

BOLT

Sub-RTT Congestion Control for Ultra-Low Latency

Serhat Arslan, Yuliang Li, Gautam Kumar, Nandita Dukkupati




Stanford
University



The 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI '23)
April, 2023

Bandwidth-Delay-Product
is increasing!

RPC Sizes in Our Data Centers



How Would a CC Oracle Behave?

Detect **precise amount and location** of congestion,

act instantaneously.

Challenges:

Fine-Grained Feedback

Minimal Feedback Delay

Minimizing the Feedback Delay

Two Components:

1. Congestion Notification Delay
2. Under-Utilization Feedback Delay

Challenges:

Fine-Grained Feedback

Minimal Congestion Notification Delay

Latent Under-Utilization Feedback Delay

“Bolt Addresses All of These”



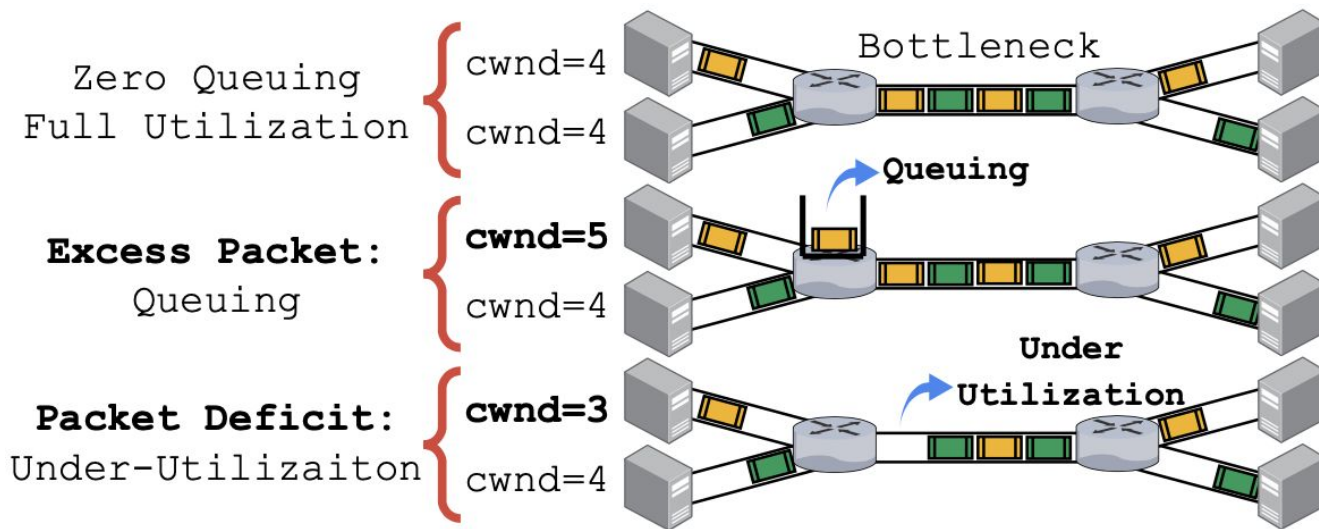
Fine-Grained Feedback

Minimal Congestion Notification Delay

Latent Under-Utilization Feedback Delay

Underlying Model

Packet Conservation Principle




Fine-Grained Feedback

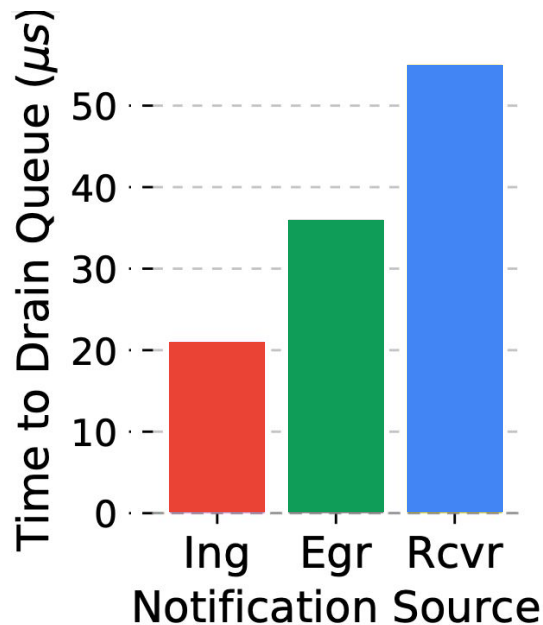
Minimal Congestion Notification Delay


Latent Under-Utilization Feedback Delay

Sub-RTT Control (SRC)



Decrease **congestion notification delay** to its absolute minimum






Fine-Grained Feedback

Minimal Congestion Notification Delay

Latent Under-Utilization Feedback Delay

Proactive Ramp-Up (PRU)



Hides under-utilization feedback delay

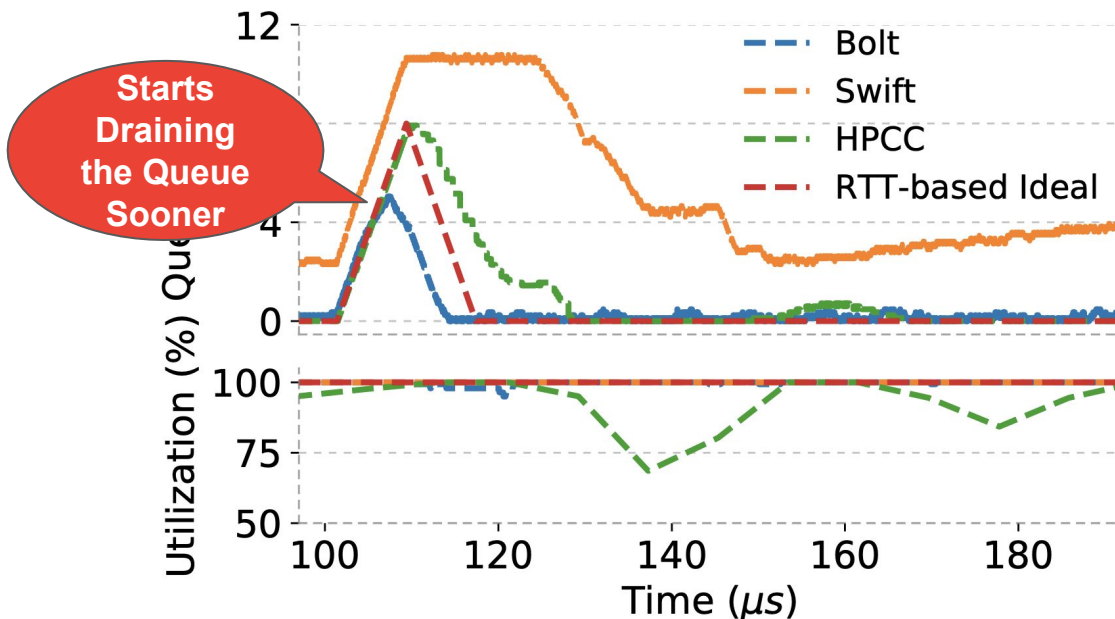
Fine-Grained Feedback

Minimal Congestion Notification Delay

Latent Under-Utilization Feedback Delay

Evaluations

Bolt's reaction to **flow arrival** versus the ideal behavior



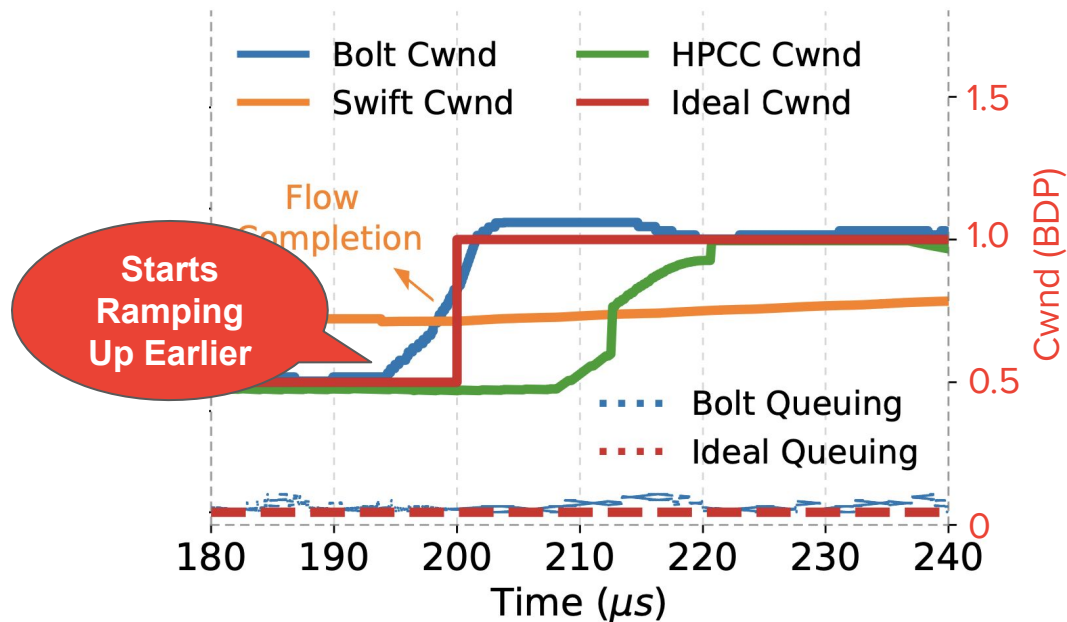
Fine-Grained Feedback

Minimal Congestion Notification Delay

Latent Under-Utilization Feedback Delay

Evaluations

Bolt's reaction to **flow completion** versus the “ideal” behavior

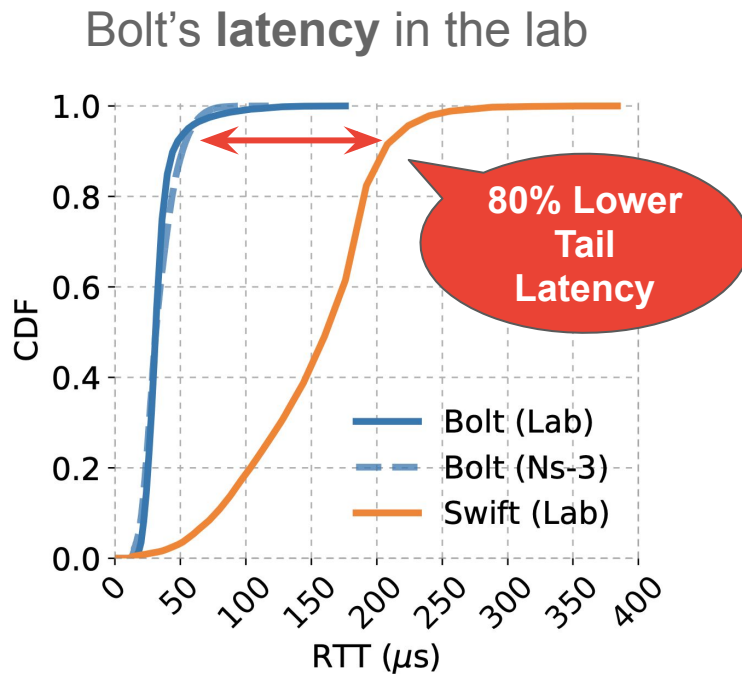


Fine-Grained Feedback

Minimal Congestion Notification Delay

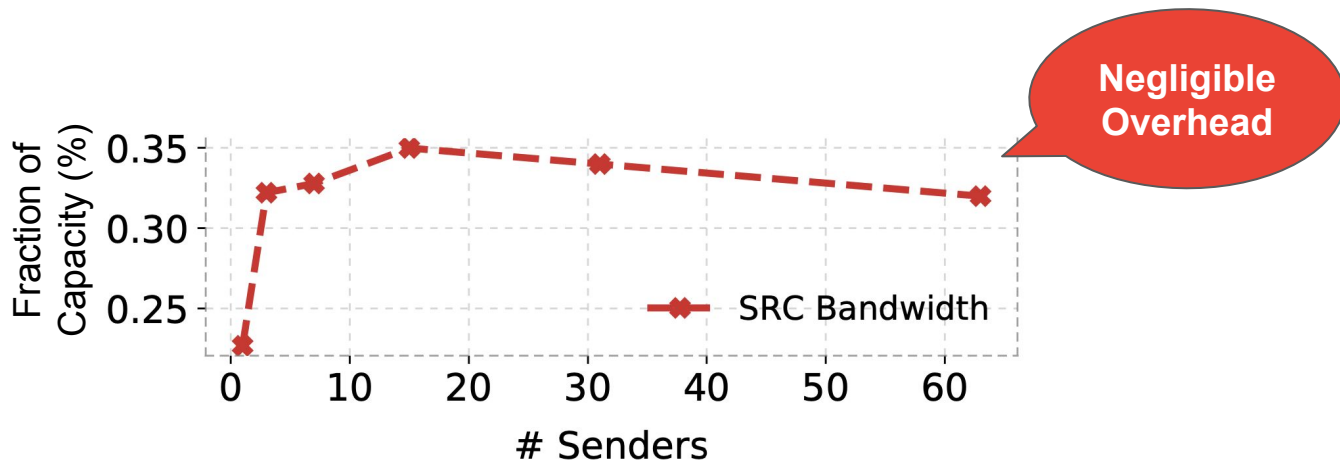
Latent Under-Utilization Feedback Delay

Evaluations



Evaluations

SRC Overhead for different levels of burstiness



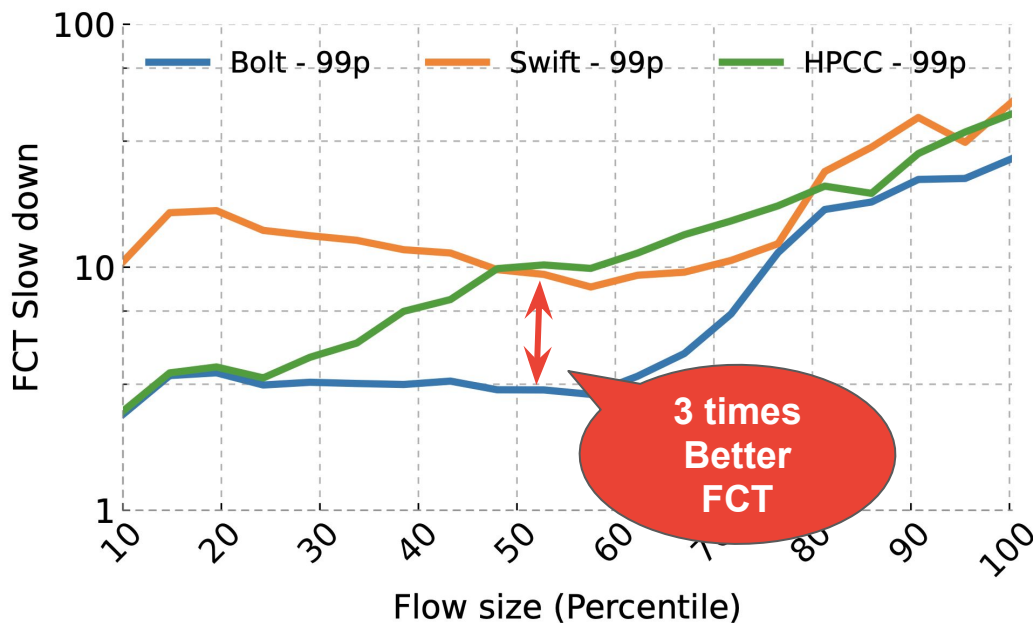
Fine-Grained Feedback

Minimal Congestion Notification Delay

Latent Under-Utilization Feedback Delay

Evaluations

Bolt's overall performance on a cluster



Fine-Grained Feedback

Minimal Congestion Notification Delay

Latent Under-Utilization Feedback Delay

Key Takeaways

- **Bolt** pushes CC towards its limits via:
 - The most **fine-grained feedback**
 - The absolute **minimum congestion notification delay** (SRC)
 - **Latent under-utilization feedback delay** (PRU)
- Results:
 - **80%** reduction in tail queuing
 - **3X** improvement in tail FCT



Serhat Arslan - sarslan@stanford.edu