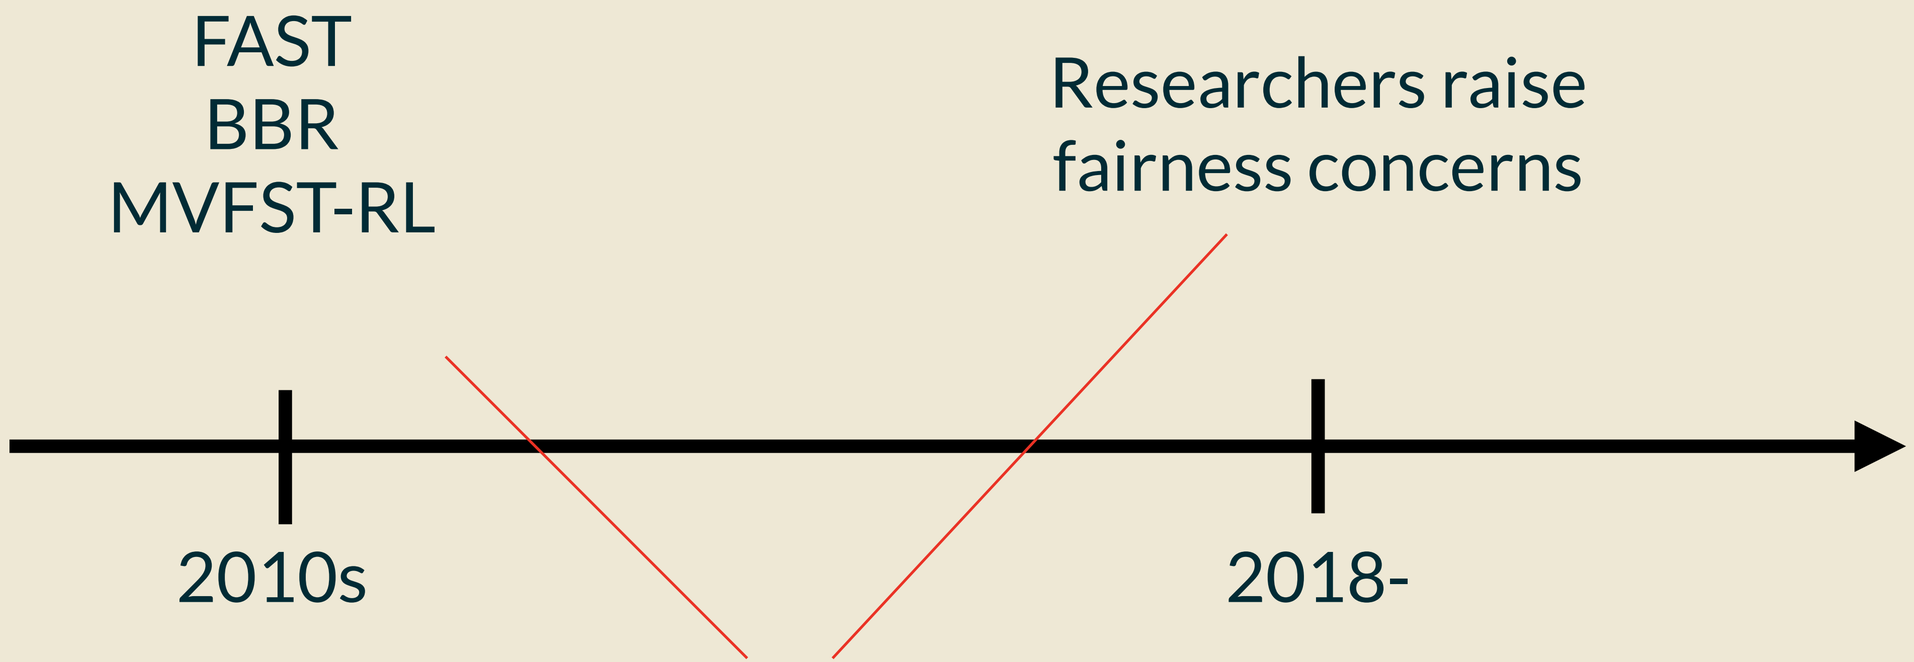

TCP Friendliness (TCPF)
concerns clearly did not prevent
new algorithms' deployment

Goal: Congestion Control Algorithm Independence (CCAI)

A stream's throughput should not depend on its choice of congestion control algorithm (CCA), relative to others' choices.

CCAI impact on congestion control design

Congestion control design concerns:

- Discovering available bandwidth
- Minimizing delay
- Ensuring fairness

CCAI impact on congestion control design

Congestion control design concerns:

- Discovering available bandwidth
- Minimizing delay
- ~~Ensuring fairness~~

CCAI allows CCA designs to focus on bandwidth-latency tradeoff

Simplest path to CCAI: isolation

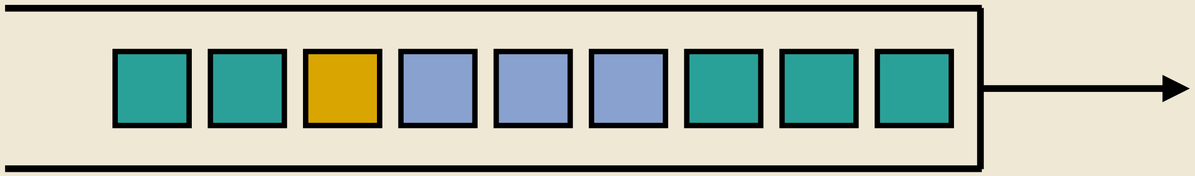

Traditional FIFO queues:
No isolation, no CCAI

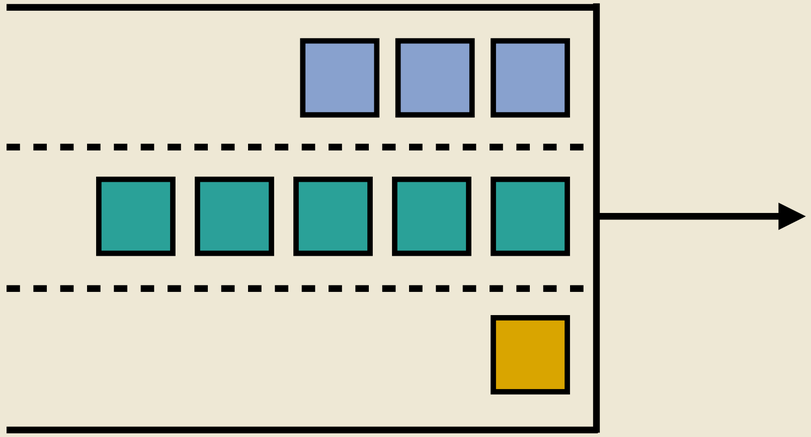

Fair queueing:
Flow isolation, CCAI

Fair queueing: not consistent with Internet economics

Internet runs on money, and flows are not economic units

"Flow Rate Fairness: Dismantling a Religion" -Bob Briscoe, IETF 2007

Internet economics: based on recursive access

Pay for **access**

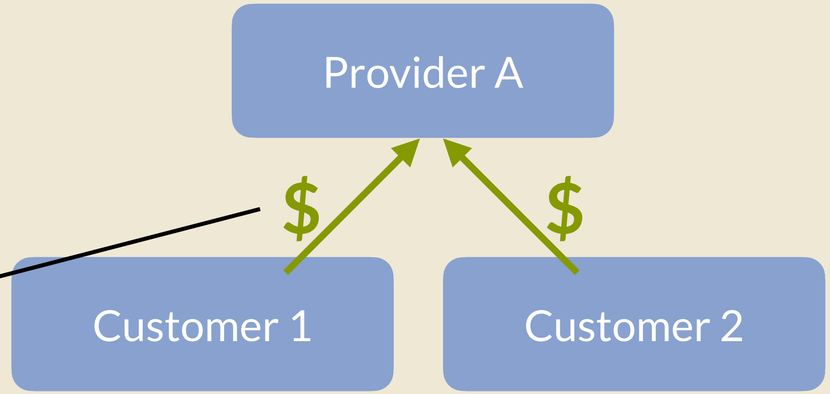

Internet economics: based on recursive access

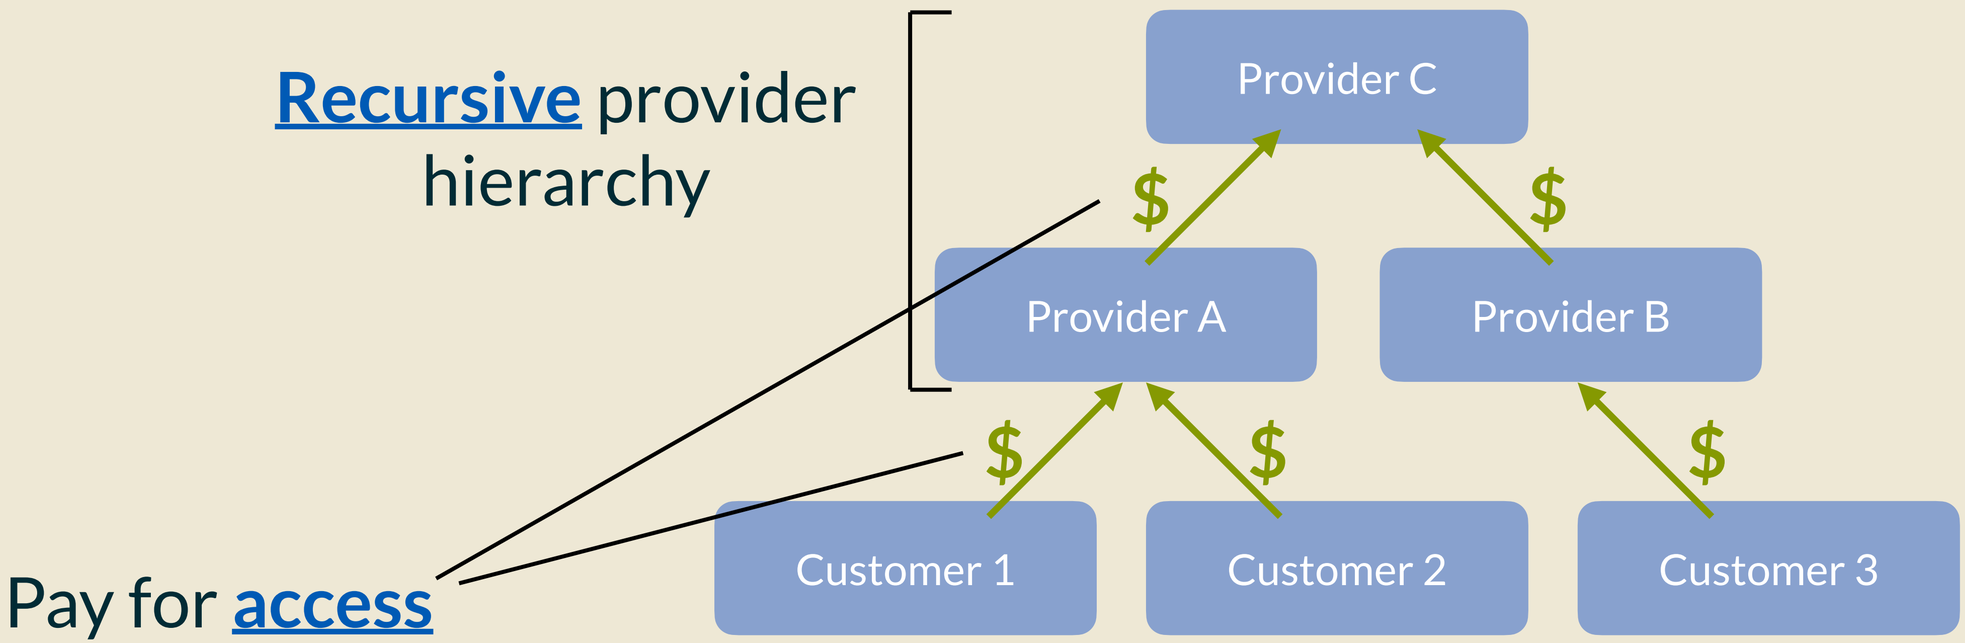

Principles for a bandwidth allocation framework

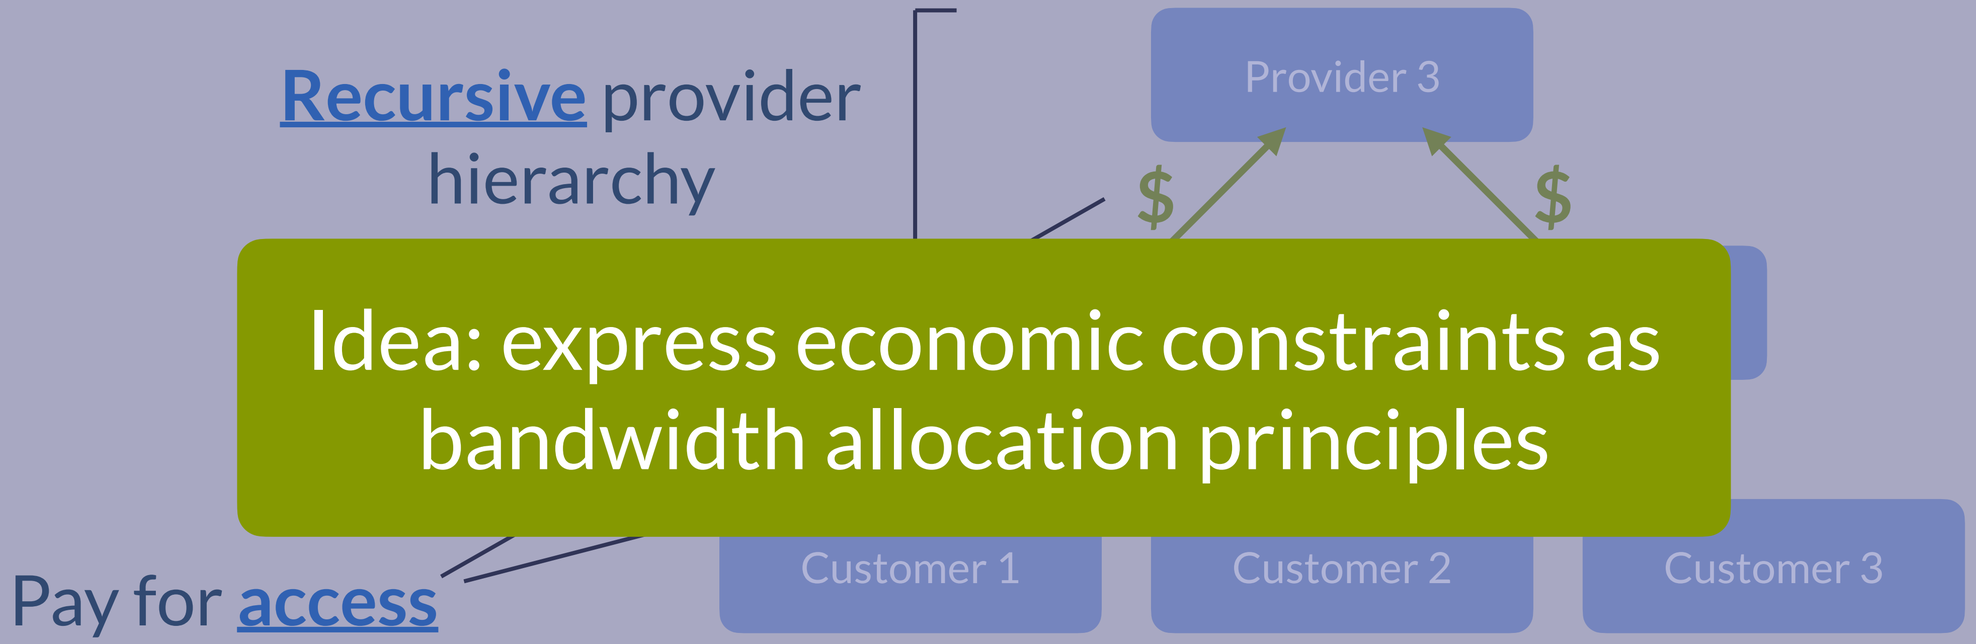

1: Relative rights

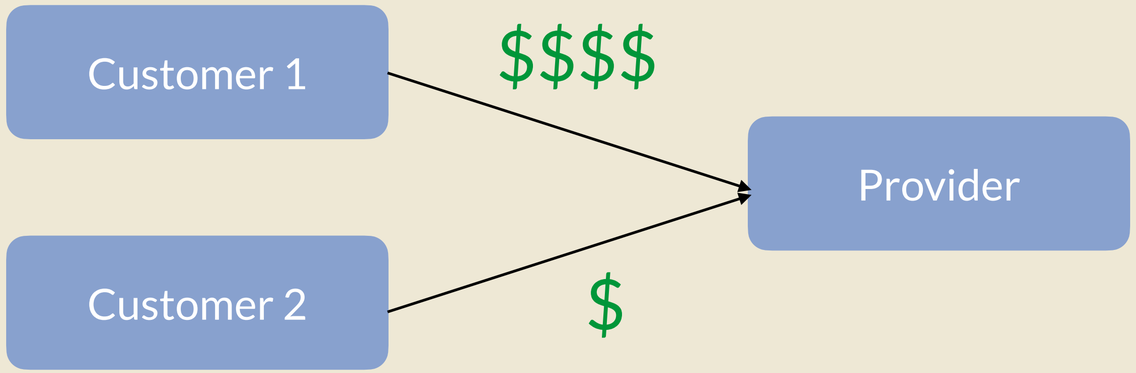

When there is congestion, allocate more bandwidth to higher-paying customers

2: Follow the money

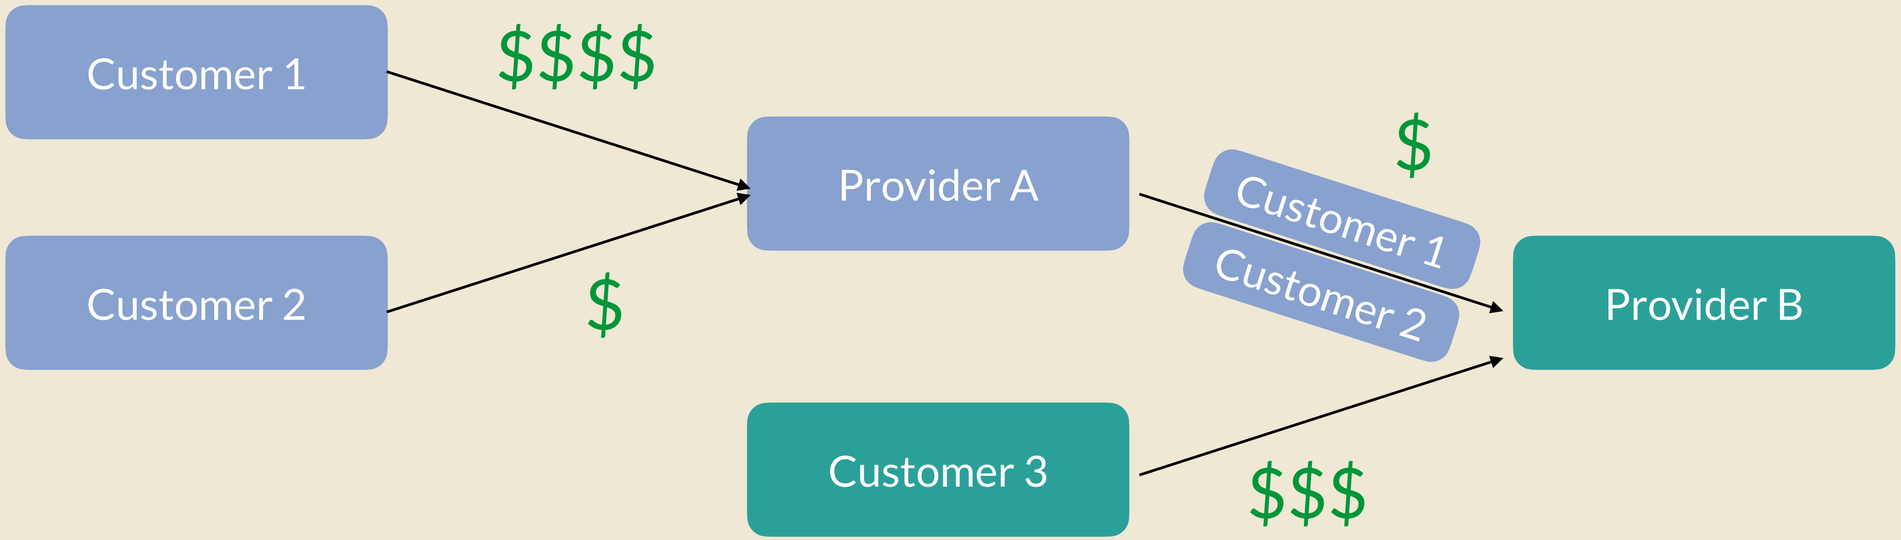

Relative rights should follow traffic through the network

2: Follow the money

1's traffic at B should depend on:
(1) 1's agreement with A
(2) A's agreement with B

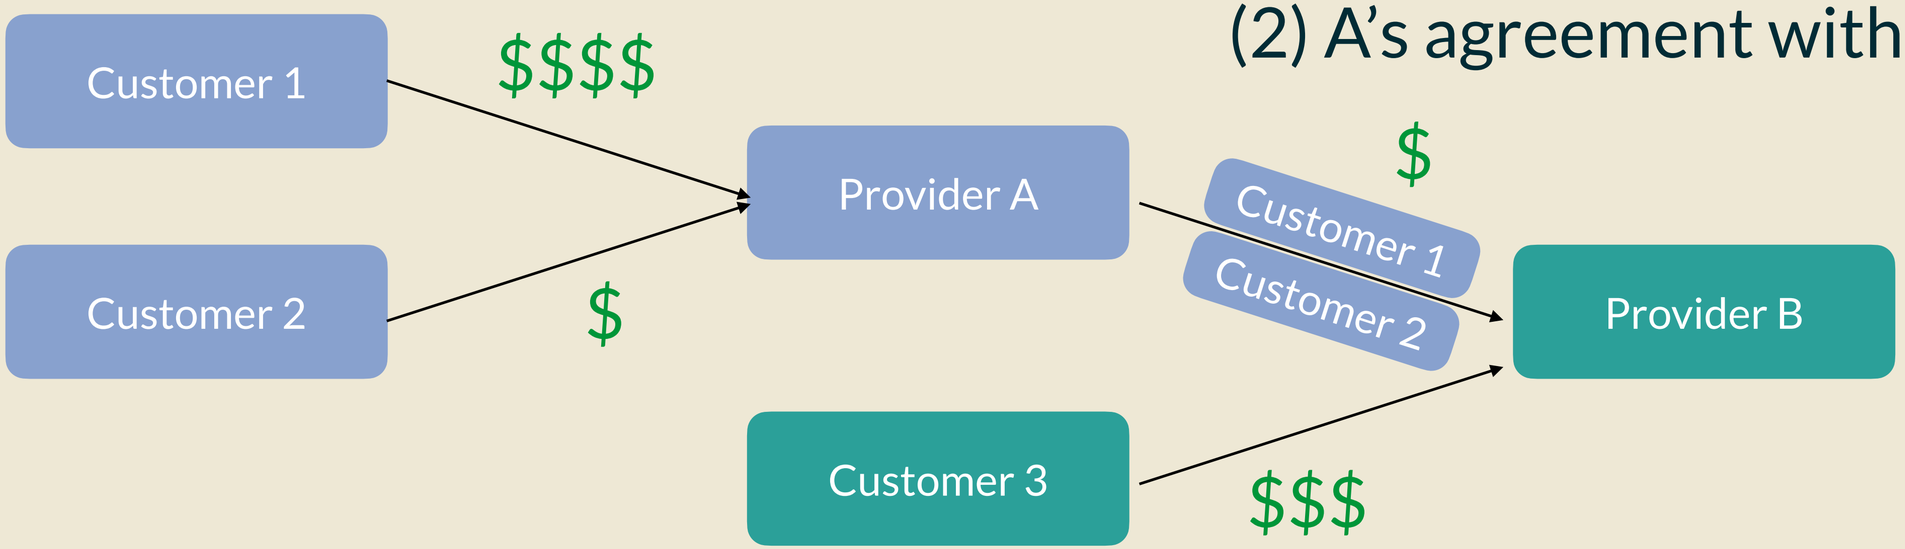

Relative rights should follow traffic
through the network

3: Endpoint control

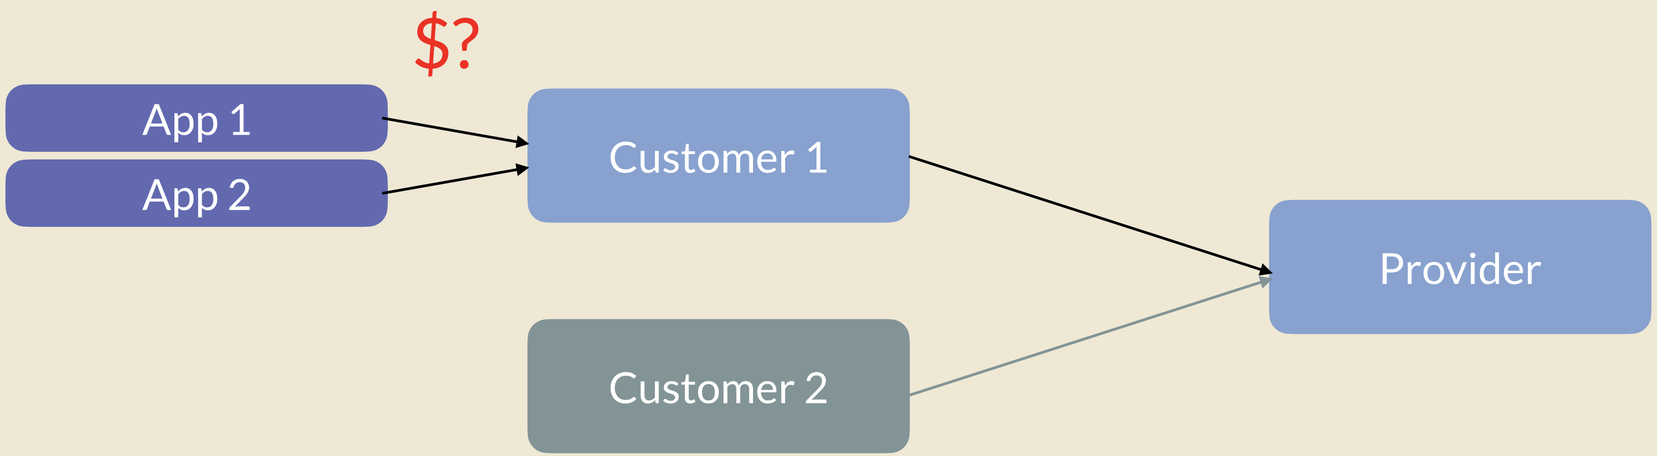

3: Endpoint control

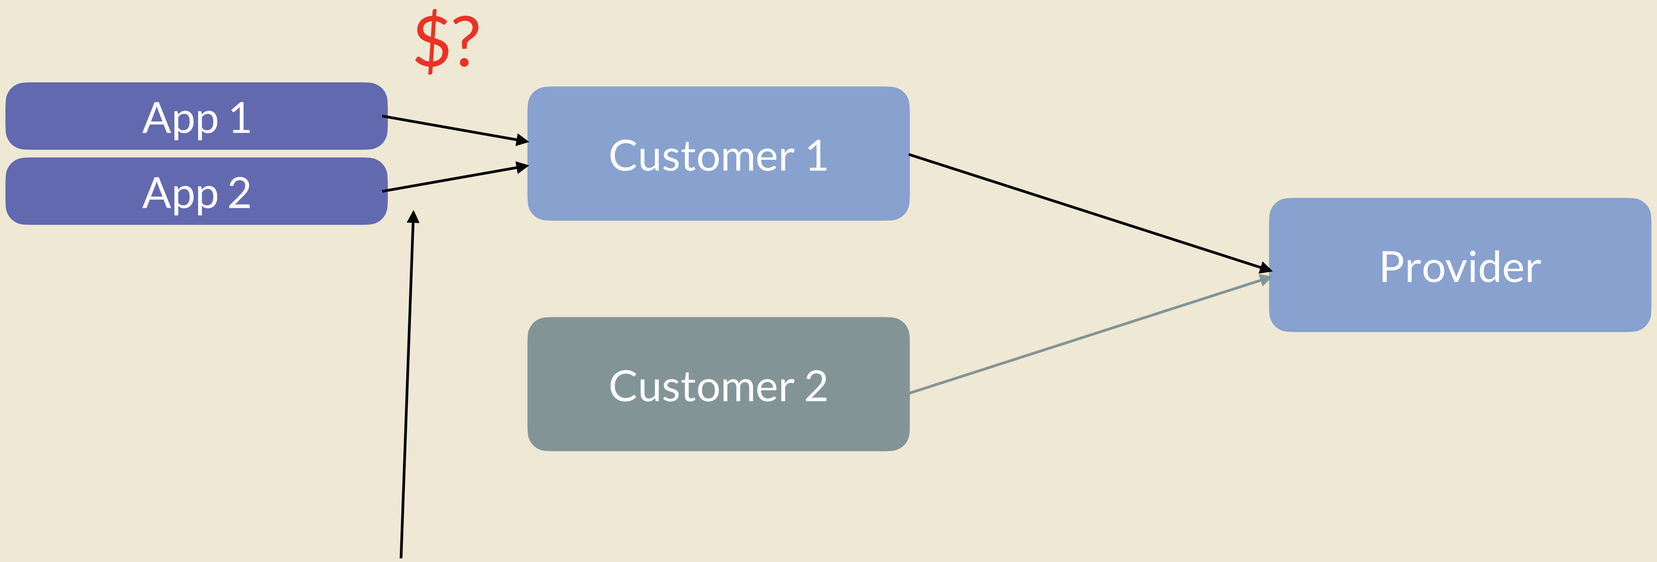

App 1 vs App 2:
Customer 1 decides at endpoints

See paper for discussion of how to achieve this

Our answer: RCS (Recursive Congestion Shares)

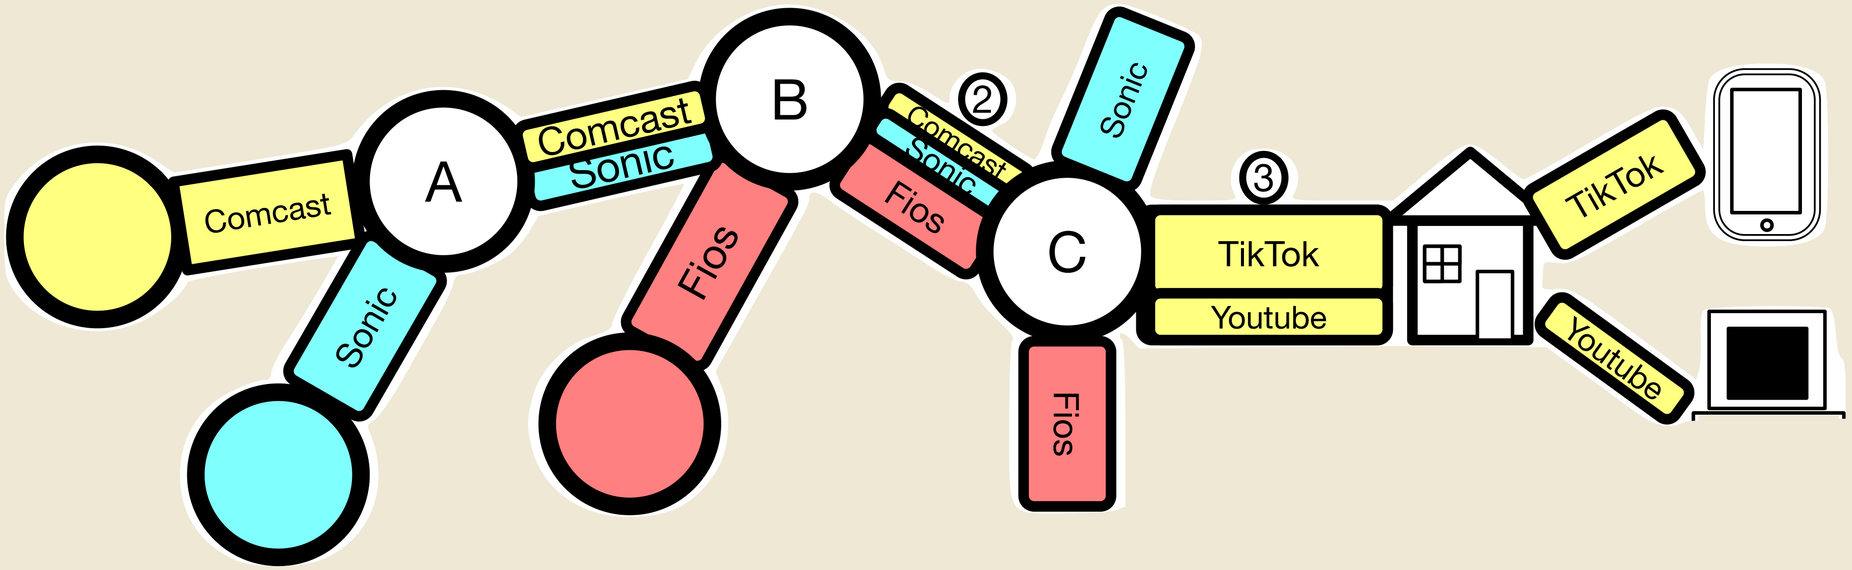

Mechanism: similar to HWFQ, but we set the weights carefully!

Evaluation: Does RCS achieve CCAI?

- Q1: Does RCS guarantee CCAI **in all cases**?
  - No, in adversarial cyclical-interaction cases (see paper for example)

Evaluation: Does RCS achieve CCAI?

- Q1: Does RCS guarantee CCAI in all cases?
  - No, in adversarial cyclical-interaction cases (see paper for example)
- Q2: Does RCS provide CCAI **in real topologies**? Yes (with high probability)
  - Model using game-theory framework

RCS achieves CCAI in most random + realistic topologies

Topology Type	CCAI	No CCAI (possible for CCAs to manipulate allocation)
Random	17,787 (95%)	936 (5%)
CAIDA-Sampled	2,897 (100%)	0 (0%)

Average BW Gain: 0.011%

Evaluation: Does RCS achieve CCAI?

- Q1: Does RCS guarantee CCAI in all cases?
  - No, in adversarial cyclical-interaction cases (see paper for example)
- Q2: Does RCS provide CCAI in real topologies? Yes (with high probability)
  - Expected gains (even in bad cases) don't merit effort
- Q3: Do these results apply in practice with real CCAs? Yes
  - Testbed emulation

Evaluation: Does RCS achieve CCAI?

- Q1: Does RCS guarantee CCAI in all cases?
  - No, in adversarial cyclical-interaction cases (see paper for example)
- Q2: Does RCS provide CCAI in real topologies? Yes (with high probability)
  - Expected gains (even in bad cases) don't merit effort
- Q3: Do these results apply in practice with real CCAs? Yes
  - Testbed emulation

RCS achieves CCAI with today's real CCAs

Do allocations remain
stable across topologies?

Allocations less stable

Allocations more stable

Setup: 10 CAIDA-sampled topologies,
10 random CCA assignments (Reno, Cubic, BBR) to each flow per topology

RCS achieves CCAI with today's real CCAs

Allocations do remain
stable across topologies

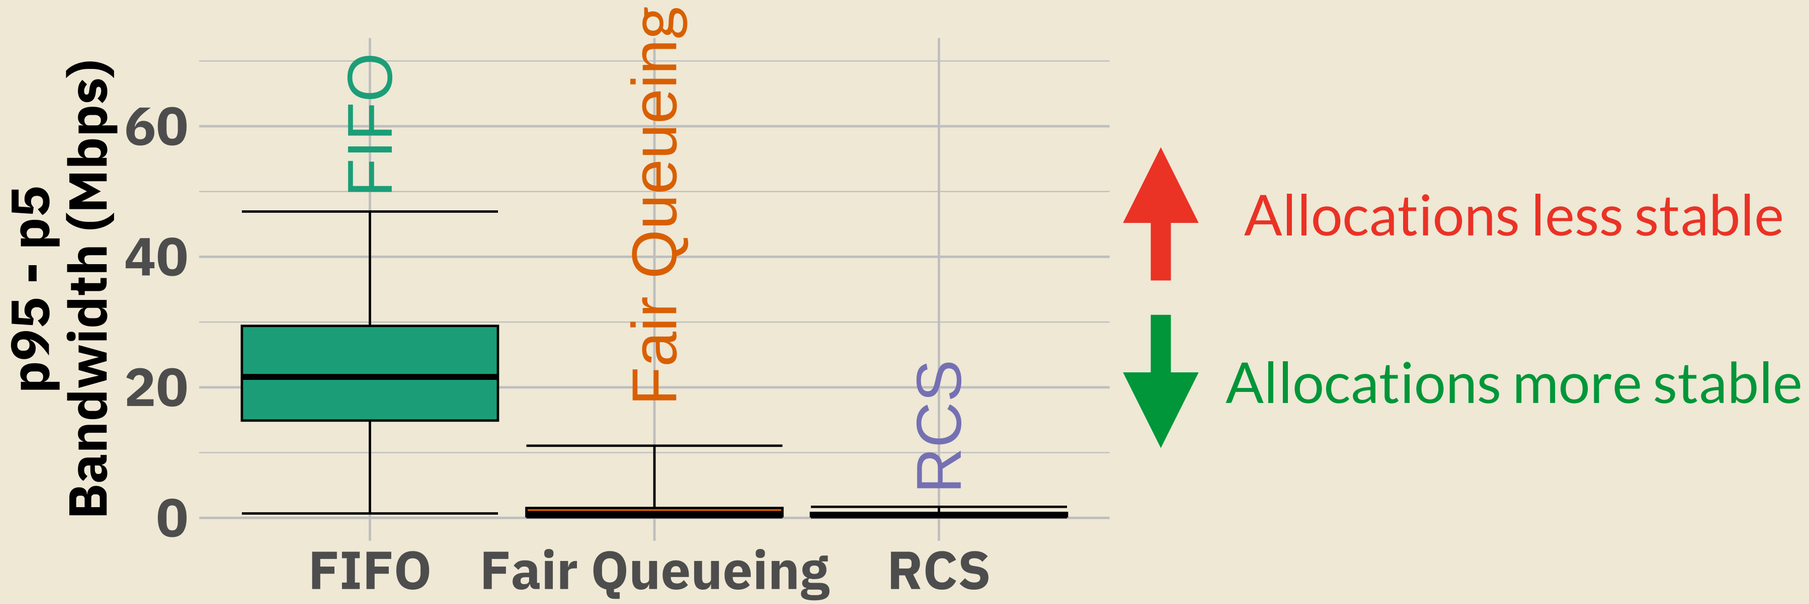

Setup: 10 CAIDA-sampled topologies,
10 random CCA assignments (Reno, Cubic, BBR) to each flow per topology

RCS achieves CCAI with today's real CCAs

Allocations do remain stable across topologies

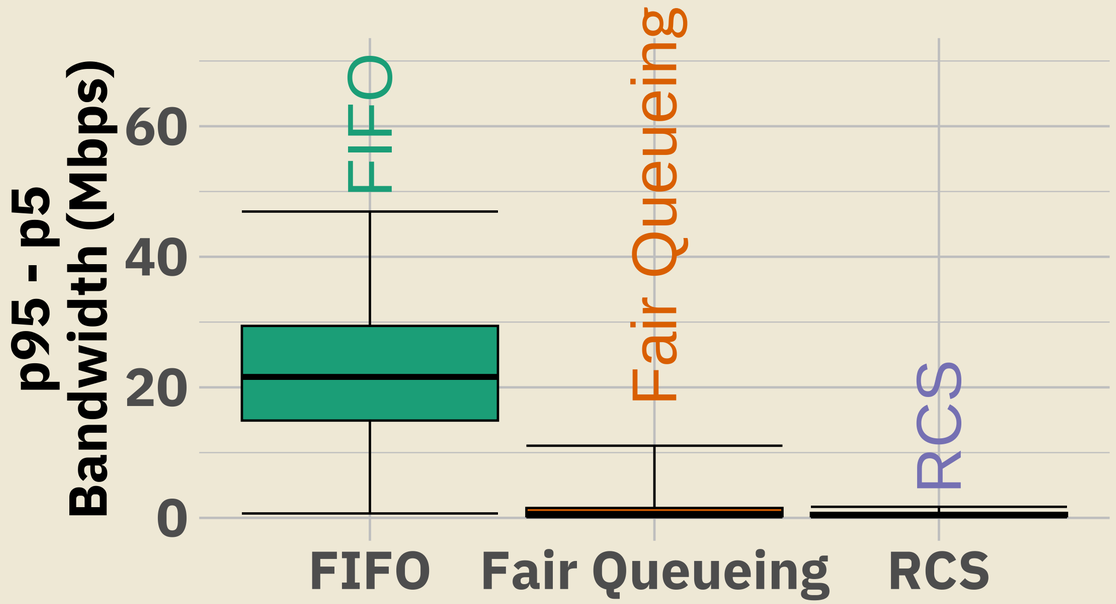

Are allocations close to what game-theoretic analysis predicted?

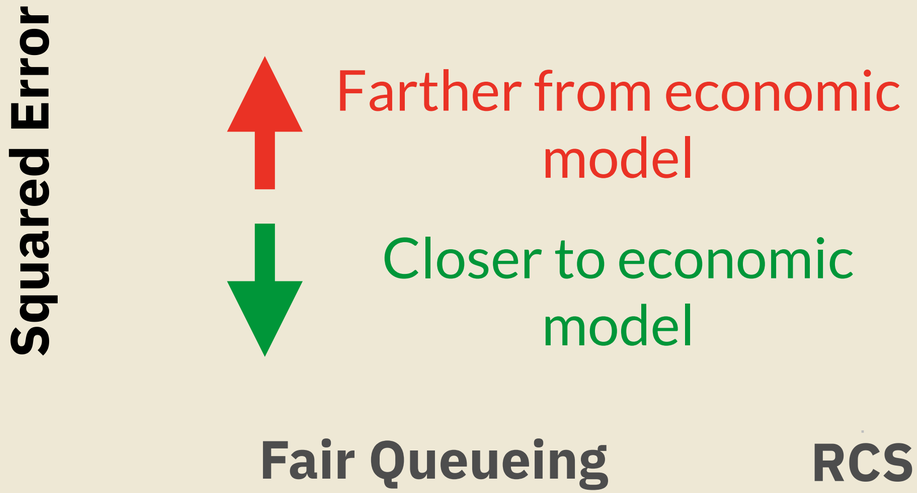

Setup: 10 CAIDA-sampled topologies,
10 random CCA assignments (Reno, Cubic, BBR) to each flow per topology

RCS achieves CCAI with today's real CCAs

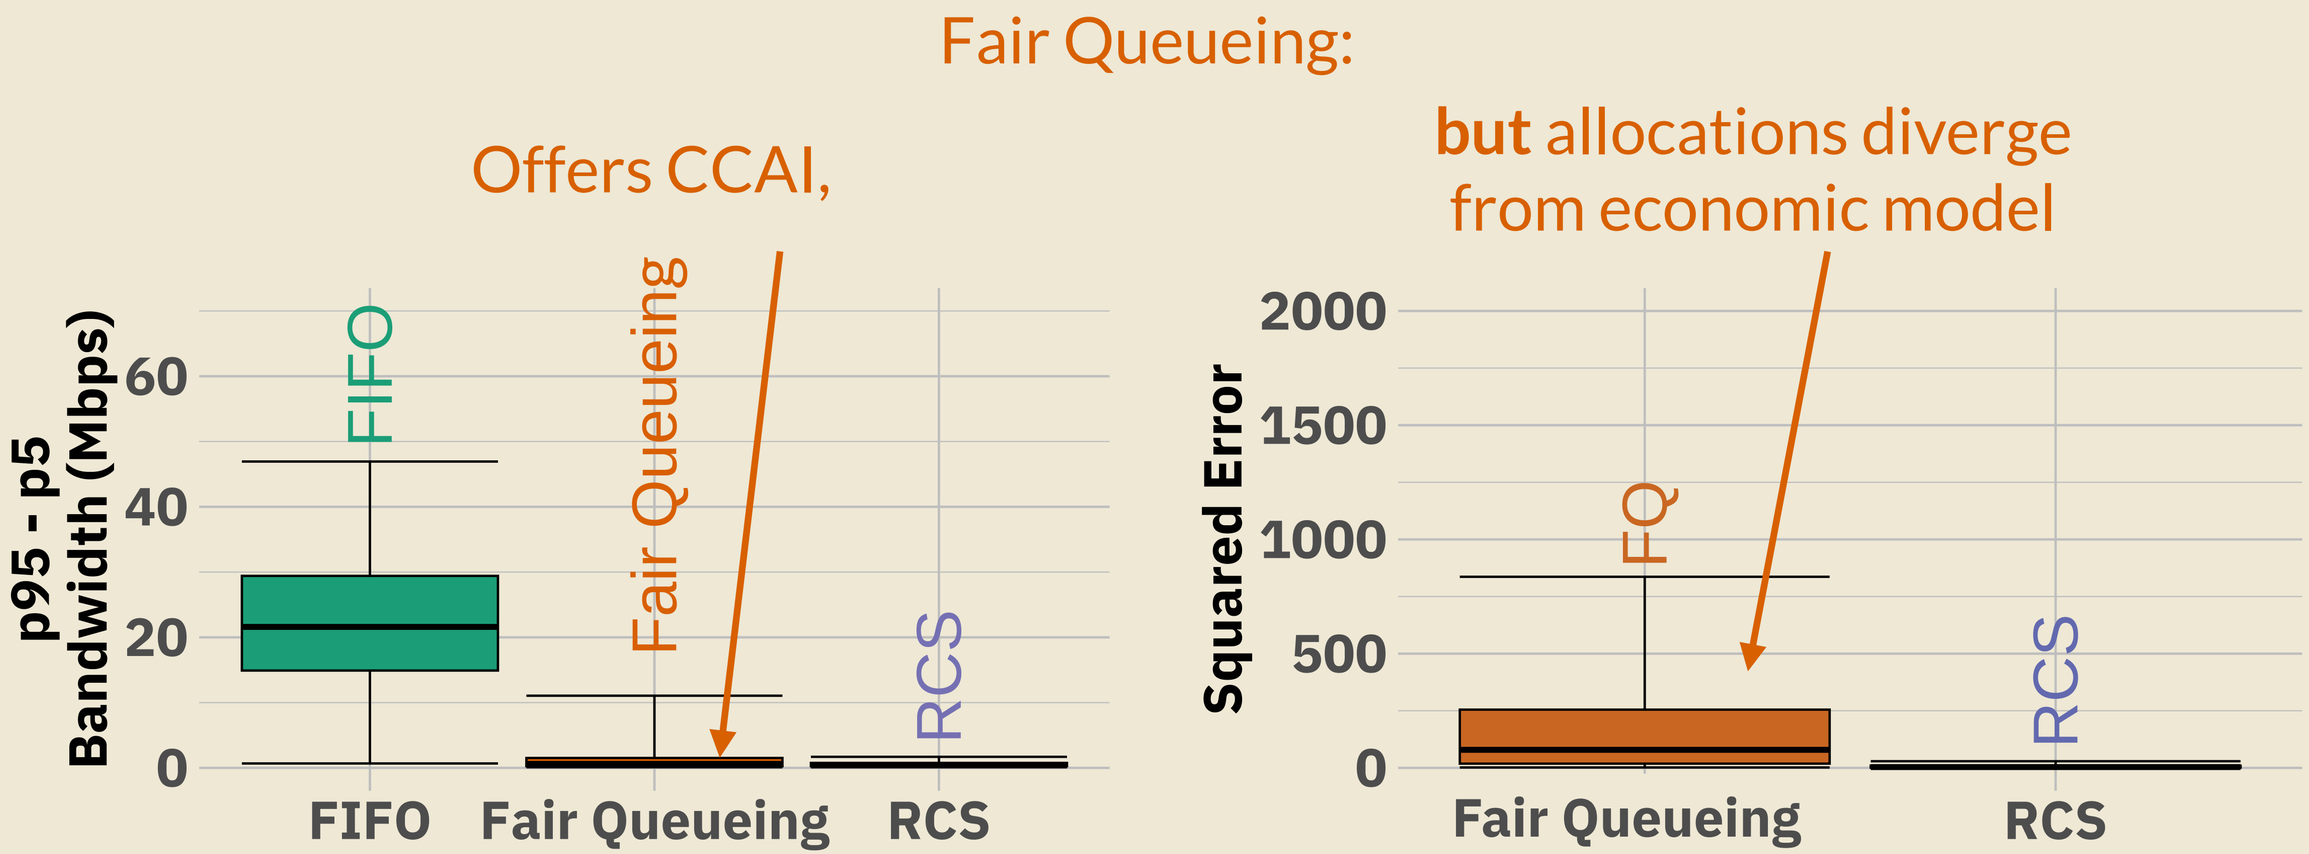

Setup: 10 CAIDA-sampled topologies,
10 random CCA assignments (Reno, Cubic, BBR) to each flow per topology

RCS achieves CCAI with today's real CCAs

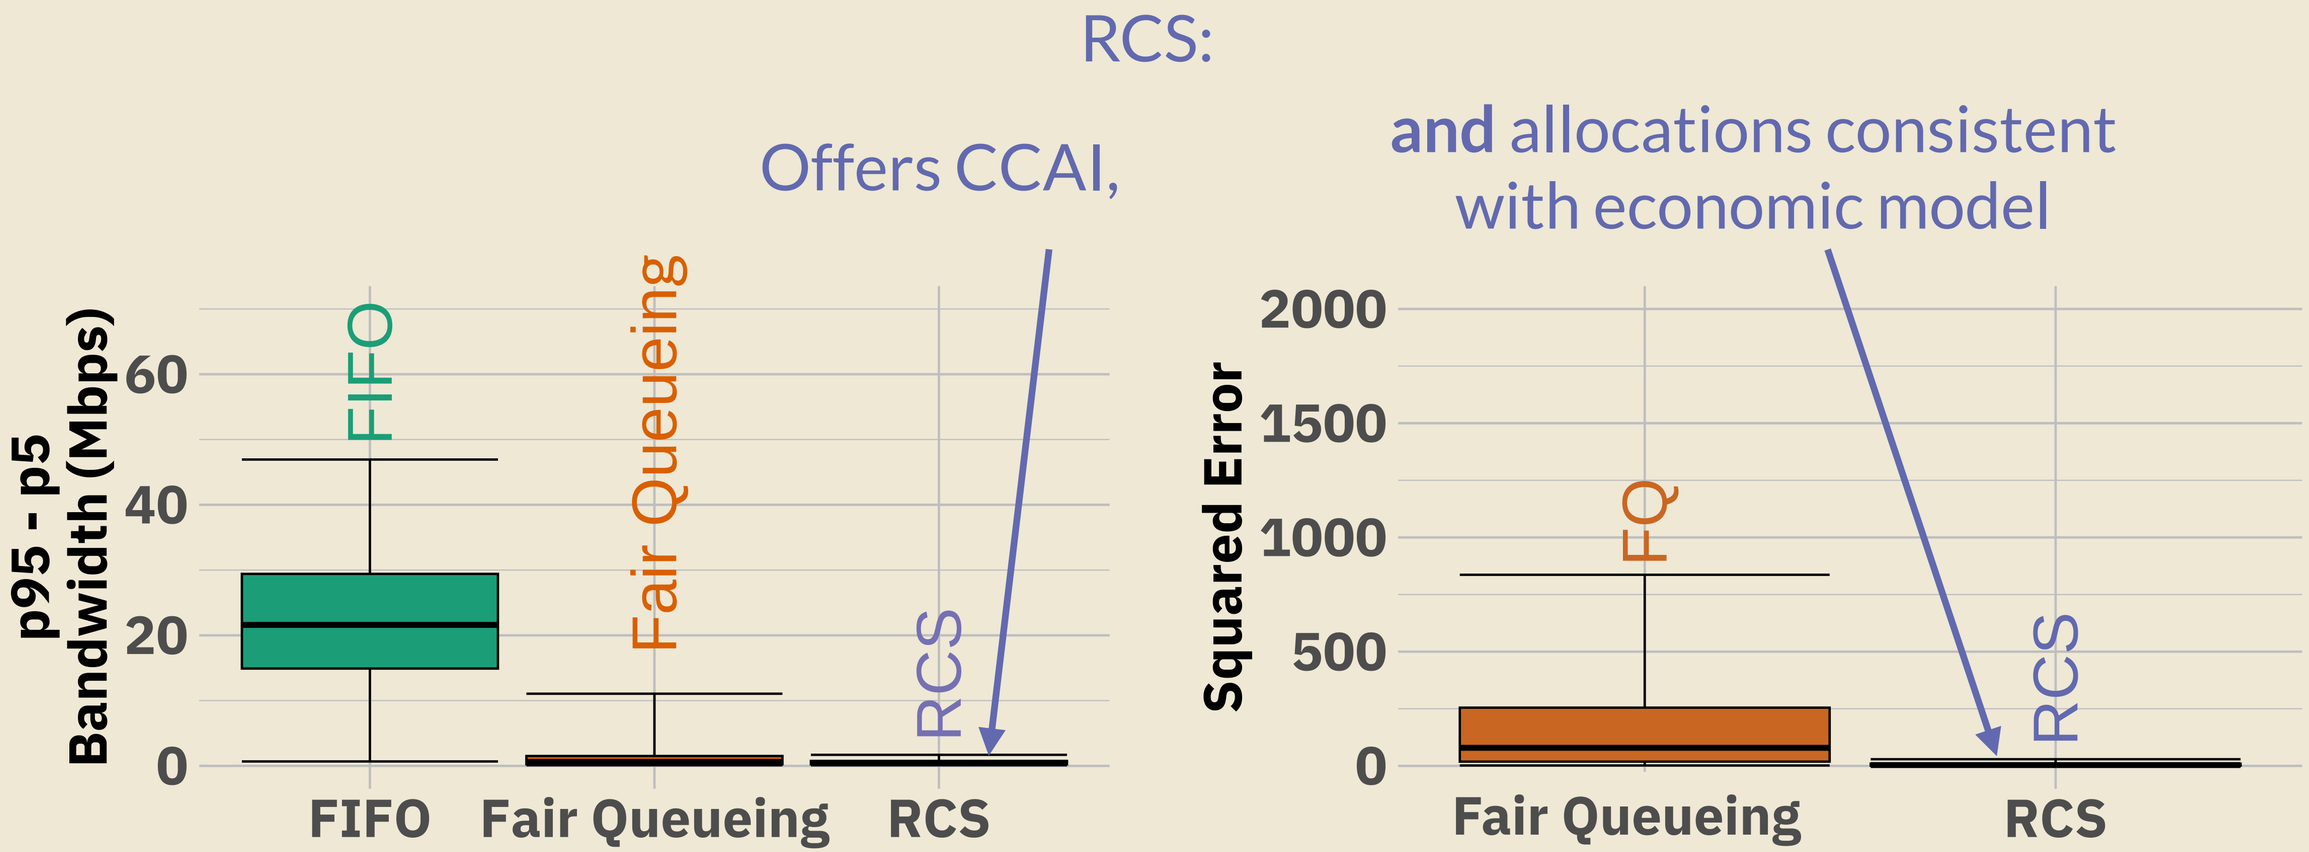

Setup: 10 CAIDA-sampled topologies,
10 random CCA assignments (Reno, Cubic, BBR) to each flow per topology

Discussion

- Implementation concerns
  - See paper: Implementing RCS is possible on modern switch hardware
- Net neutrality concerns
  - See paper: RCS conforms to widely accepted net neutrality definitions
- Adoption incentives
  - RCS offers incremental benefits scaling with deployment

RCS: Bandwidth Allocation Framework for Congestion Control Algorithm Independence

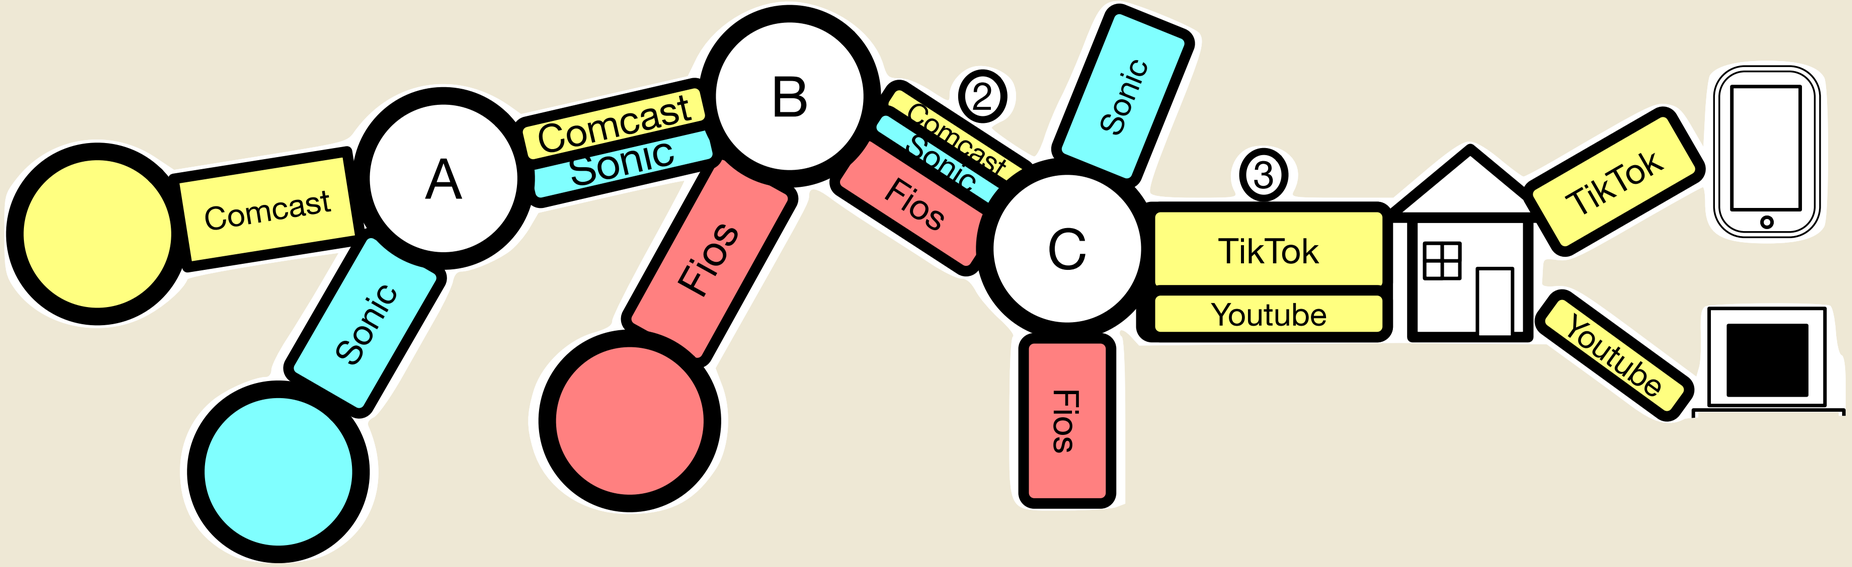

Akshay Narayan - akshayn@brown.edu