HOW I LEARNED TO STOP WORRYING ABOUT CCA CONTENTION

Lloyd Brown, Yash Kothari, Akshay Narayan, Arvind Krishnamurthy, Aurojit Panda, Justine Sherry, Scott Shenker

HotNets 2023

Berkeley UNIVERSITY OF CALIFORNIA
Carnegie Mellon University
W UNIVERSITY of WASHINGTON
NYU

Does CCA contention typically determine a flow's bandwidth allocation in the Internet today?

Does CCA contention typically determine a flow's bandwidth allocation in the Internet today?

Congestion: Offered load > capacity

CCA contention:
Higher sending rate leads to
higher bandwidth allocation
+ lower allocation for other flows

Conventional CCA model of contention

Internet CCA model:
Zero-sum bandwidth competition

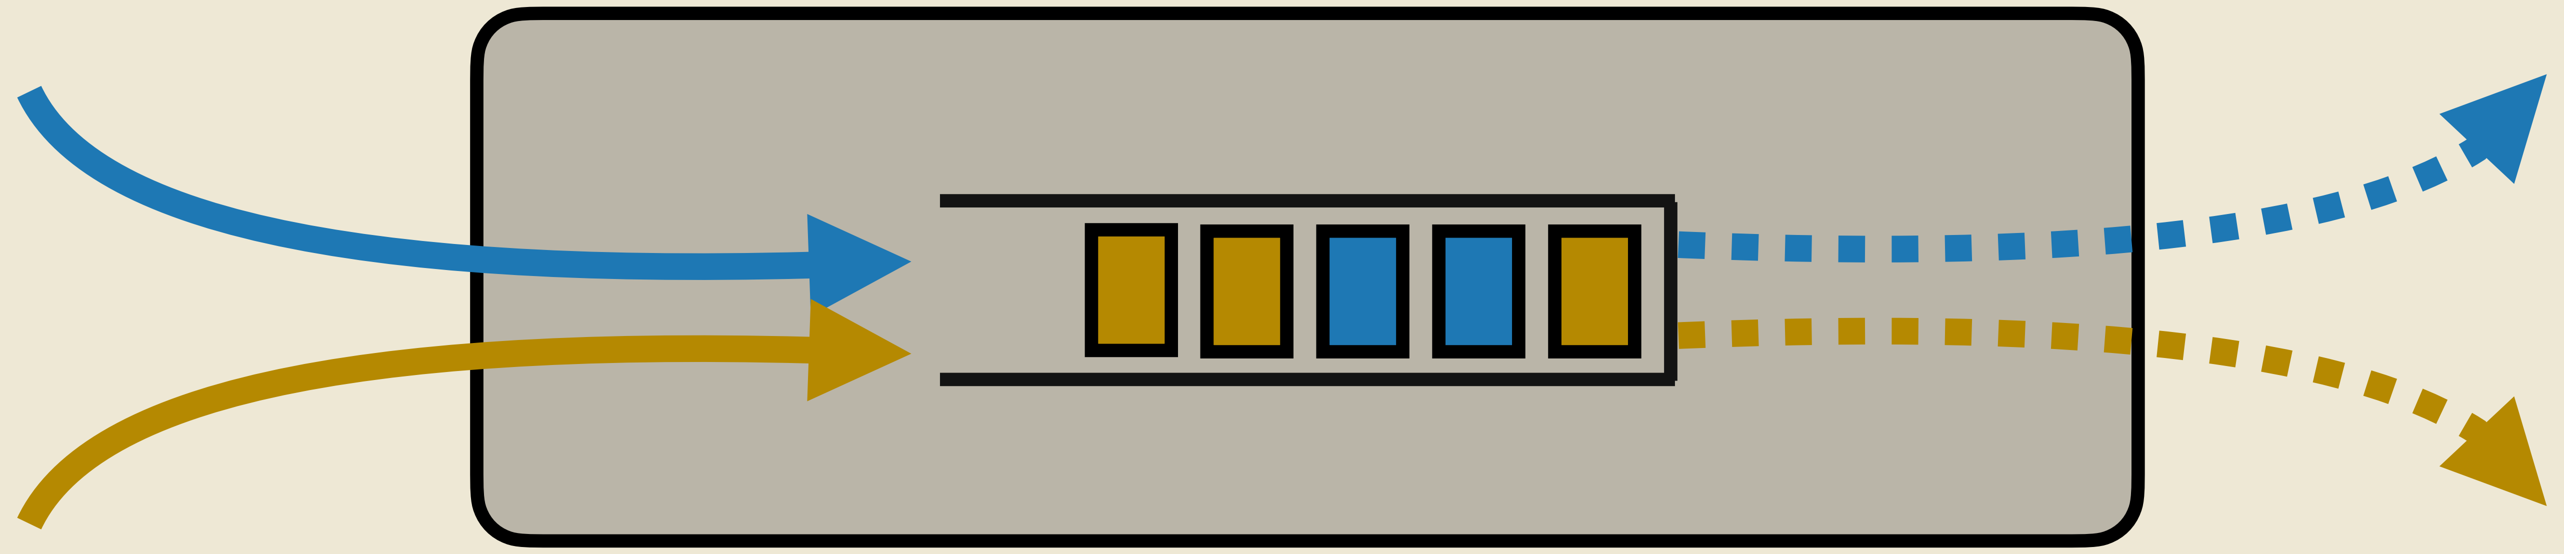

Traditional Protocols:
Prescriptive AI-MD
“TCP-Friendliness”

→

Modern Algorithms:
Low Delay
High Throughput

Expectation: Contention between modern CC algorithms determines their bandwidth allocations

CCA Contention properties: longstanding analysis target!

SIGCOMM 2022
Best Student Paper

Starvation in End-to-End Congestion Control

Venkat Arun, Mohammad Alizadeh, Hari Balakrishnan
{venkatar, alizadeh, hari}@csail.mit.edu
MIT Computer Science and Artificial Intelligence Lab, Cambridge, MA

ABSTRACT

To overcome weaknesses in traditional loss-based congestion control algorithms (CCAs), researchers have developed and deployed several delay-bounding CCAs that achieve high utilization without bloating delays (e.g., Vegas, FAST, BBR, PCC, Copa, etc.). When run on a path with a fixed bottleneck rate, these CCAs converge to a small delay range in equilibrium. This paper proves a surprising result: although designed to achieve reasonable inter-flow fairness, current methods to develop delay-bounding CCAs cannot always avoid starvation, an extreme form of unfairness. Starvation may occur when such a CCA runs on paths where non-congestive network delay variations due to real-world factors such as ACK aggregation and end-host scheduling exceed double the delay range that the CCA converges to in equilibrium. We provide experimental evidence for this result for BBR, PCC Vivace, and Copa ...

1 INTRODUCTION

With the rise of interactive and real-time applications, the ...

2019 IRTF
Applied Networking
Research Prize

Beyond Jain's Fairness Index: Setting the Bar For The Deployment of Congestion Control Algorithms

Ranysha Ware, Carnegie Mellon University, rware@cs.cmu.edu
Matthew K. Mukerjee, Nefeli Networks, mukerjee@nefeli.io
Srinivasan Seshan, Carnegie Mellon University, srini@cs.cmu.edu
Justine Sherry, Carnegie Mellon University, sherry@cs.cmu.edu

ABSTRACT

The Internet community faces an explosion in new congestion control algorithms such as Copa, Sprout, PCC, and BBR. In this paper, we discuss considerations for deploying new algorithms on the Internet. While past efforts have focused on achieving 'fairness' or 'friendliness' between new algorithms and deployed algorithms, we instead advocate for an approach centered on quantifying and limiting harm caused by the new algorithm on the status quo. We argue that a harm-based approach is more practical, more future proof, and handles a wider range of quality metrics than traditional notions of fairness and friendliness.

outcomes between a collection of flows, all using some CCA α, we need to analyze what happens when a new CCA α is deployed on a network with flows using some legacy CCA β. Is α's impact on the status quo acceptable?

Our community has traditionally analyzed inter-CCA competition in two ways, which we refer to as 'fairness' and 'mimicry.' While both approaches are insightful, we argue that neither is a sound basis for a deployment threshold.

A throughput allocation is fair if it maximizes every users utility function given limited link capacity [21]. A end-host CCA, typically defines users as flows, aiming to maximize util- ...

Congestion Control Safety via Comparative Statics

Pratiksha Thaker, School of Computer Science, Carnegie Mellon University, USA, prthaker@cmu.edu
Matei Zaharia, Department of Computer Science, Stanford University, USA, matei@cs.stanford.edu
Tatsunori Hashimoto, Department of Computer Science, Stanford University, USA, thashim@stanford.edu

***Abstract*—When congestion control algorithms compete on shared links, unfair outcomes can result, especially between algorithms that aim to prioritize different objectives. For example, a throughput-maximizing application could make the link completely unusable for a latency-sensitive application. In order to ...**

the behavior of these algorithms through measurement after deployment [6], [9].

In this paper, we present a framework that allows protocol designers to reason directly about the performance objectives that applications aim to optimize, rather than relying on experiments to validate heuristic algorithms post-hoc. Taking the perspective ...

Axiomatizing Congestion Control

DORON ZARCHY, Hebrew University of Jerusalem
RADHIKA MITTAL, University of Illinois at Urbana-Champaign
MICHAEL SCHAPIRA, Hebrew University of Jerusalem
SCOTT SHENKER, UC Berkeley, ICSI

The overwhelmingly large design space of congestion control protocols, along with the increasingly diverse range of application environments, makes evaluating such protocols a daunting task. Simulation and experiments are very helpful in evaluating the performance of designs in *specific* contexts, but give limited insight into the more general properties of these schemes and provide no information about the *inherent* limits of congestion control designs (such as, which properties are simultaneously achievable and which are mutually exclusive). In contrast, traditional theoretical approaches are typically focused on the design of protocols that achieve to specific, predetermined objectives (e.g., network utility maximization), or the analysis of specific protocols (e.g., from control-theoretic perspectives), as opposed to the inherent tensions/derivations between desired properties.

To complement today's prevalent experimental and theoretical approaches, we put forth a novel principled framework for reasoning about congestion control protocols, which is inspired by the axiomatic approach from social choice theory and game theory. We consider several natural requirements ("axioms") from congestion control protocols – e.g., efficient resource-utilization, loss-avoidance, fairness, stability, and TCP-friendliness – and investigate which combinations of these can be achieved within a single design. Thus, our framework allows us to investigate the fundamental tradeoffs between desiderata, and to identify where existing and new congestion control architectures fit within the space of possible outcomes.

On the Future of Congestion Control for the Public Internet

Lloyd Brown[1] Ganesh Ananthanarayanan[2] Ethan Katz-Bassett[3] Arvind Krishnamurthy[4]
Sylvia Ratnasamy[1] Michael Schapira[5] Scott Shenker[1,6]
[1] UC Berkeley [2] Microsoft Research [3] Columbia University [4] University of Washington
[5] Hebrew University of Jerusalem [6] ICSI

Abstract

The conventional wisdom requires that all congestion control algorithms deployed on the public Internet be TCP-friendly. If universally obeyed, this requirement would greatly constrain the future of such congestion control algorithms. If partially ignored, as is increasingly likely, then there could be significant inequities in the bandwidth received by different flows. To avoid this dilemma, we propose an alternative to the TCP-friendly paradigm that can accommodate innovation, is consistent with the Internet's current economic model, and is feasible to deploy given current usage trends.

includes all endpoint CCAs that are deployed on hosts in private networks but which communicate with other endpoints that are reached by crossing the public Internet, because such flows may interact with flows using CCAs chosen by others. In this open setting, where there is no central control over which CCAs are deployed, we consider the question of how to tolerate diversity in CCAs, which is necessary for enabling congestion control innovations to be freely deployed.

To review, the modern era of congestion control started with the seminal works of Jacobson [25, 26] and Ramakrishnan and Jain [27], and has culminated with TCP Cubic [28] (the default in Linux) and ...

Does CCA contention typically determine a flow's bandwidth allocation in the Internet today?

Our Hypothesis: Not typically

- Operators' capacity planning usually prevents congestion
- Most flows are app-limited or operator-throttled
- We need active measurement to settle the question

Later: What are the implications?

Operators’ capacity planning usually prevents congestion

Flow Perspective

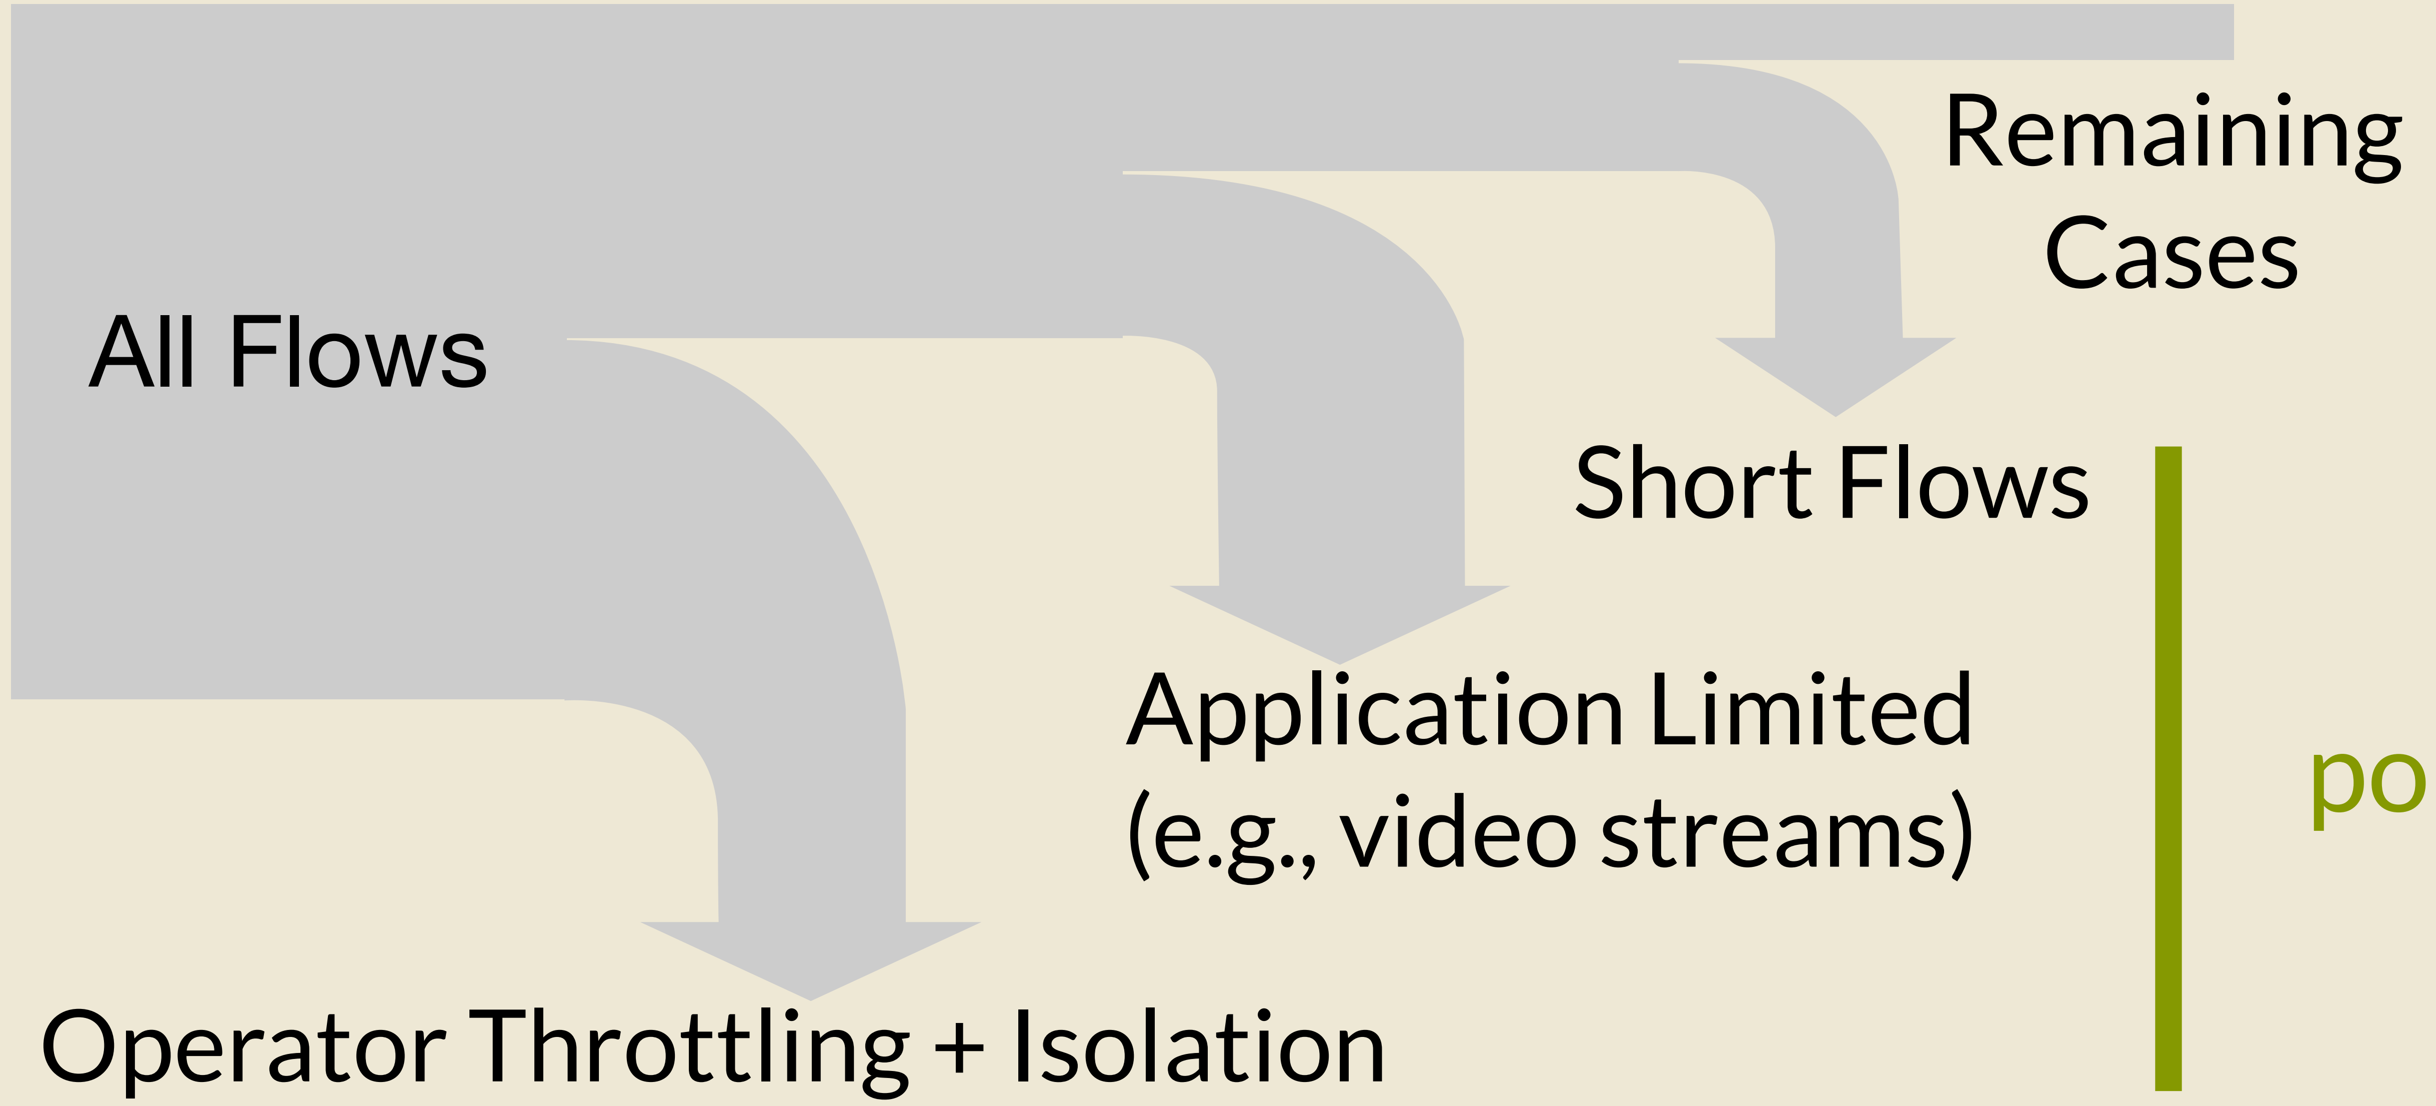

Measurement data points to these cases being common

Most flows either throttled or have bandwidth demand satisfied

Operators' capacity planning usually prevents congestion

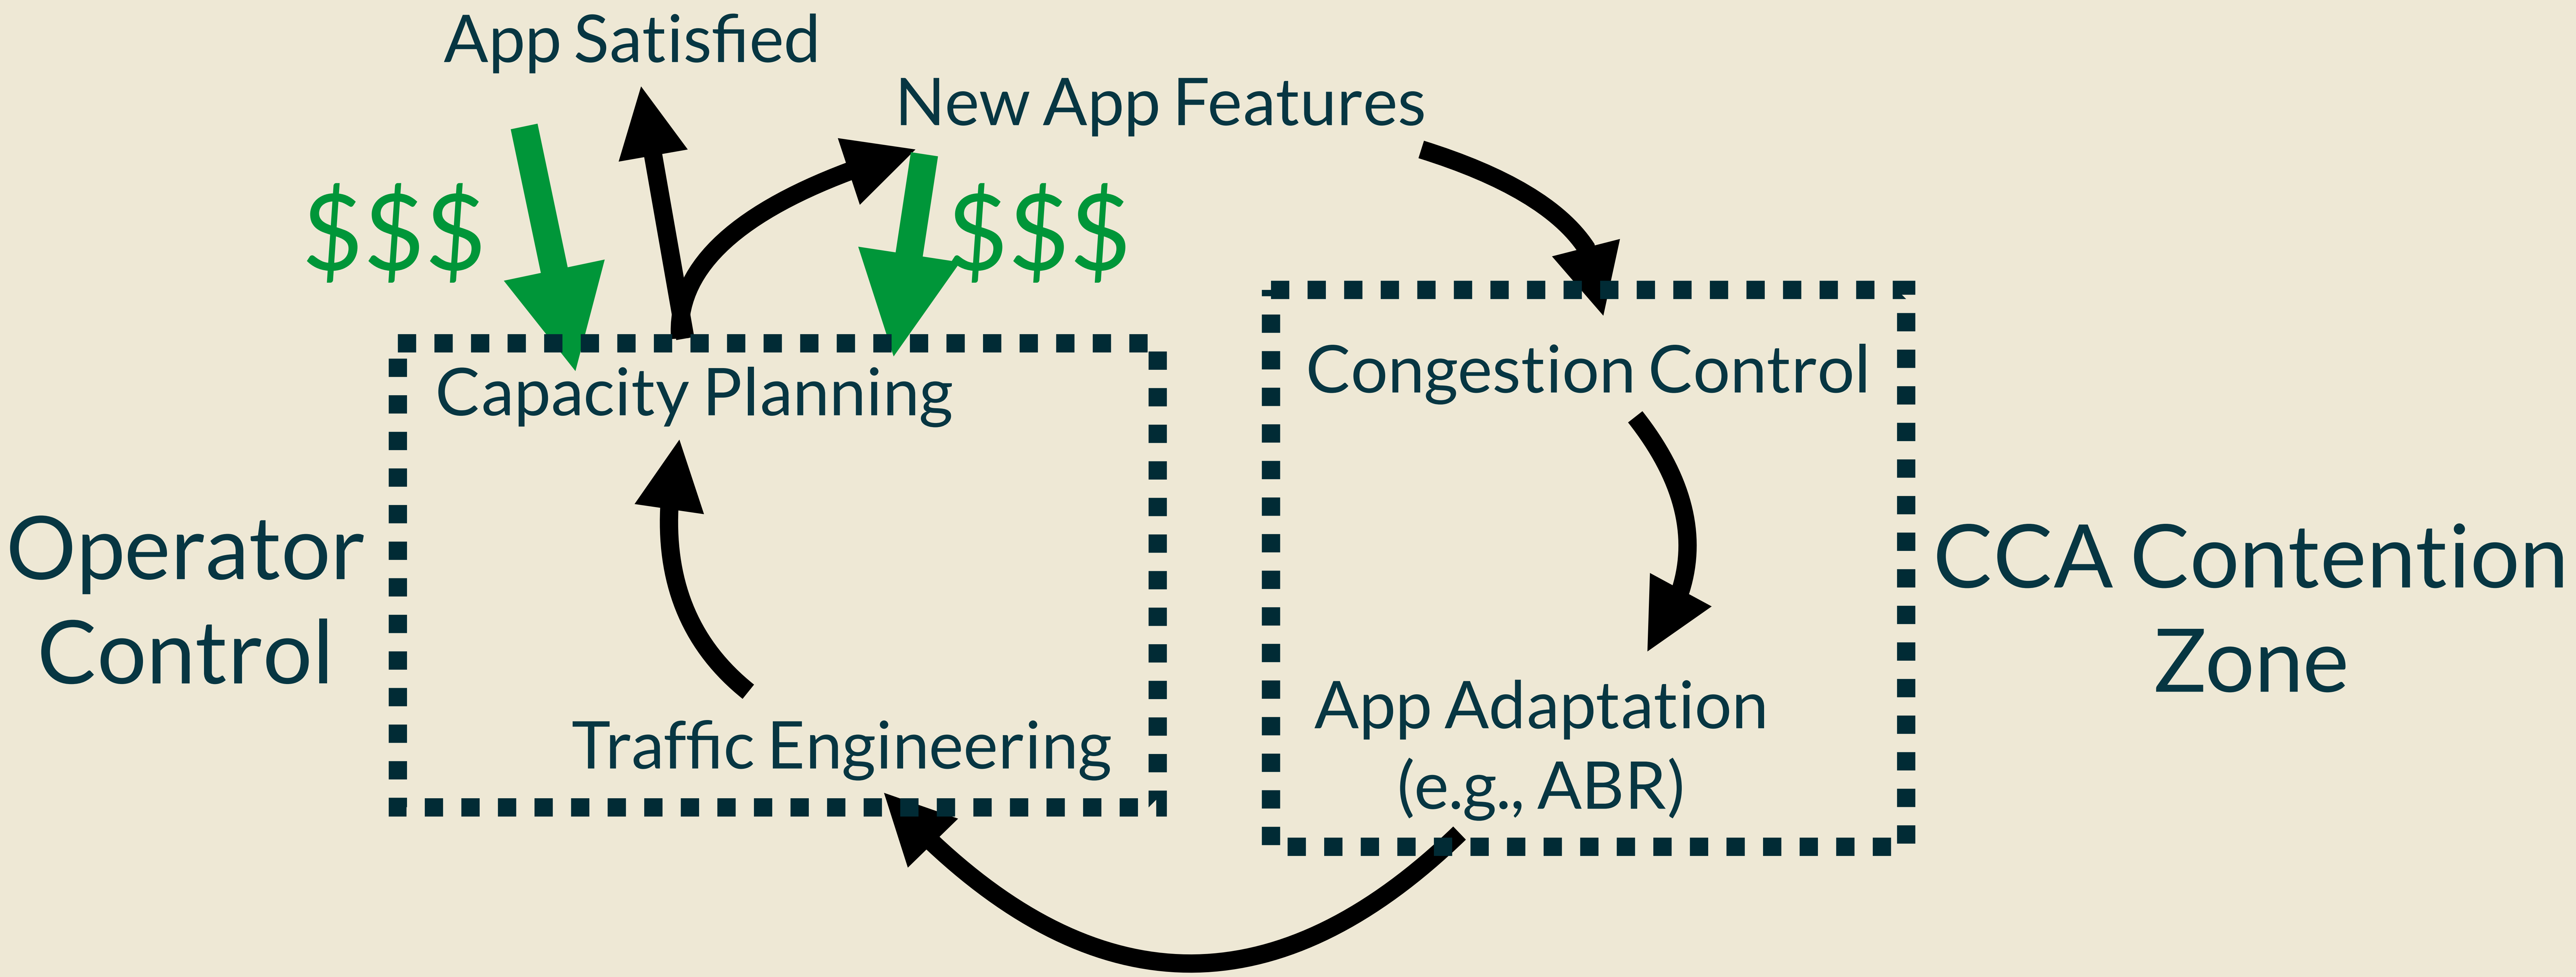

Operators care about money, not bandwidth contention:

If many flows don't have bandwidth demands satisfied, operators upgrade capacity

Current trends point to reduced contention*

Access links:
Yang et al. (2022): In 64% of cases, Wi-Fi is as much of a bottleneck as home users' access links

↓

Remaining cases:
Can use well-known fair queueing and isolation mechanisms

Low-Bandwidth Scenarios:
Possible for commonly application-limited traffic (e.g., video streaming) to contend for limited bandwidth

↓

Wireless (cellular/satellite) access:
isolation already common

*Further study required

Does CCA contention typically determine a flow's bandwidth allocation in the Internet today?

Our Hypothesis: Not typically

- Operators' capacity planning usually prevents congestion
- Most flows are app-limited or operator-throttled
- We need active measurement to settle the question

Passively gathered M-Lab NDT Data:
does not indicate contention, but not conclusive either way
(see paper)

Active measurements can directly measure CCA contention

Any approach using passive measurements must infer CCA contention, so is likely to be inconclusive

Idea: Adapt Nimbus (SIGCOMM 2022), which actively determines cross-traffic's aggressiveness

See paper for details

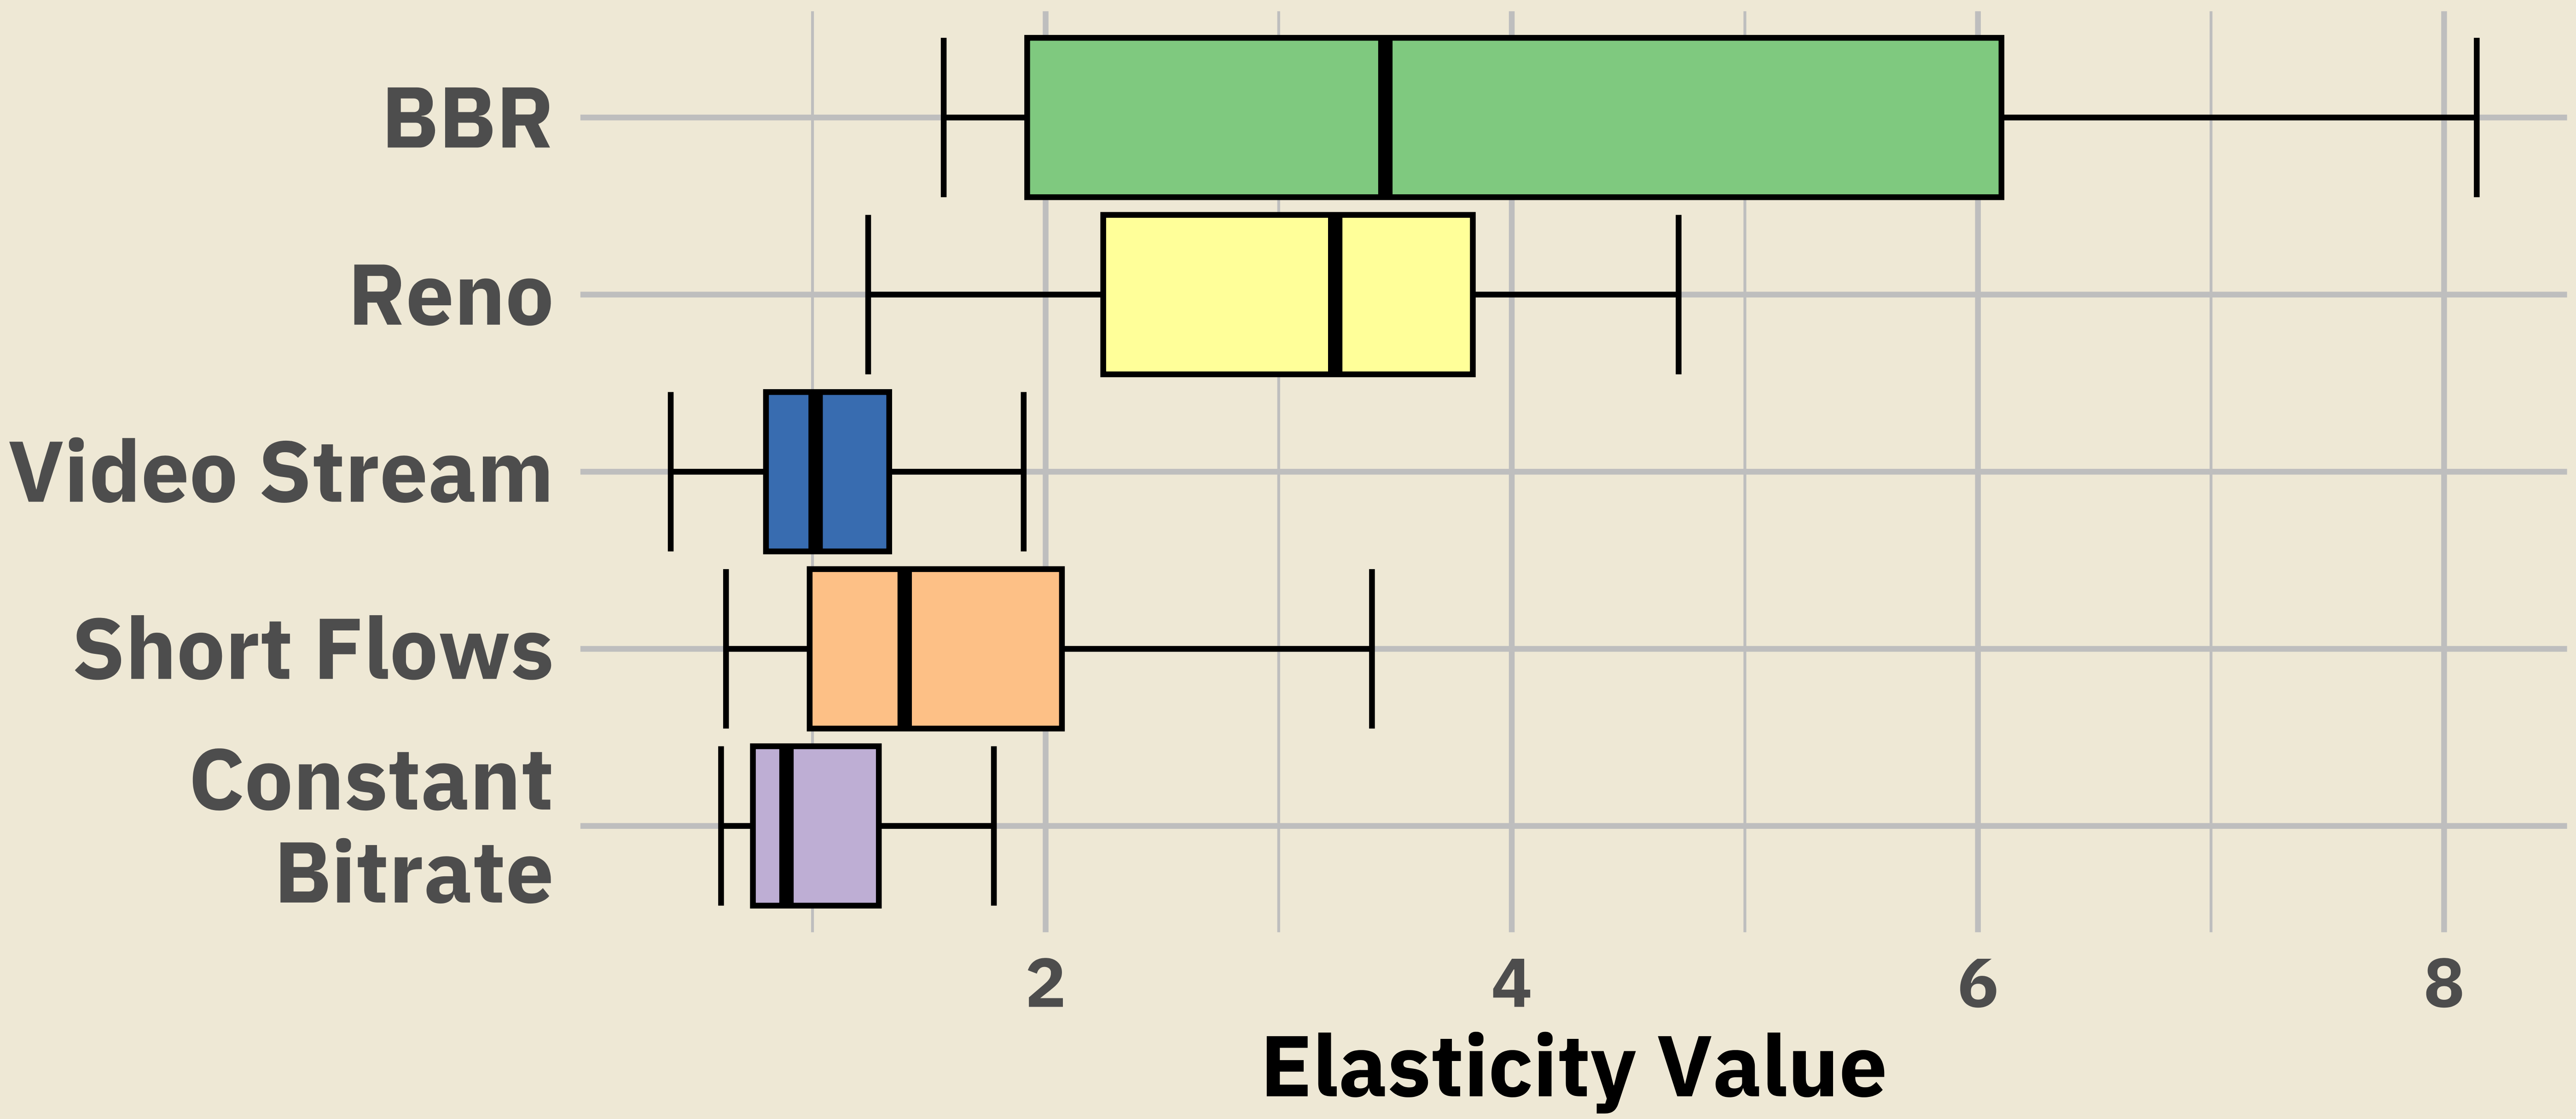

Does CCA contention typically determine a flow's bandwidth allocation in the Internet today?

Our Hypothesis: Not typically

- Operators' capacity planning usually prevents congestion
- Most flows are app-limited or operator-throttled
- We need active measurement to settle the question

If true, what are the implications?

Future CCAs can focus on bandwidth-latency tradeoff

Less concern over contention properties can make future CCAs simpler, but CCAs will continue to have a rich design space

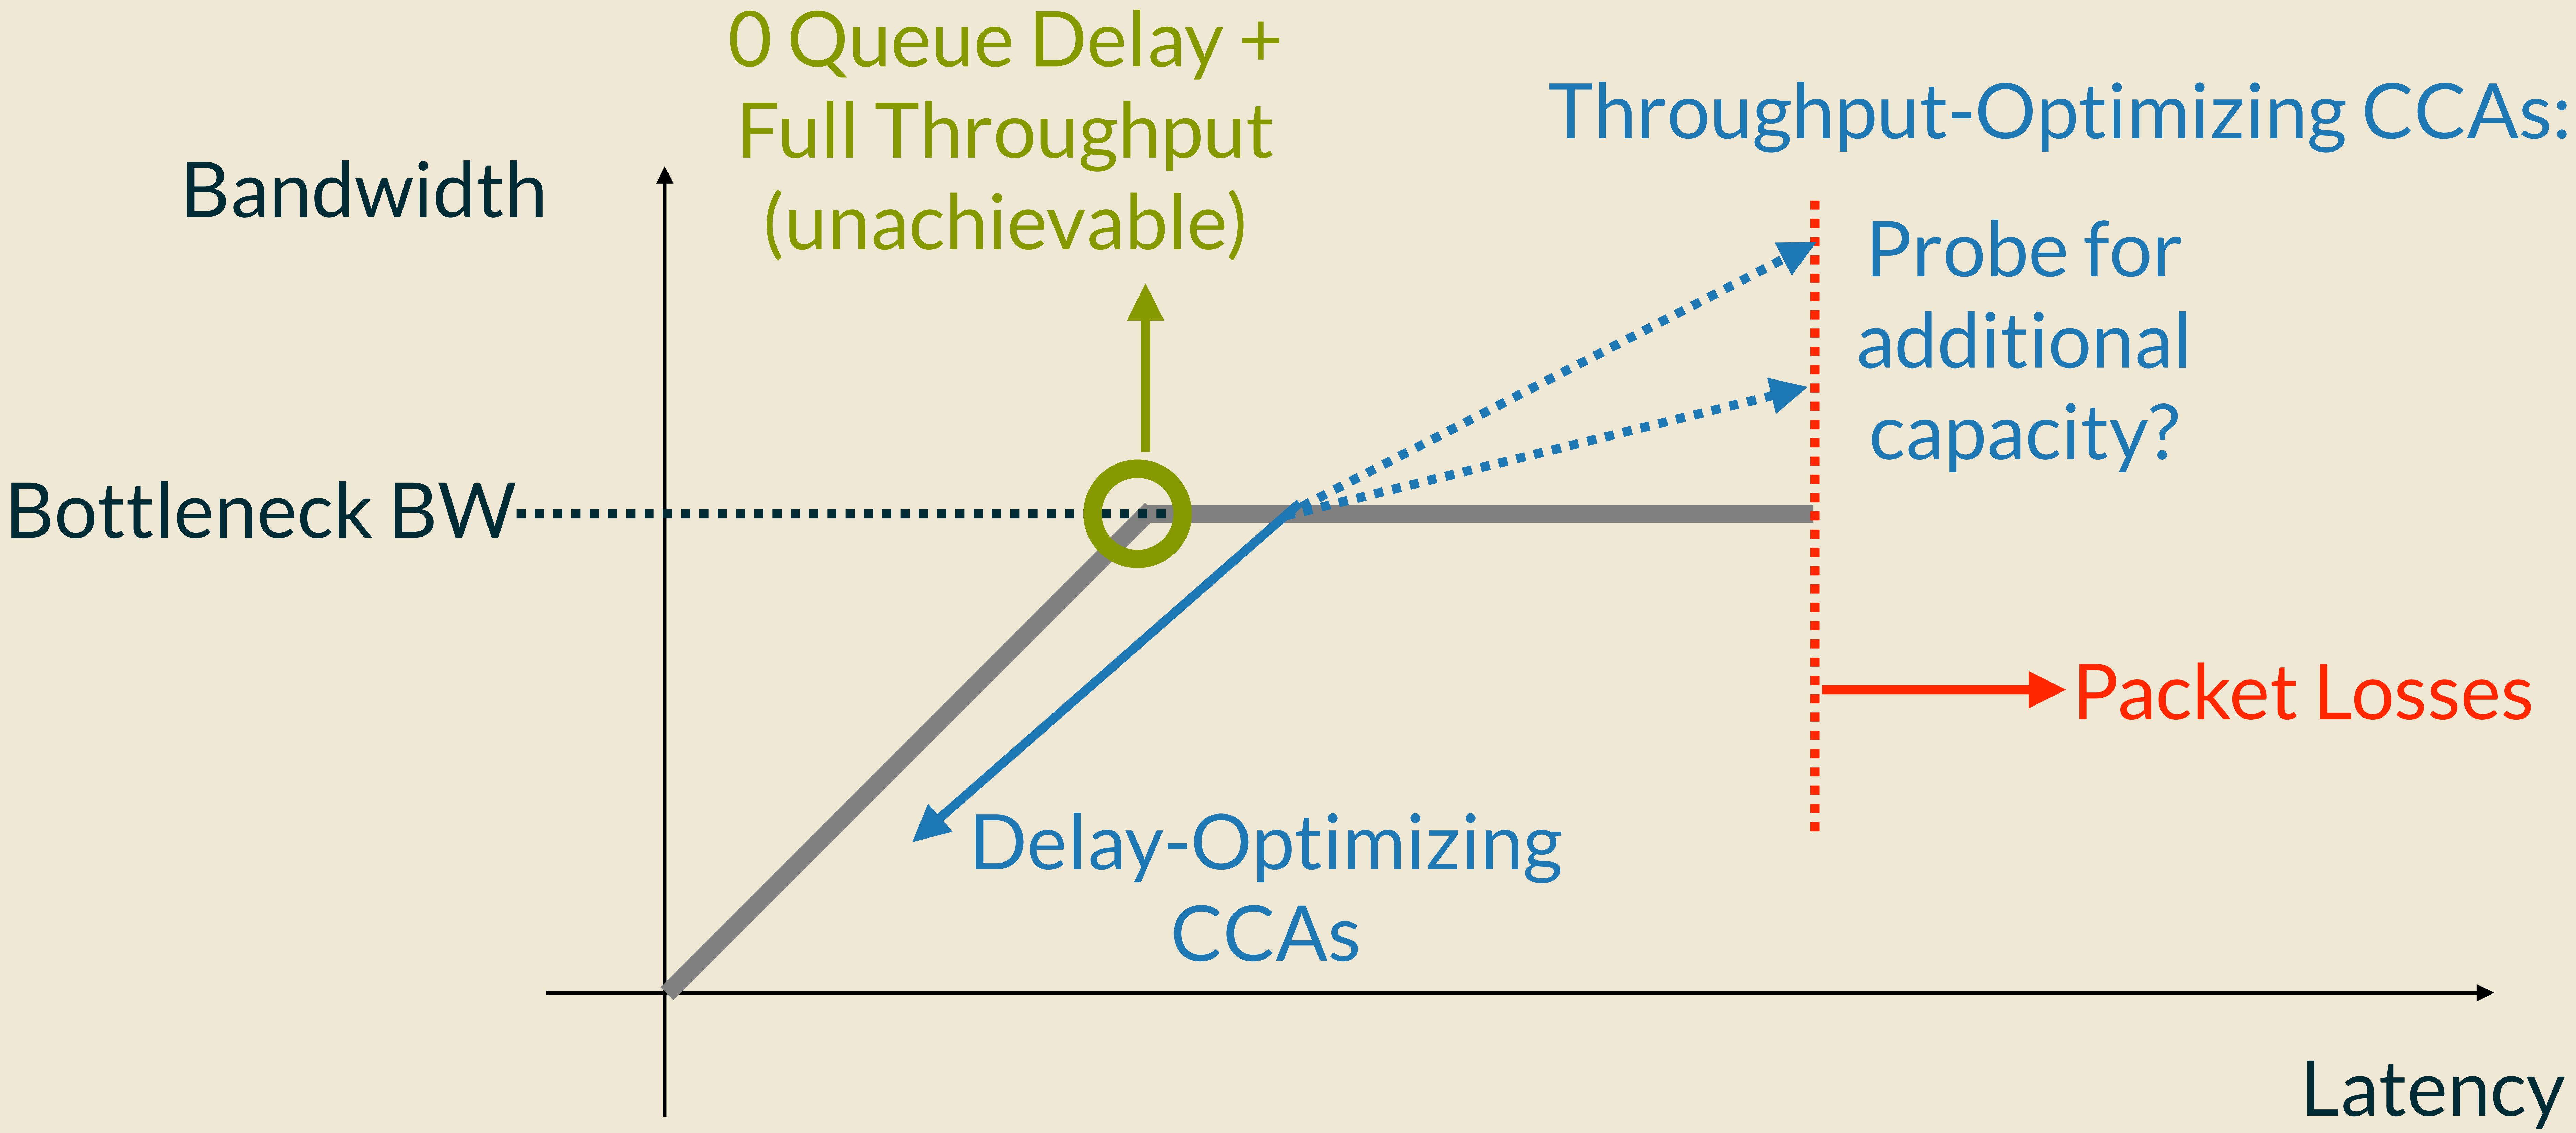

Does CCA contention typically determine a flow's bandwidth allocation in the Internet today?

Not typically:

- Operators' capacity planning usually prevents congestion
- Most flows are app-limited or operator-throttled
- We need active measurement to settle the question

↳ Join Us!

akshay@eecs.berkeley.edu
@akshayn@discuss.systems
https://akshayn.xyz