

BCN Calibration Simulation Results


Zhi-Hern Loh
zloh@fulcrummicro.com
10/12/2006



Workload

- **Traffic Type: 100% UDP (or raw Ethernet) Traffic**
- **Destination: EP0-EP3 sending to EP4**
- **Frame Size Distribution: 1500 byte fixed**
- **Arrival Distribution: Bernoulli temporal distribution**
- **Offered load at endpoint = 50%**

Topology



- Link capacity 10Gbps
- Core switch egress port buffer size infinite
- Rate limiter queue buffer size 150KB
- Switch latency (1 us)
- Link length (not modelled, 0 latency)
- Endpoint response time (not modelled, 0 latency)

BCN Parameters

- **Qeq 375 * 64 byte pages**
- **Frame Sampling 150KB +/- 5KB (random jitter)**
- **W = 2**
- **Gi = 5.3 x 10⁻¹**
- **Gd = 2.6 x 10⁻⁴**
- **Ru = 1 Mbps**

BCN Mod 1

- **What if we measured the offered load and use that as basis for rate decrease?**

- May help with latency in BCN reaction, rate limiters do not need to always start from 100%.

- **Pseudo code**

- If we need to do rate decrease ($F_b < 0$)
 - Calculate new rate based on measured rate at rate limiter instead of previously set rate of rate limiter.
 - Measured rate = actual sending rate at rate limiter
 - Could be implemented with exponential moving average

Simulations

- **Parameters**

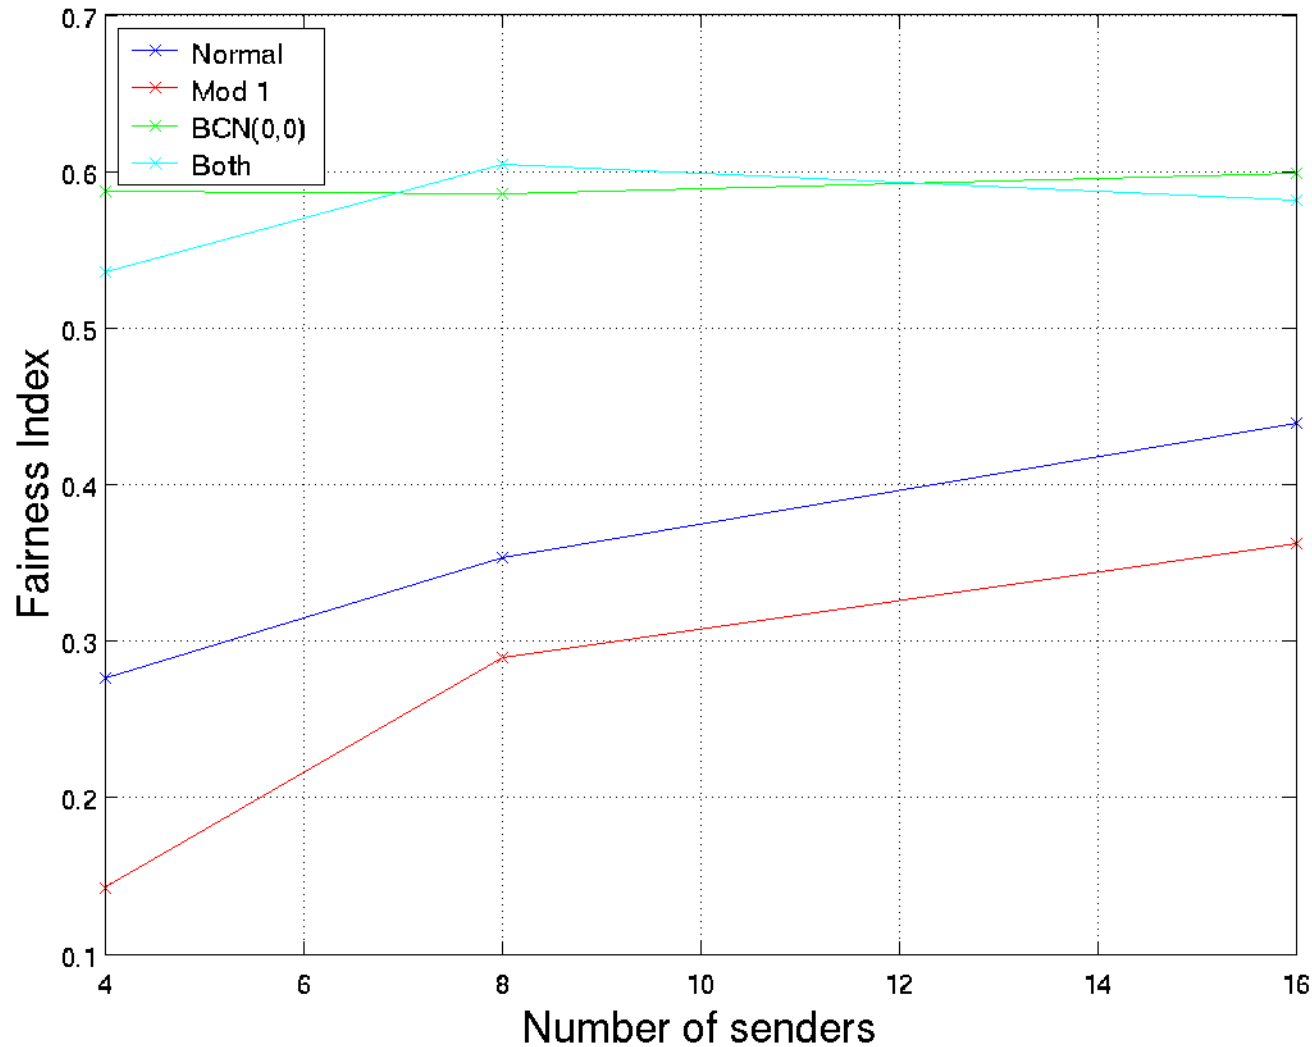
- No limits on max queue length at congestion point
- Data sets
 - Normal BCN with default parameters, $W=2$, $G_i=5.3 \times 10^{-1}$, $G_d=2.6 \times 10^{-4}$, $R_u=1 \text{ Mbps}$
 - Mod 1
 - BCN(0,0)
 - $Q_{sc} = 125 \text{ KB}$, $R_{min} .1 \text{ Gbps}$, $T_{max} 100 \mu\text{s}$, Drift disabled.
 - Mod 1 + BCN(0,0)

Sweep over increasing “congestion”

- **Sweeping over number of senders, as means of increasing congestion.**
- **Start with 4 senders (Same as baseline simulation topology).**
- **Next data set using 8 senders with half of previous load, 25% instead of 50%.**
- **Next data set using 16 senders with quarter of (4 sender) load, 12.5%.**
- **Each data point created from 25 runs.**

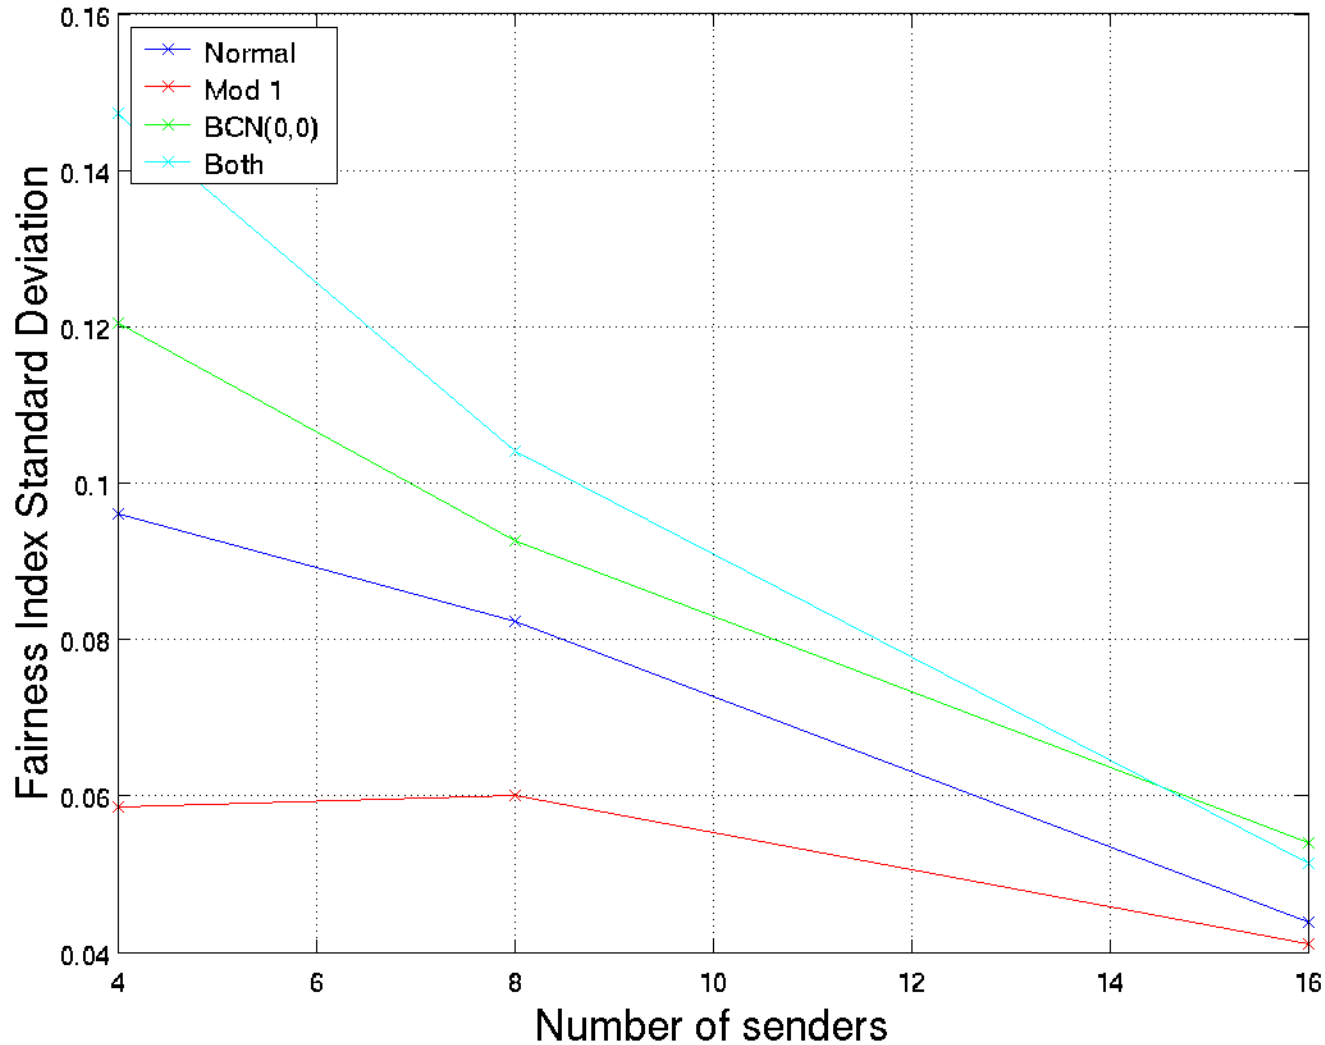
Results – Fairness Index

Mean FI vs num senders

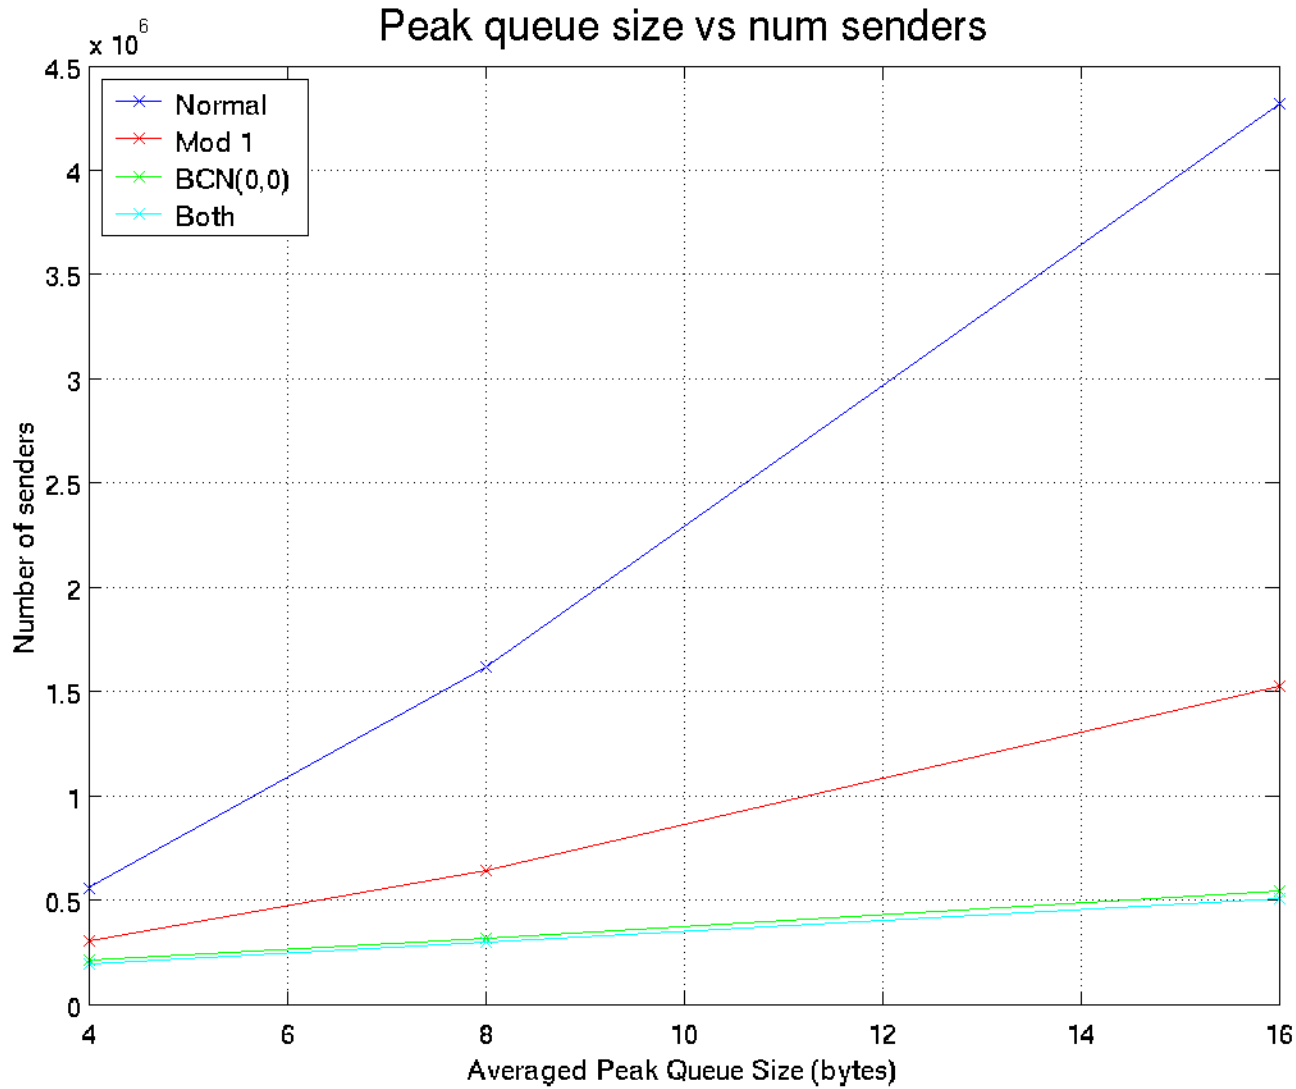


Results – FI Standard Deviation

FI Std. vs num senders



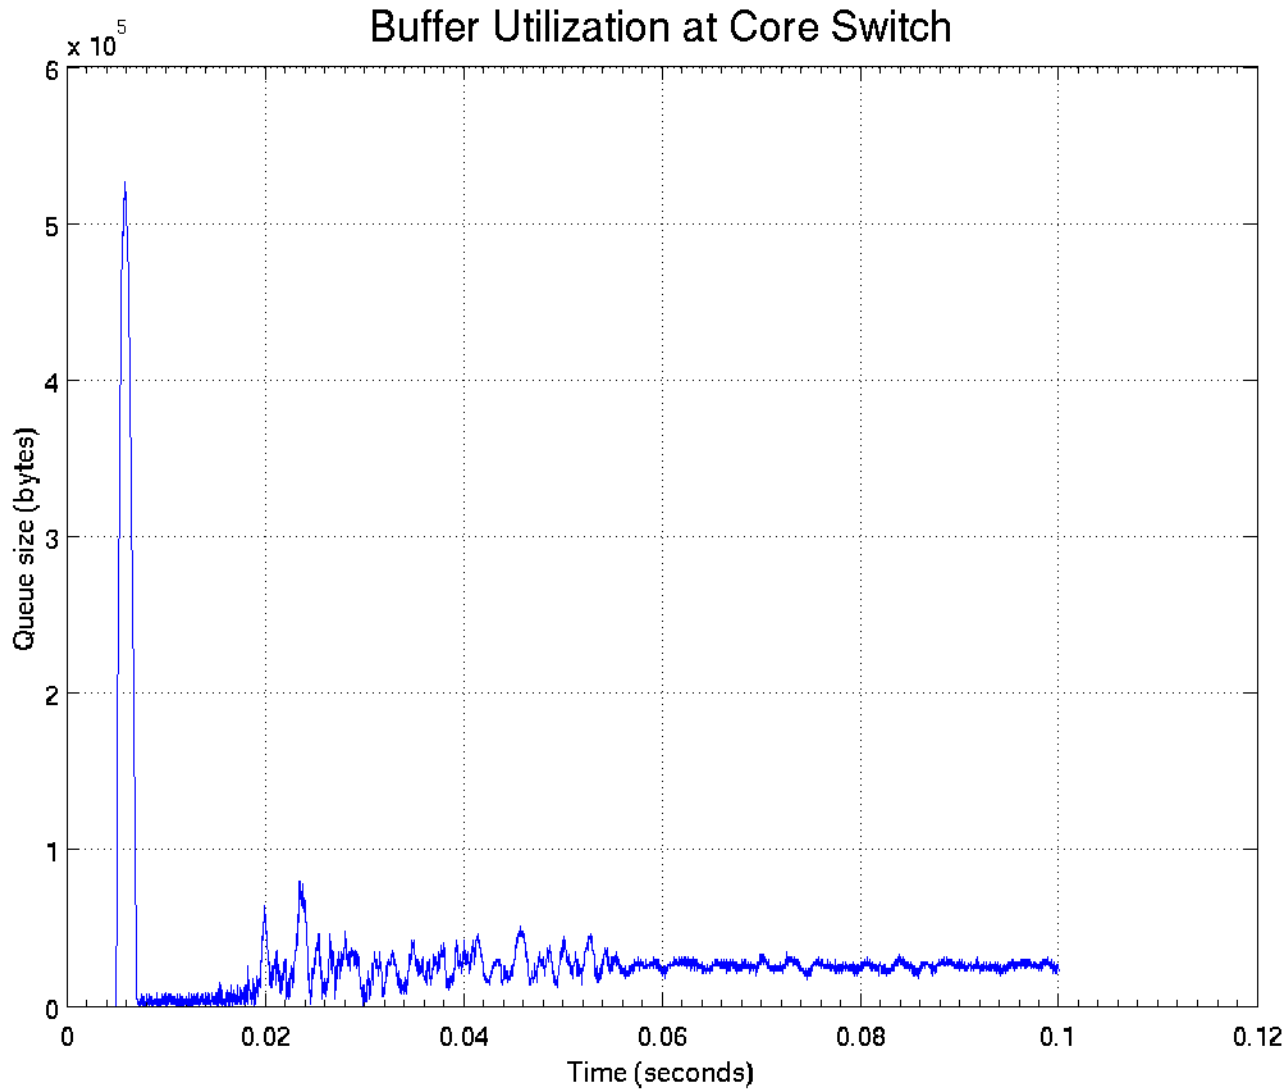
Results Averaged Peak Queue Size



Results


- **Previous values of Rmin was 1 Gbps**
 - Does not work when we have more than 10 to 1 congestion.
 - For this simulation sweep, .1 Gbps was used for Rmin
- **BCN(0,0)**
 - Introduces more unfairness. Some senders stop while others continue.
 - Long time to recovery, already saw this in earlier results from Raj Jain and Davide Bergamasco.

BCN(0,0) + MOD 1 Run 16 Senders



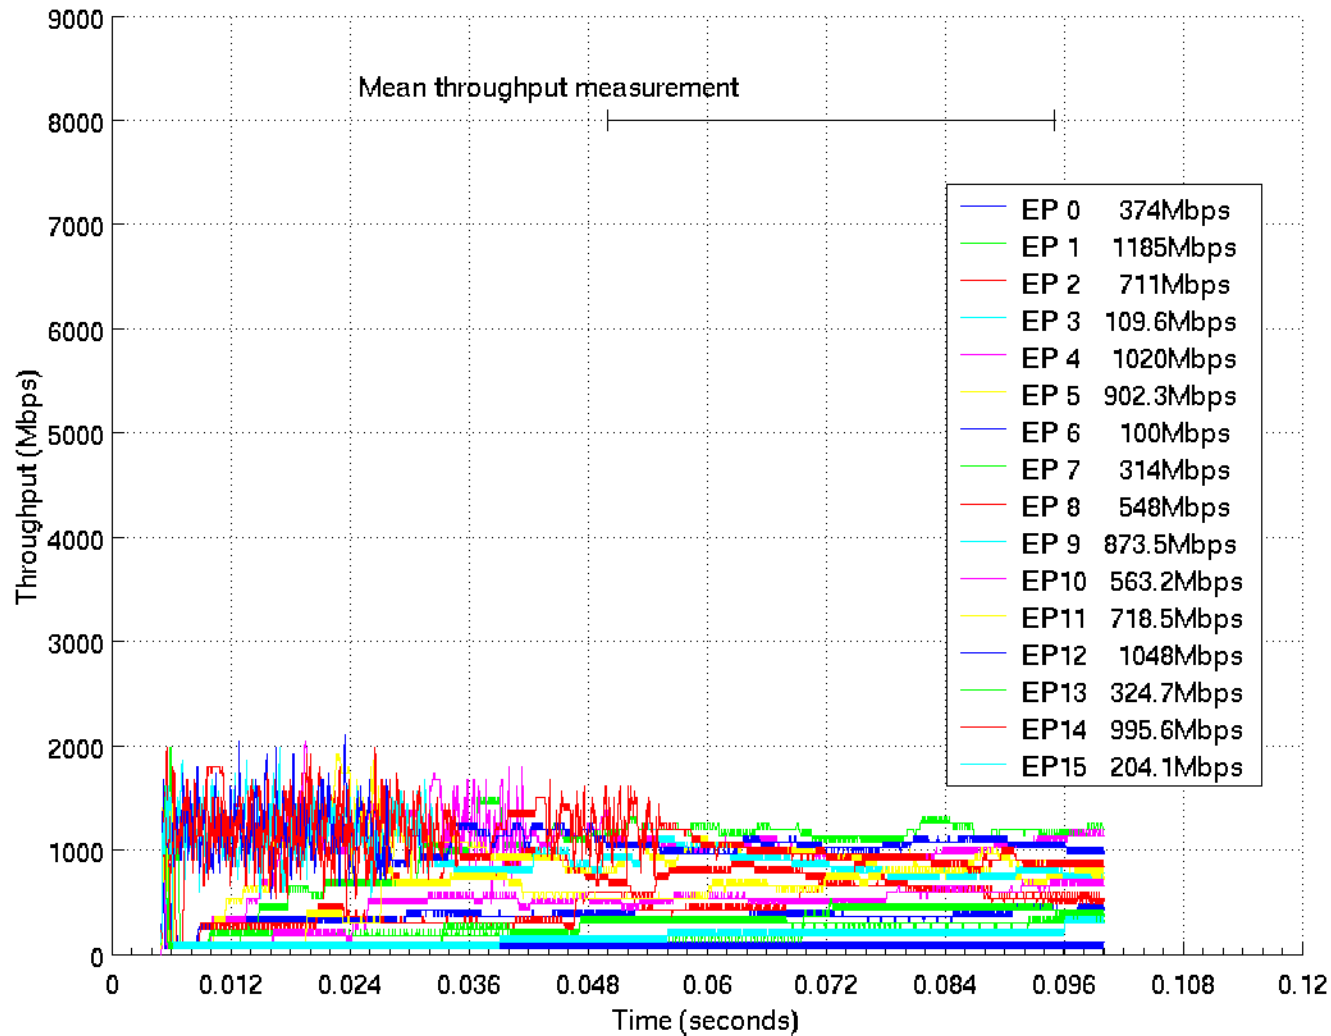
BCN(0,0) + MOD 1 Run 16 Senders

RLQ Rate vs Time



BCN(0,0) + MOD 1 Run 16 Senders

Accepted Throughput



BCN(0,0) + MOD 1 Run 16 Senders

Egress Throughput vs Time

