

UPDATE

SEARCH – a New Slow Start Algorithm for TCP and QUIC

Jae Chung
Feng Li

Maryam Ataei Kachooei
Mark Claypool

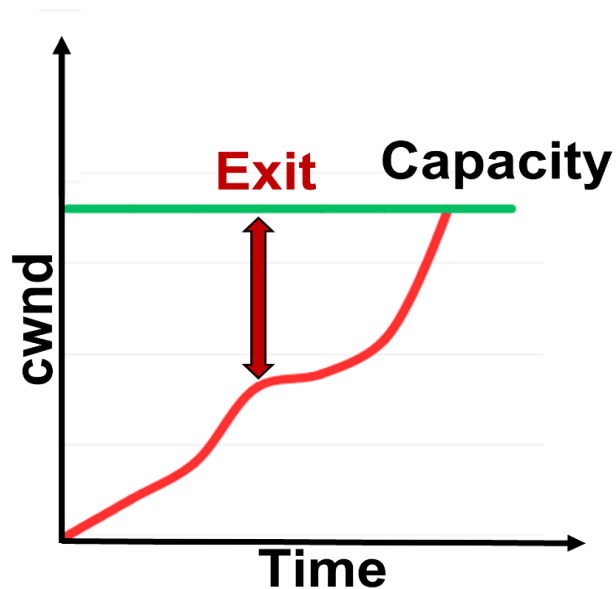
IETF CCWG

Dublin, Ireland
November 2024







Motivation



Exit Too Early



Exit Too Late




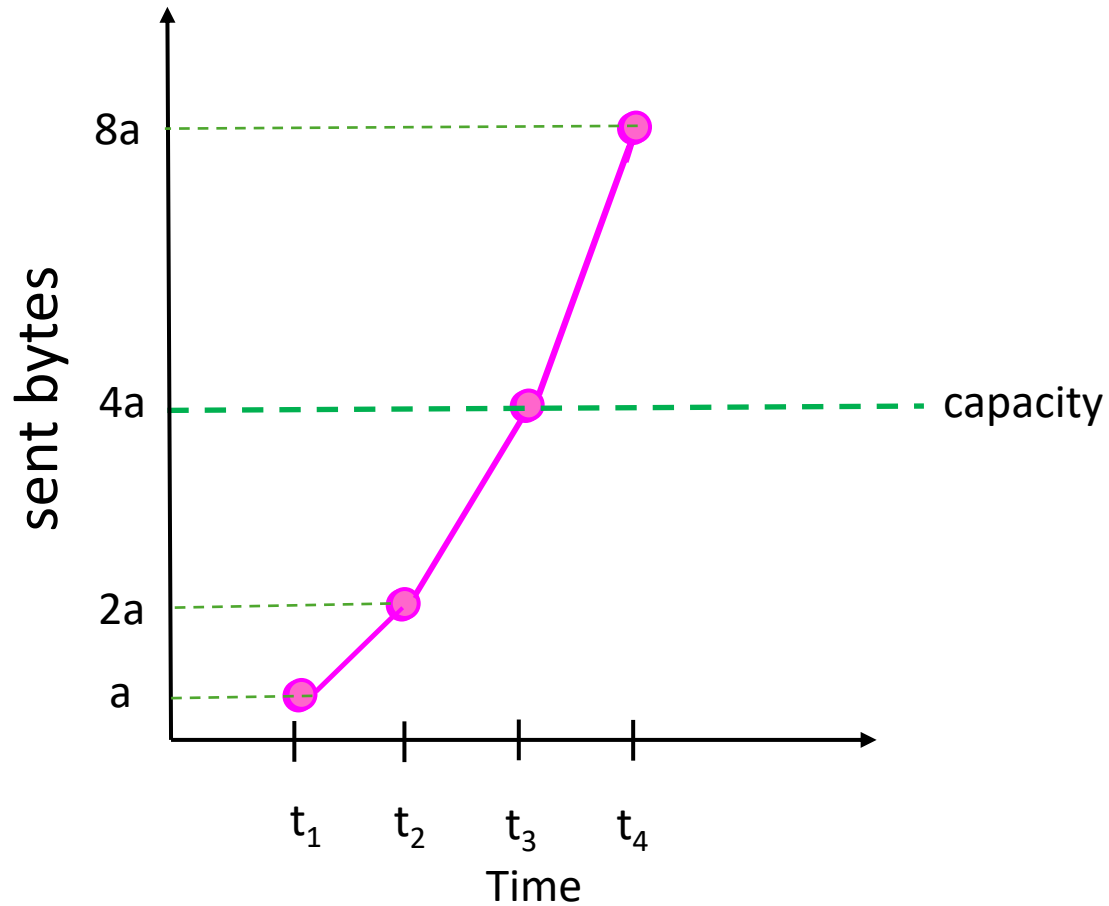
Exit at Choke Point

HyStart over wireless

SEARCH



SEARCH – Slow start Exit at Right CHokepoint





SEARCH – Slow start Exit at Right CHokepoint


$$\text{sent}' = 2 \cdot \text{delv}_{\text{previous}}$$

$$\text{diff} = \text{sent}' - \text{delv}_{\text{now}}$$

$$\text{normalized_diff} = \text{diff} / \text{sent}'$$

$$\text{normalized_diff} \geq \text{threshold?}$$

→ exit slow start





Outline

- SEARCH Review (done)
- Updates (next)
 - Algorithm
 - Reduce Bytes per Flow
 - Test Rate Limited Flows
- Next Steps



Algorithm Update

0	1	2	3	4	5
10	20	10	20	30	20

bins are **bytes received**
during that interval

0	1	2	3	4	5
10	30	40	60	90	110

bins are **cumulative bytes received**
up to that interval

Bytes delivered from **bin 1** to **bin 4**?


$$\sum_{i=1}^4 \text{bin}[i] = \text{bin}[1] + \text{bin}[2] + \text{bin}[3] + \text{bin}[4]$$

$$\text{bin}[4] - \text{bin}[0]$$




Suggested Updates from Last IETF Meeting

Per-flow memory use



Rate-limited flows





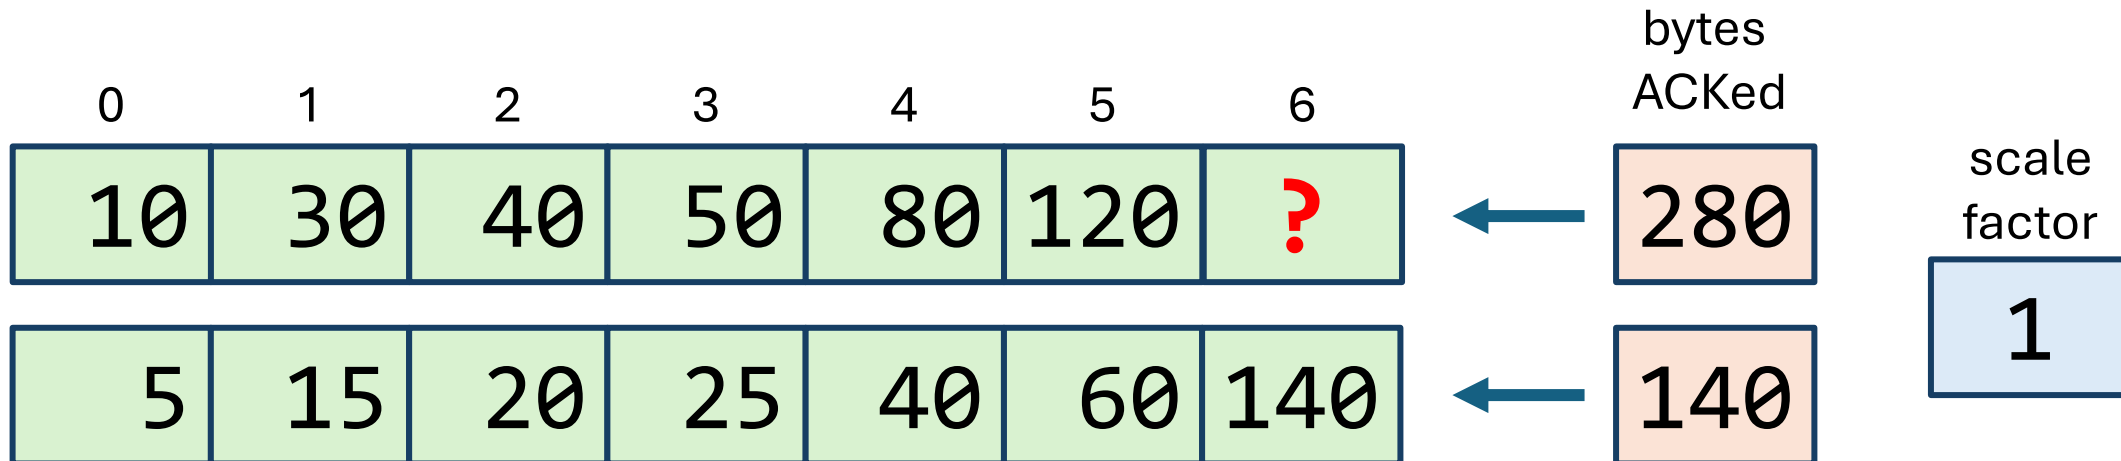
Reduce Bytes Per Flow

SEARCH-specific
per-flow data

- Reduce number of bits for each bin

u32 → u8

```
u32 bin[TOTAL_BINS]; // array of bins
u32 bin_duration_us; // duration of each bin
u32 bin_end_us; // end time of latest bin
s32 curr_idx; // total number of bins
u8 scale_factor; // divisions by 2 (shifts)
```





Error with Reduced Bits Per Bin

$$\text{sent}' = 2 \cdot \text{delv}_{\text{previous}}$$
$$\text{diff} = \text{sent}' - \text{delv}_{\text{now}}$$

$$\text{normalized_diff} = \text{diff} / \text{sent}'$$

$\text{normalized_diff} \geq \text{threshold}?$

→ exit slow start

`u32 bin[TOTAL_BINS];` normalized_diff



`u8 bin[TOTAL_BINS];` normalized_diff



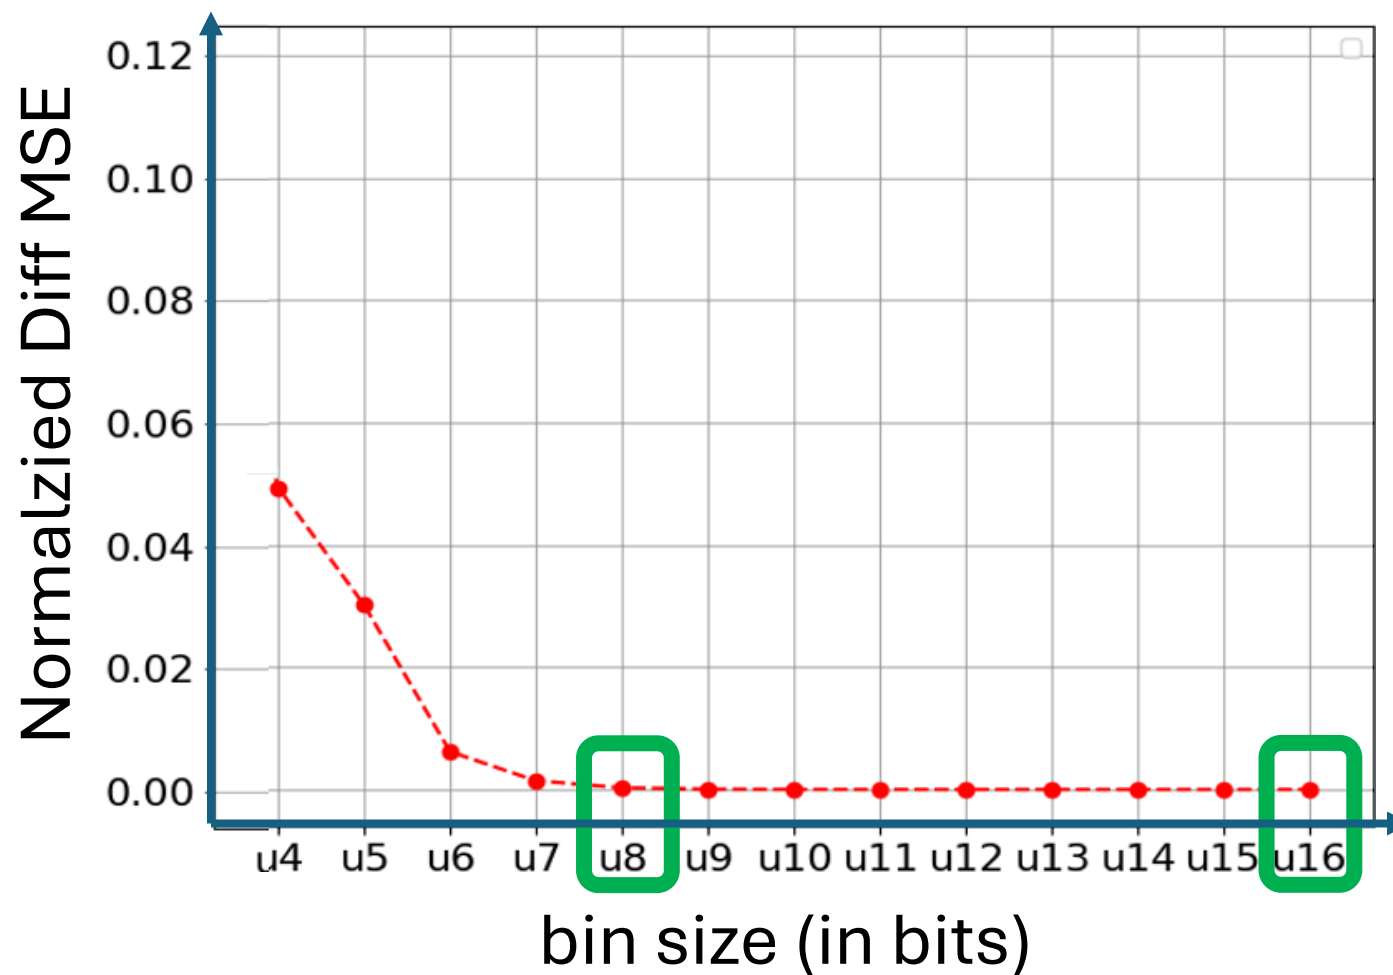
Error with Reduced Bits Per Bin

```
u32 bin[TOTAL_BINS];  
normalized_diff
```



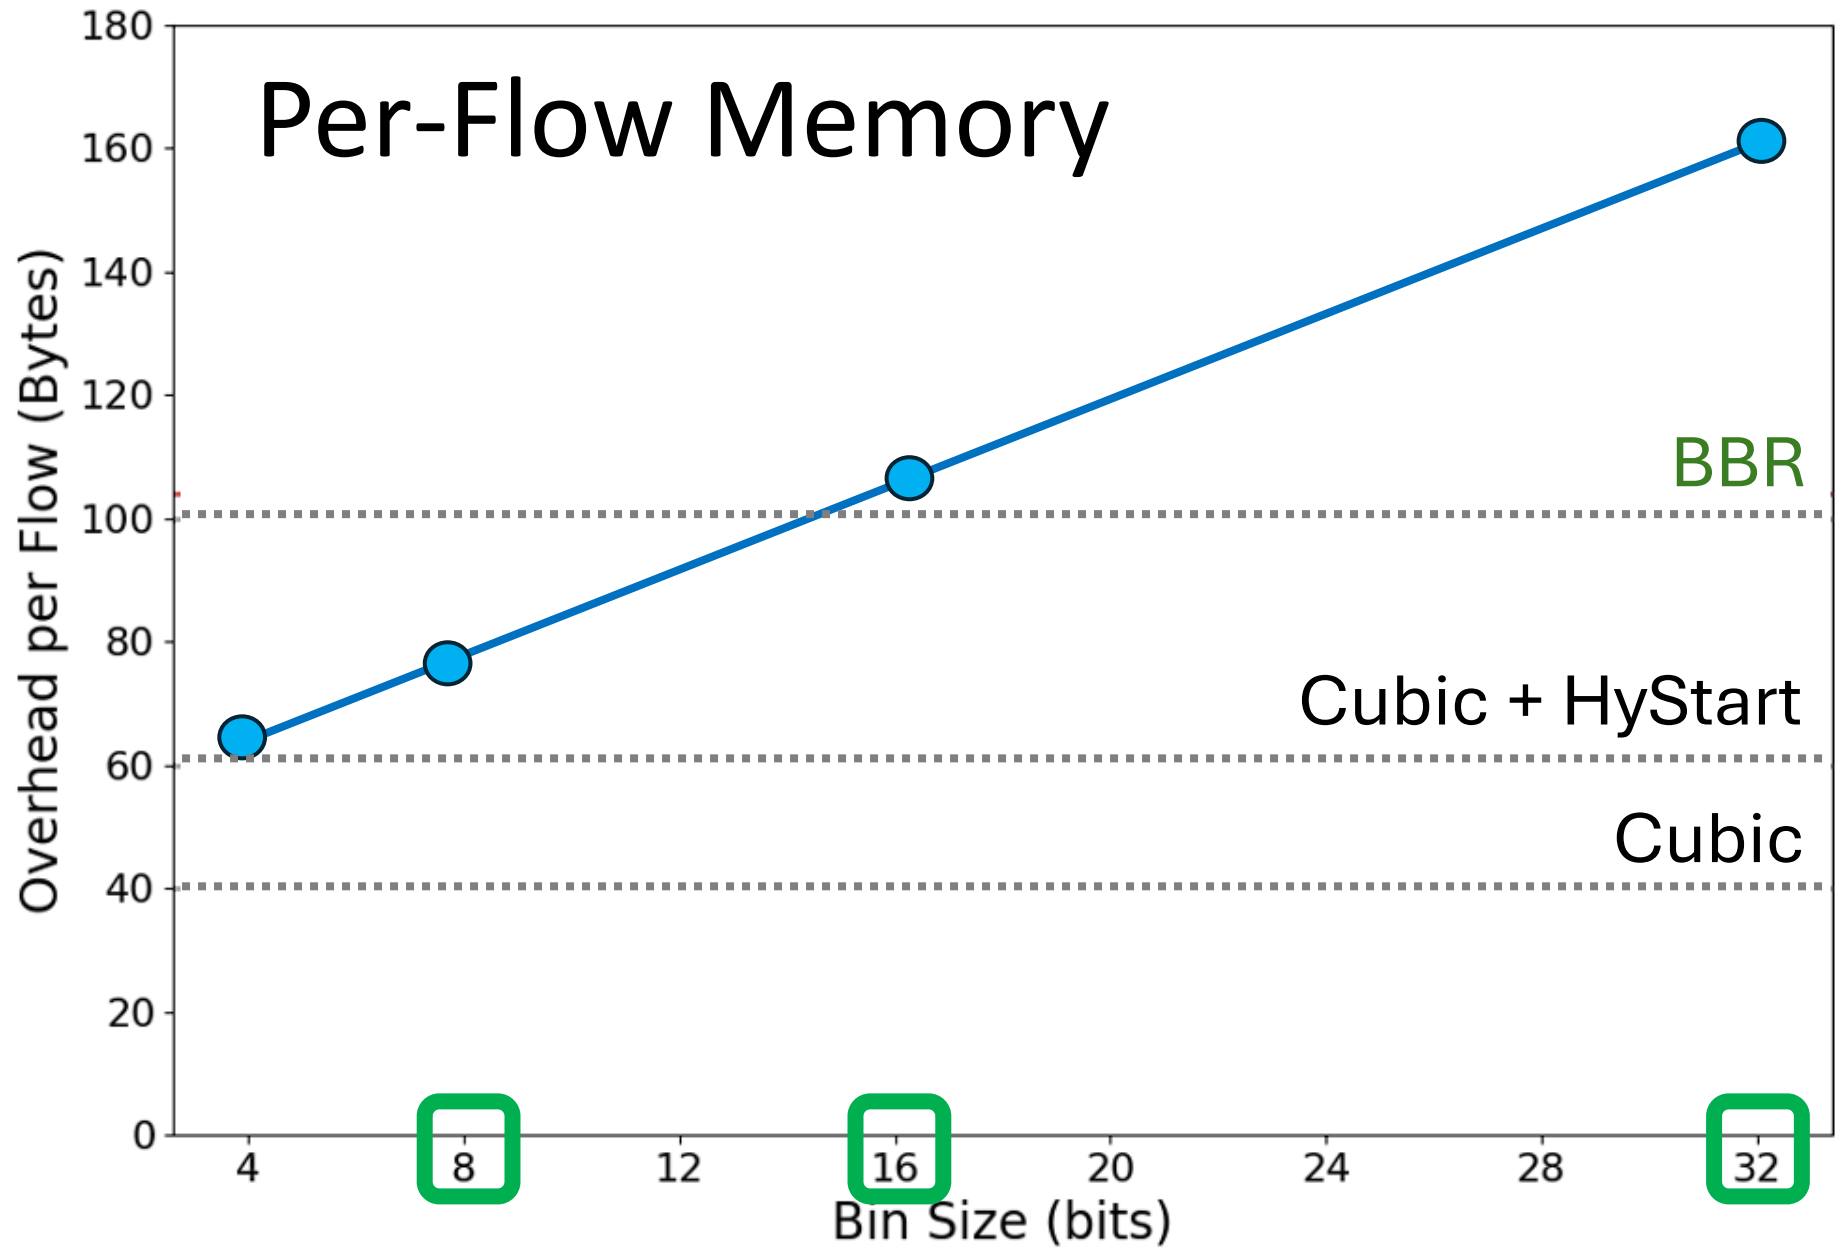
normalized_diff

```
u8 bin[TOTAL_BINS];
```



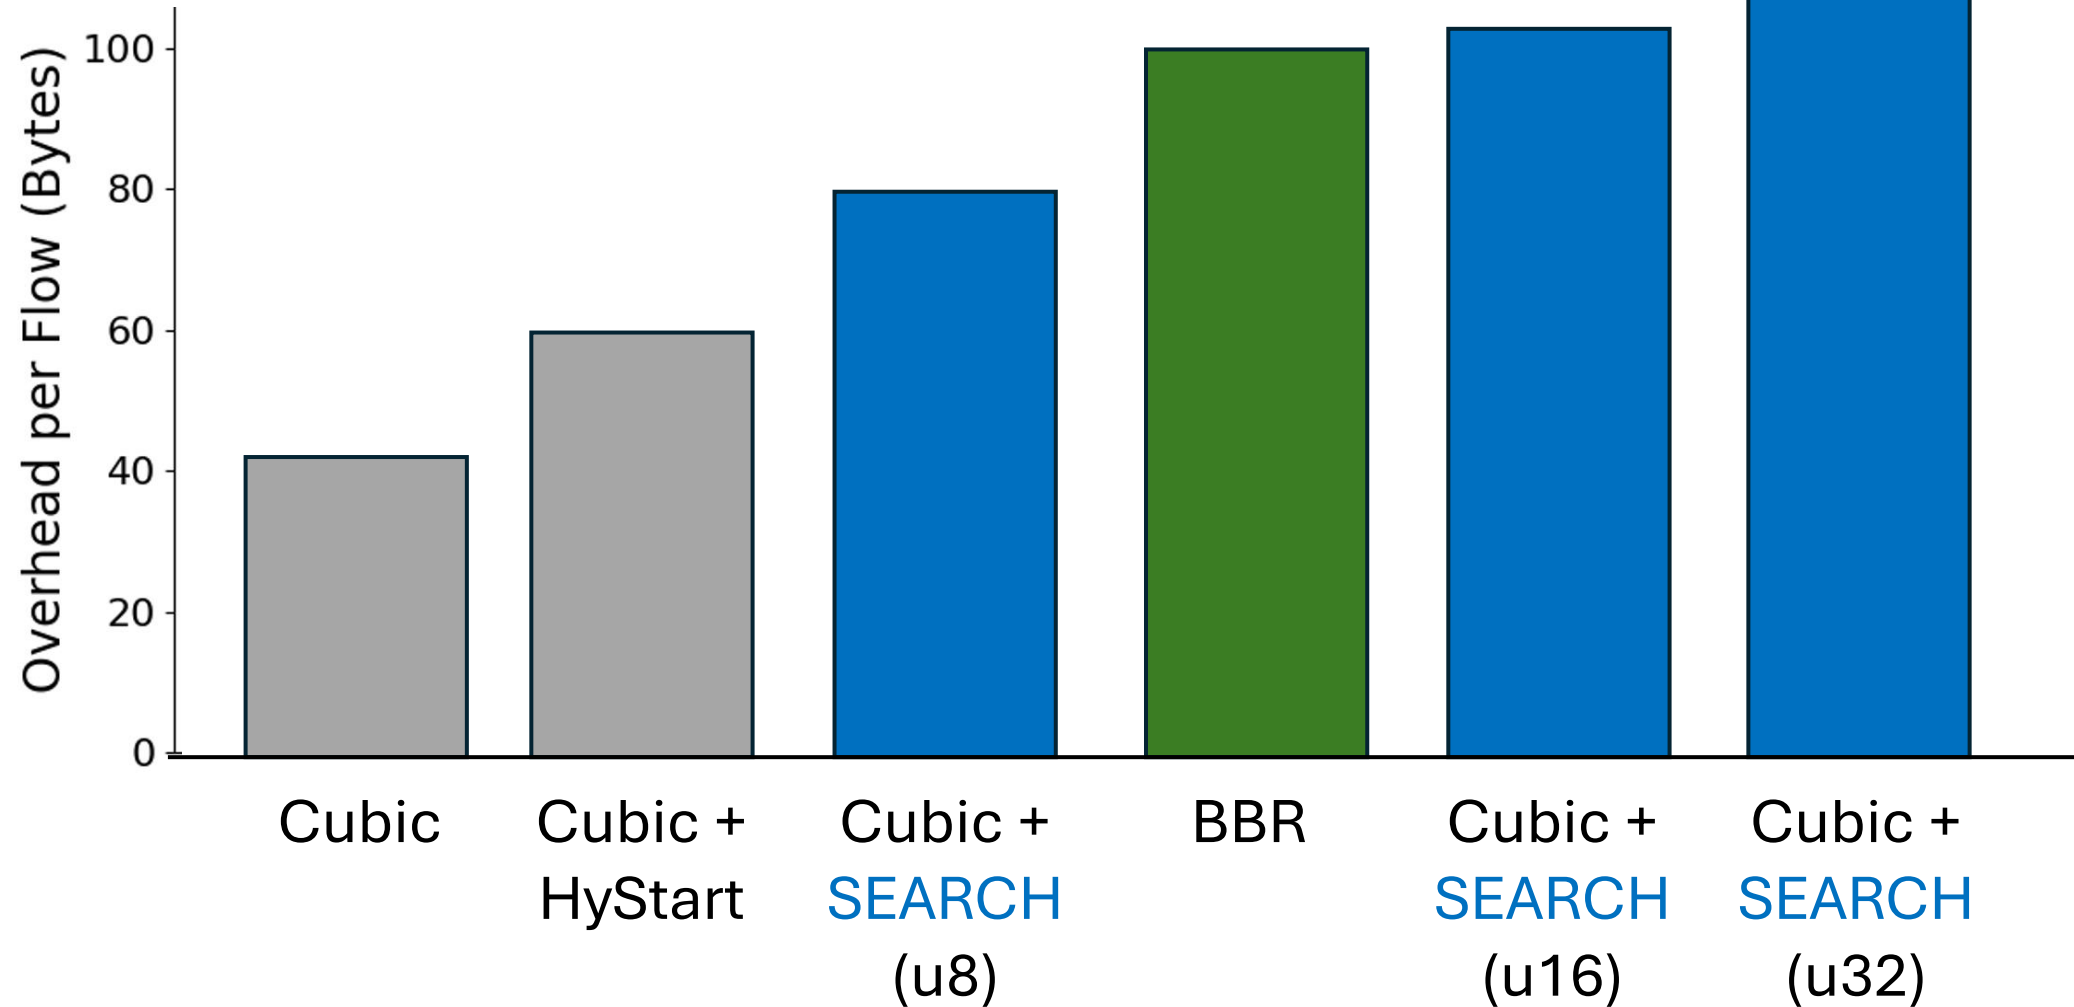


Linux
Private
ICSK_CA_PRIV
104 bytes






Per-Flow Memory






Suggested Updates from Last IETF Meeting

Per-flow memory use




Rate-limited flows





Test Rate Limited Flows


- Application limited TCP flow
 - **Periodic** – e.g., HTTP streaming
 - **Intermittent** – e.g., Web browsing
 - Application rate limited – e.g., CBR traffic
- Receiver / Sender limited TCP flow
 - e.g., buffer size limits





Test Rate Limited Flow – Application

- AWS EC2 Linux
- netem to add 200 ms latency
- Application limited flow (e.g., HTTP server)
 - 3x 32 MB chunks at 5 s intervals
 - 3x 64 MB chunks at 5 s interval





Outline

- SEARCH Review (done)
- Updates (done)
 - Algorithm (done)
 - Bytes per Flow (done)
 - Rate Limited Flows (done)
- Next Steps (next)



Next Steps

- CCWG seeks **empirical evidence** of safety
- Most impactful → data from **deployments** in the field
- Looking for volunteers!
- We will help with SEARCH deployment
 - Provide implementation (currently, **Linux** and **Quicly**, working on **FreeBSD** version)
 - Help with installation and trouble shooting
 - Give assistance with data collection and analysis



Summary



SEARCH

<https://search-ss.wpi.edu/>

- **SEARCH**

- Determines “choke point” from expected delivered bytes
- Exits slow start after congestion point, before loss

- Updates (now version 3.0)

- Algorithm tweak (cumulative bytes)
- Reduced bytes per flow
- App-limited flows

- Looking for volunteers to try it out!

Thank-you for your attention!

SEARCH – a New Slow Start Algorithm for TCP and QUIC

Jae Chung

Feng Li

Maryam Ataei Kachooei

Mark Claypool

IETF CCWG

Dublin, Ireland

November 2024





References

- Improving TCP Slow Start Performance in Wireless Networks with SEARCH
 - *IEEE World of Wireless, Mobile and Multimedia Networks (WoWMoM)*
 - Perth, Australia, June 2024
- Improving QUIC Slow Start Behavior in Wireless Networks with SEARCH
 - *IEEE Local and Metropolitan Area Networks (LANMAN)*
 - Boston, Massachusetts, USA, July 2024
- Implementation of the SEARCH Slow Start Algorithm in the Linux Kernel
 - *0x18 NetDev Conference*
 - Santa Clara, California, USA, July 2024