BCN Simulation Results
Innocent Flows With Varying Hot Spot Degree Experiment
IEEE Plenary Dallas

Zhi-Hern Loh
zloh@fulcrummicro.com
Nov 2006

FULCRUM microsystems

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%**

Baseline Topology

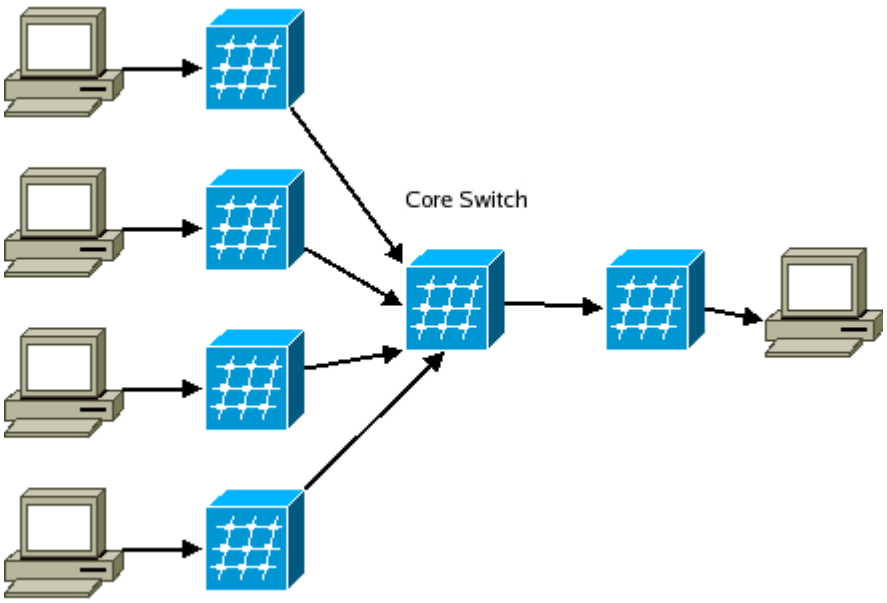

- Link capacity 10Gbps
- Core switch egress port buffer size infinite
- Rate limiter queue buffer size 150KB
- Switch latency (1 us)
- Link length (not modelled, 0 latency)
- Endpoint response time (not modelled, 0 latency)

Baseline 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

Innocent Flow Topology

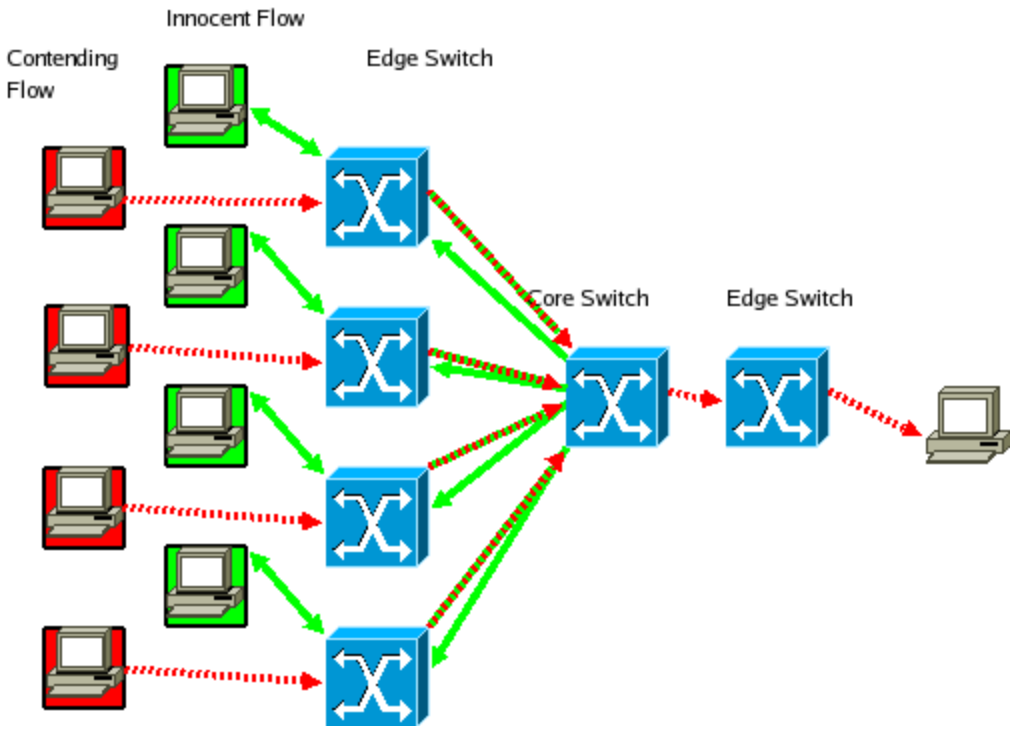

FULCRUM microsystems

Simulation Setup

- **Normal BCN, pause enabled (no BCN(0,0) BCN(max))**
  - Pause On 140K, Pause Off 130K, Sample Period 100K
- **Run with increasing hotspot degree, i.e. increasing number of senders 2, 4, 6 ...**
- **Keep hotspot severity constant, total congesting throughput = 20Gbps**

Simulation Setup ...

- **Toplogy A**
  - 1 Innocent Flow Per Edge Switch
  - 1 Contending Flow Per Edge Switch
- **Topology B**
  - 5 Innocent Flows Per Edge Switch
  - 5 Contending Flows Per Edge Switch
- **Topology C**
  - 10 Innocent Flows Per Edge Switch
  - 10 Contending Flows Per Edge Switch

Define Metric: Time To Stability

- **Want to measure length of congestion period**
  - Indicates how long pause will be asserted when pause enabled.
  - Rough measure of drop period if dropping frames.
- **Metric**
  - Time To Stability = Simulation time when queue length first falls below 2 * Qeq
  - Rough metric, not meaningful if system is not stable

Hot Spot Degree: 22 Senders

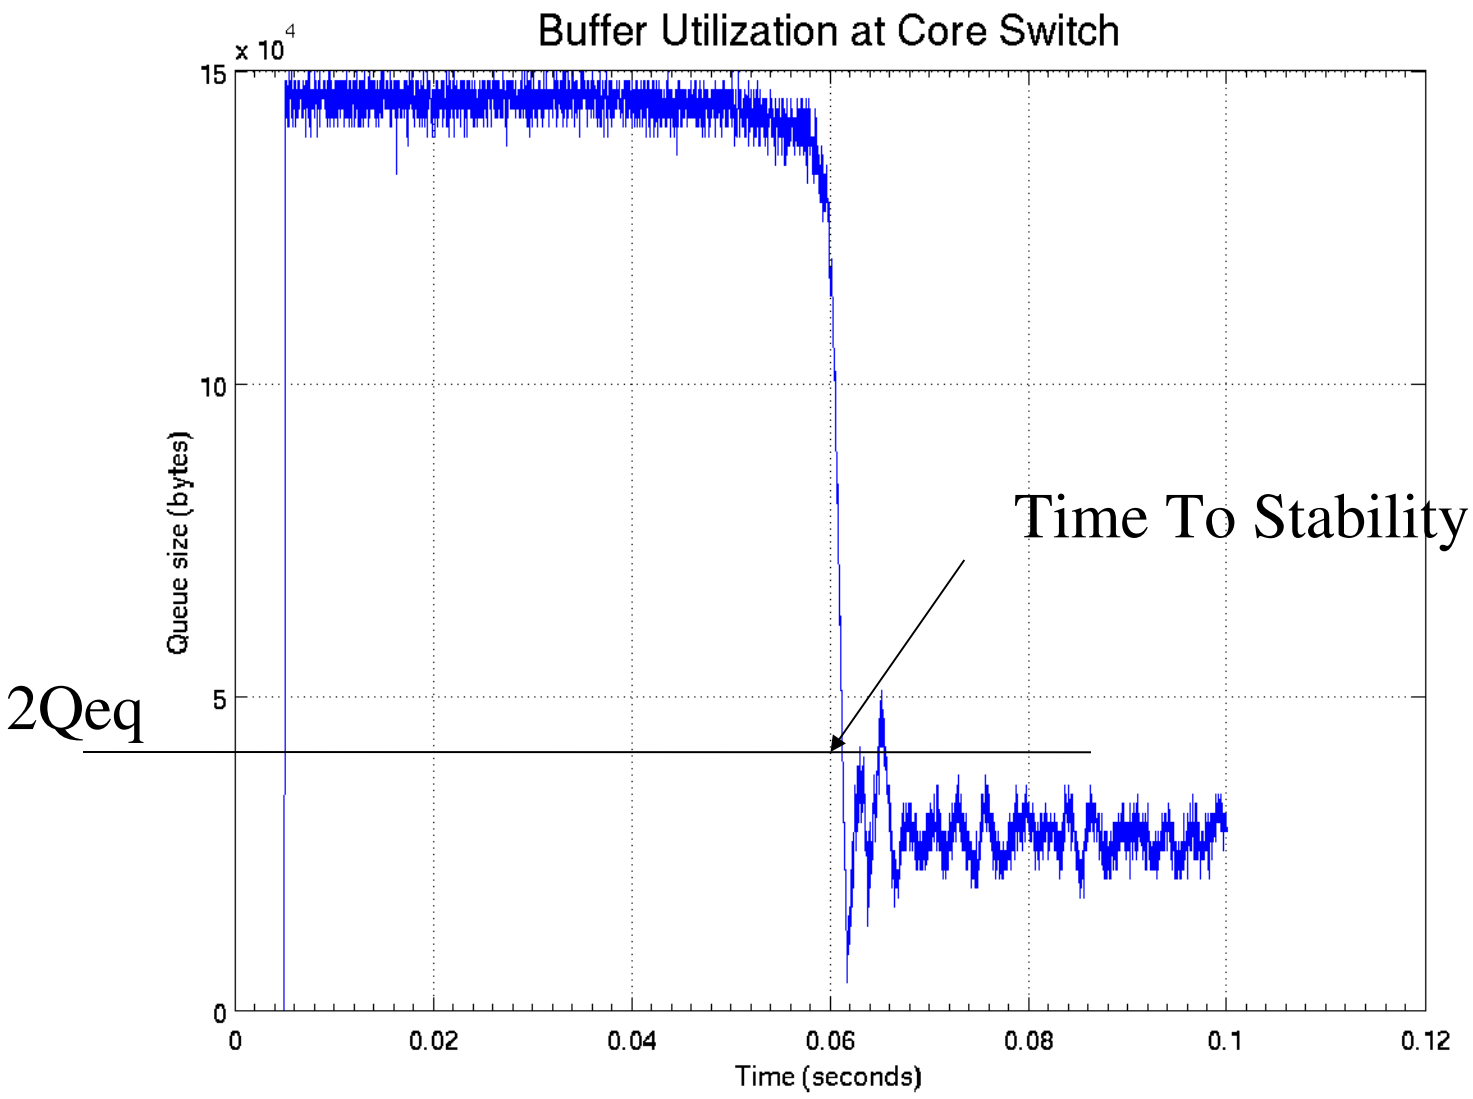

FULCRUM microsystems

Hotspot Degree 22: Innocent Flow Thpt.

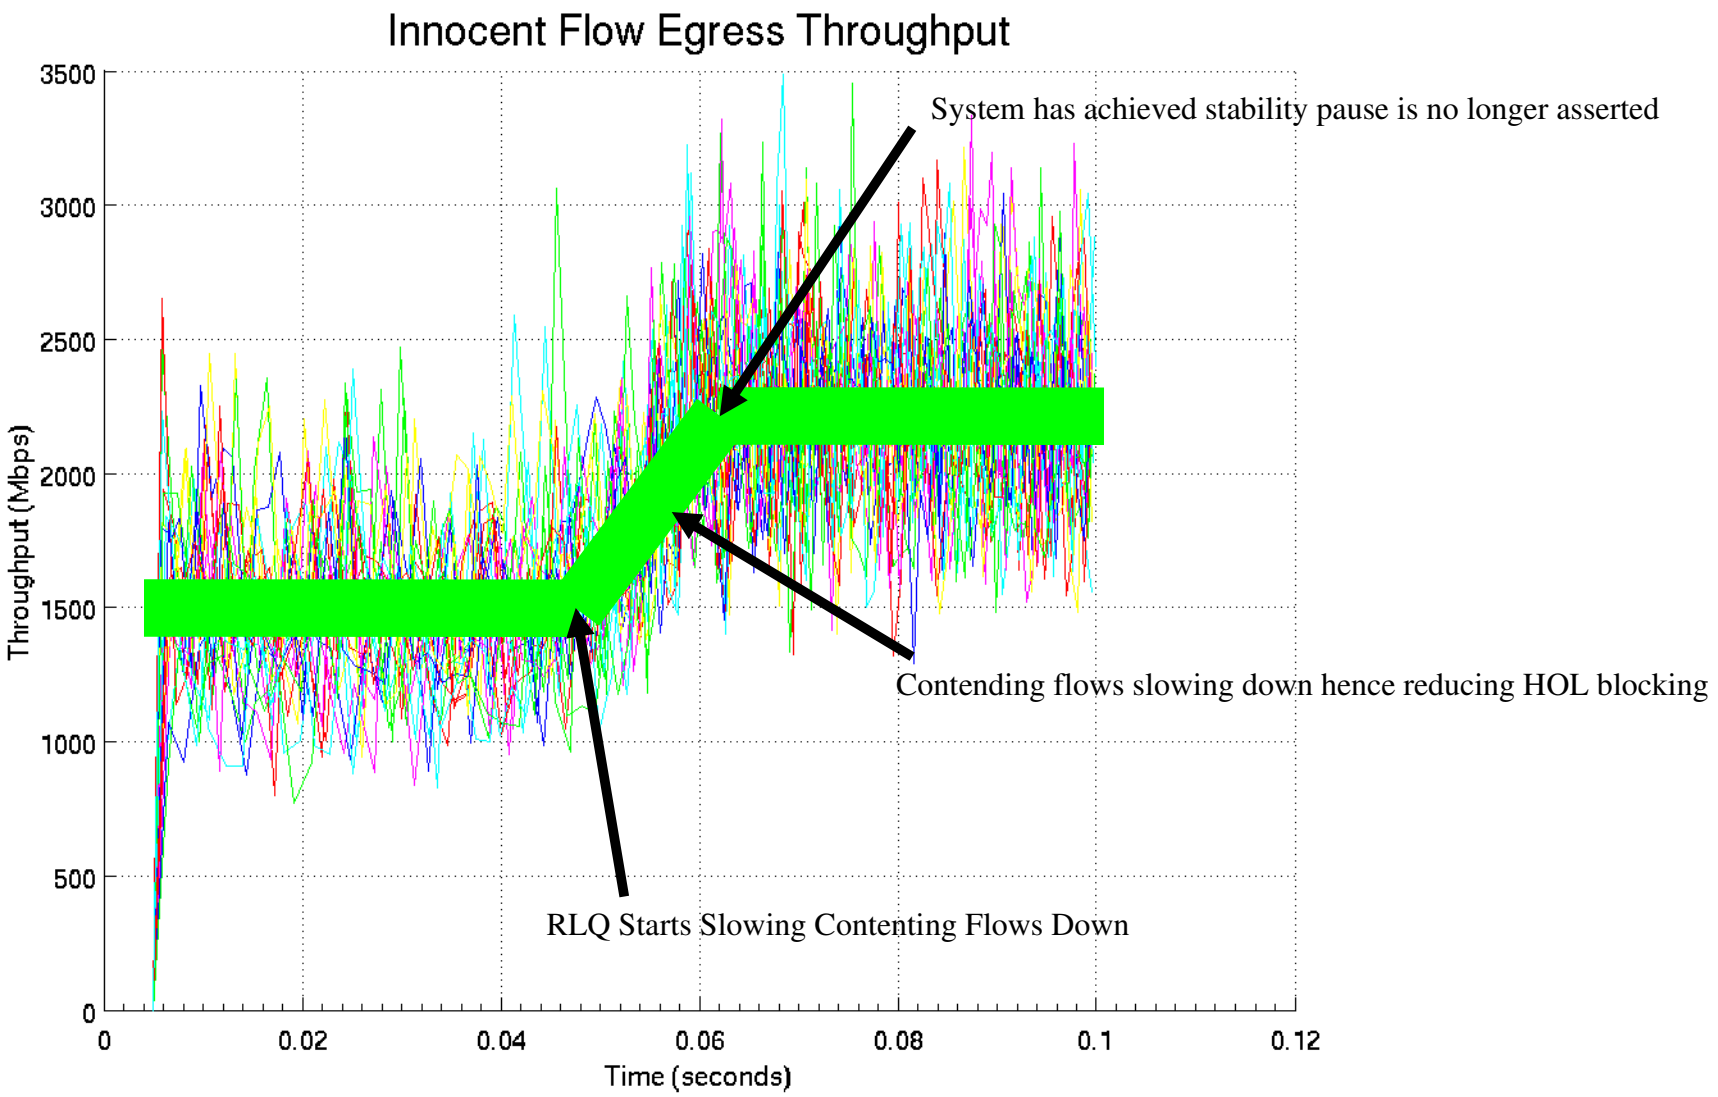

FULCRUM microsystems

Hotspot Degree vs Time To Stability

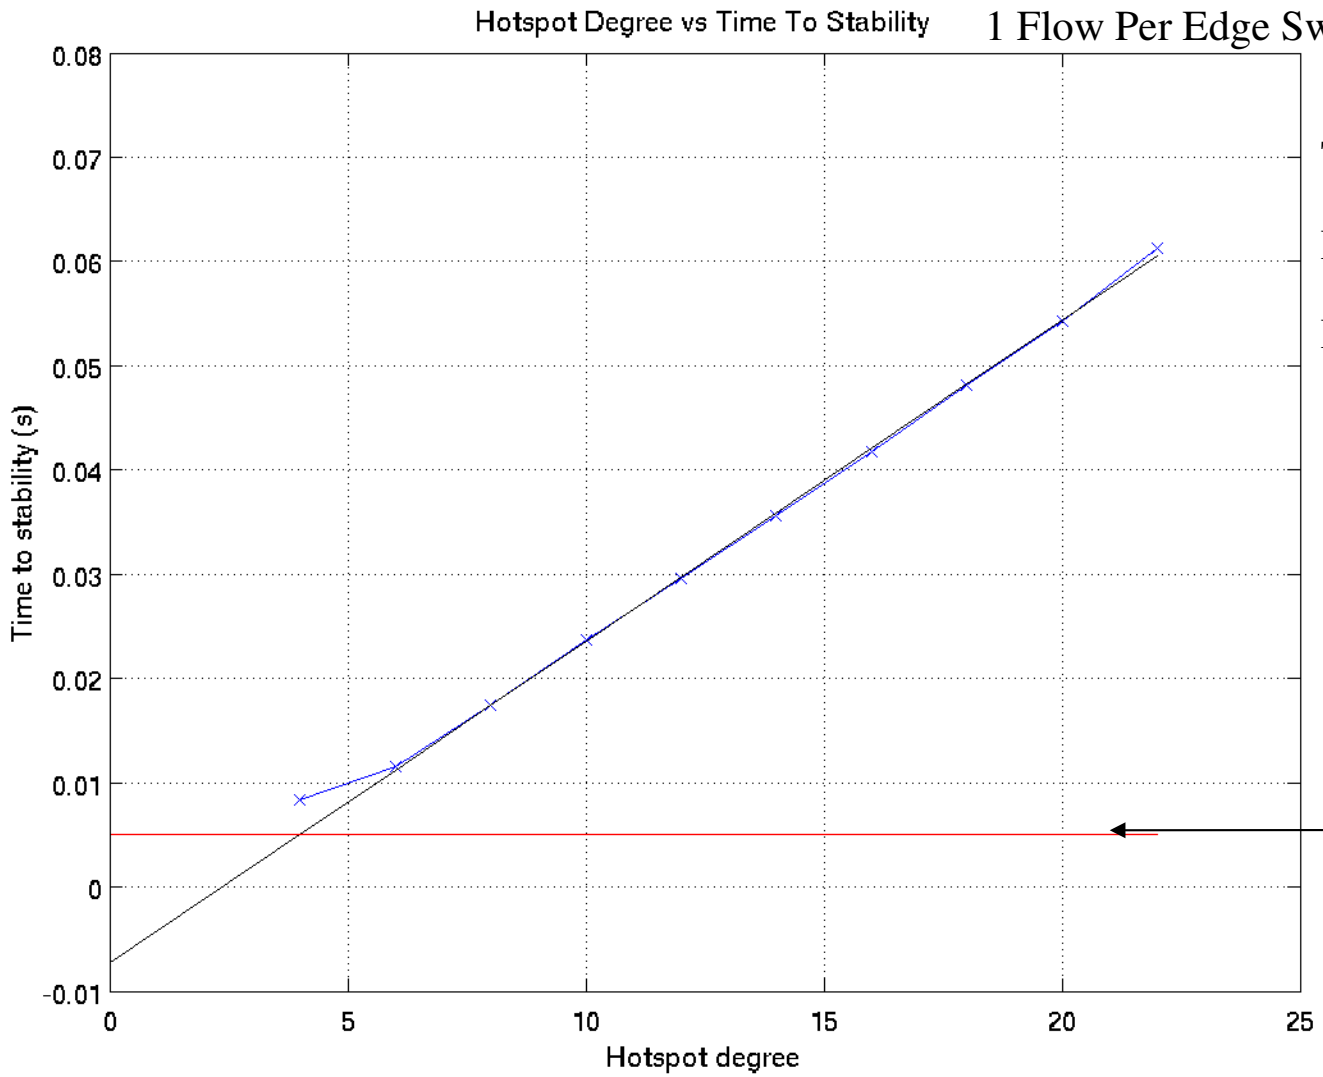

1 Flow Per Edge Switch

Time To Stability

Linear w.r.t. Hot Spot

Degree

Sim Start

Time 5ms

1 Run per datapoint

FULCRUM microsystems

Time To Stability (5 Flows Per Sw.)

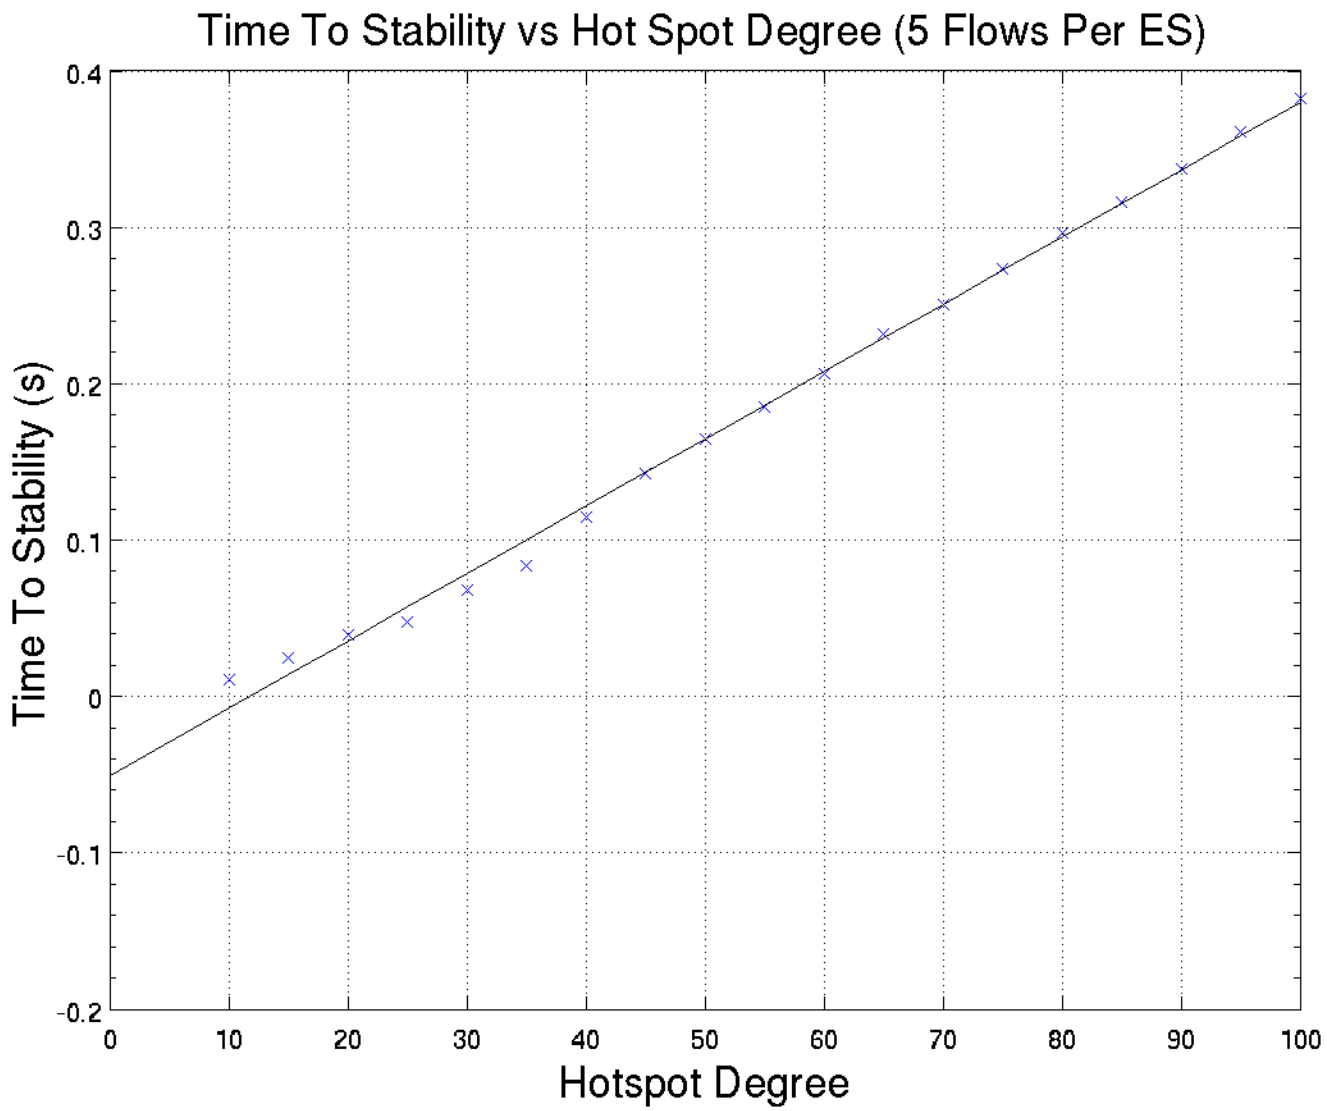

FULCRUM microsystems

Time To Stability (1 and 5 Flows per Sw.)

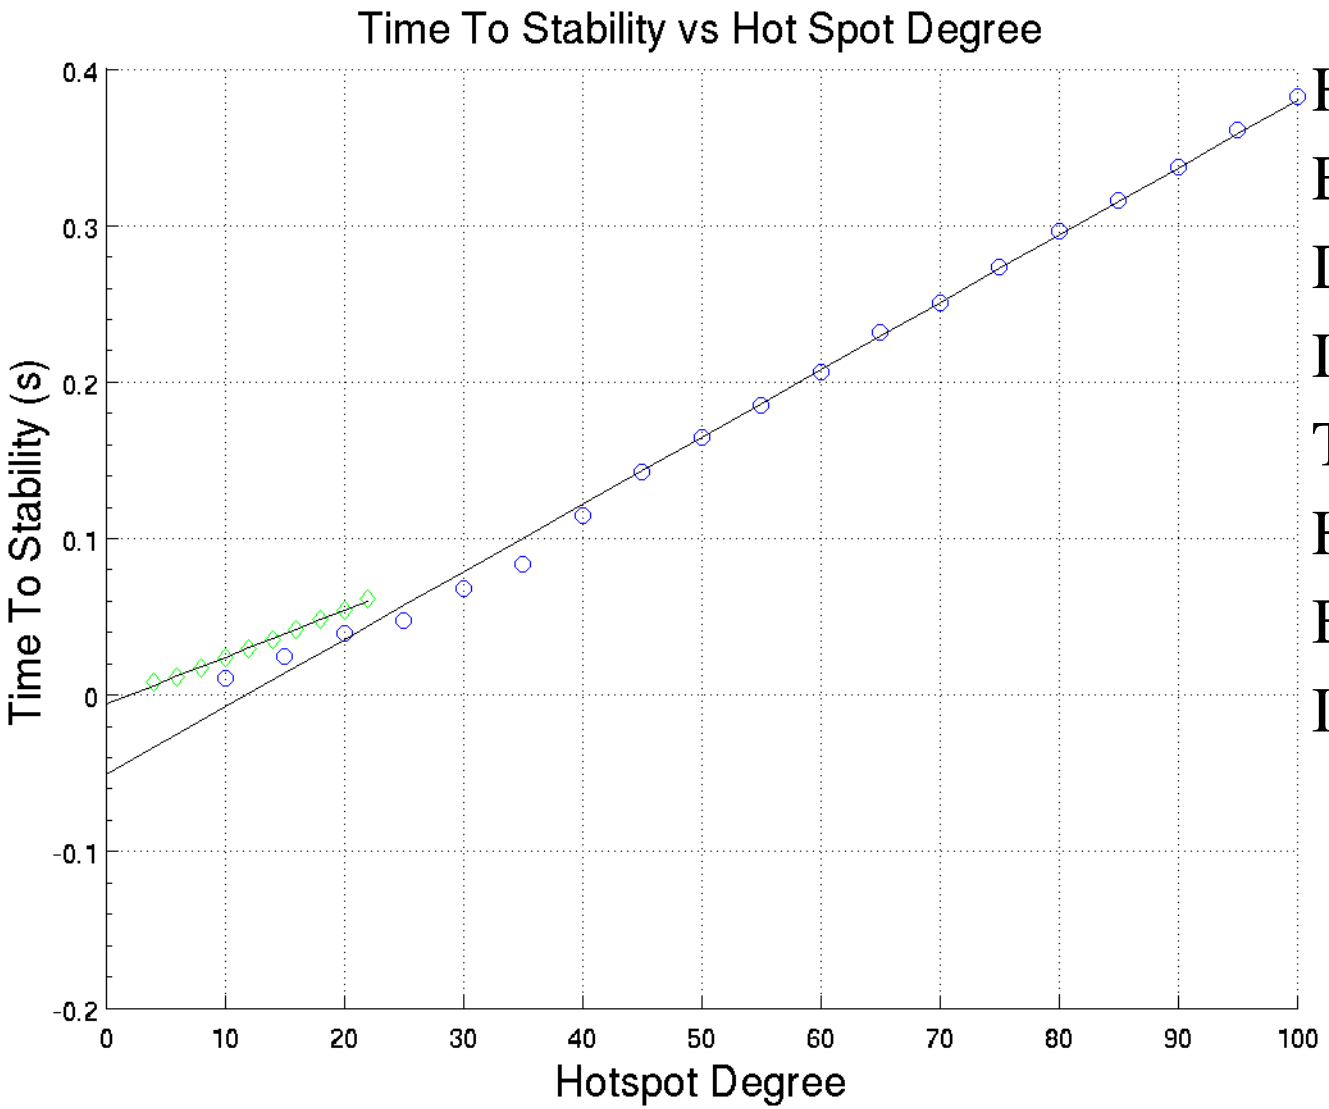

- HOL Blocking on Edge Switch Due to Pause Improves Time To Stability
- Hidden Cost Is In Backpressuring Innocent Flows

FULCRUM microsystems

Time To Stability (1, 5, 10 Flows per Sw.)

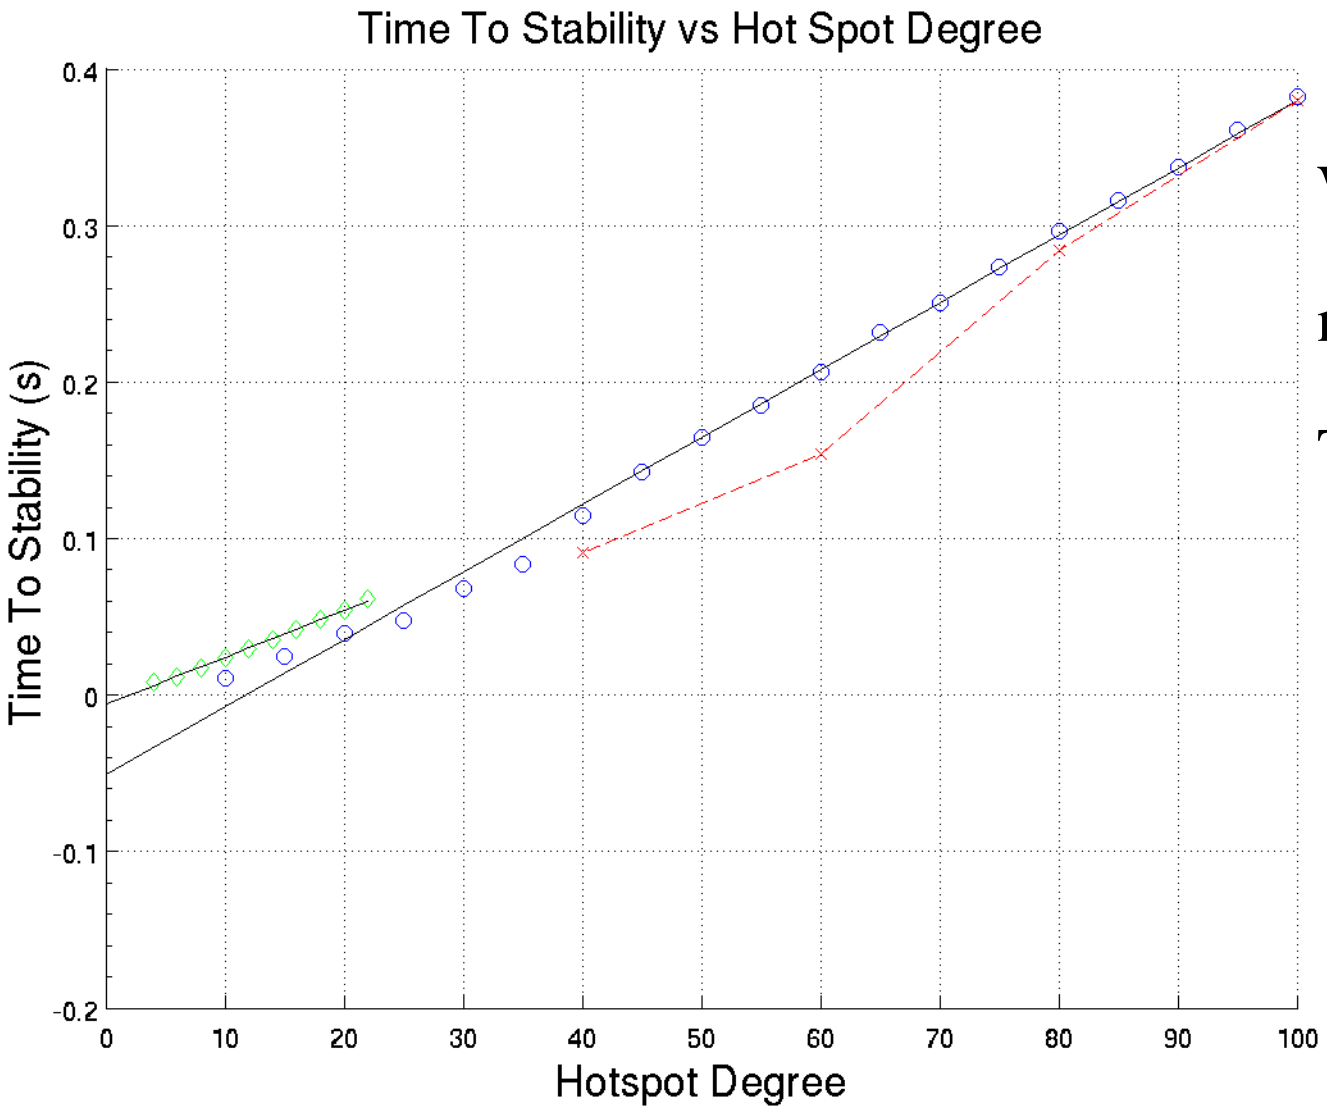

With large number of flows TTS converges

FULCRUM microsystems

Large Number of Flows: CP Occupancy

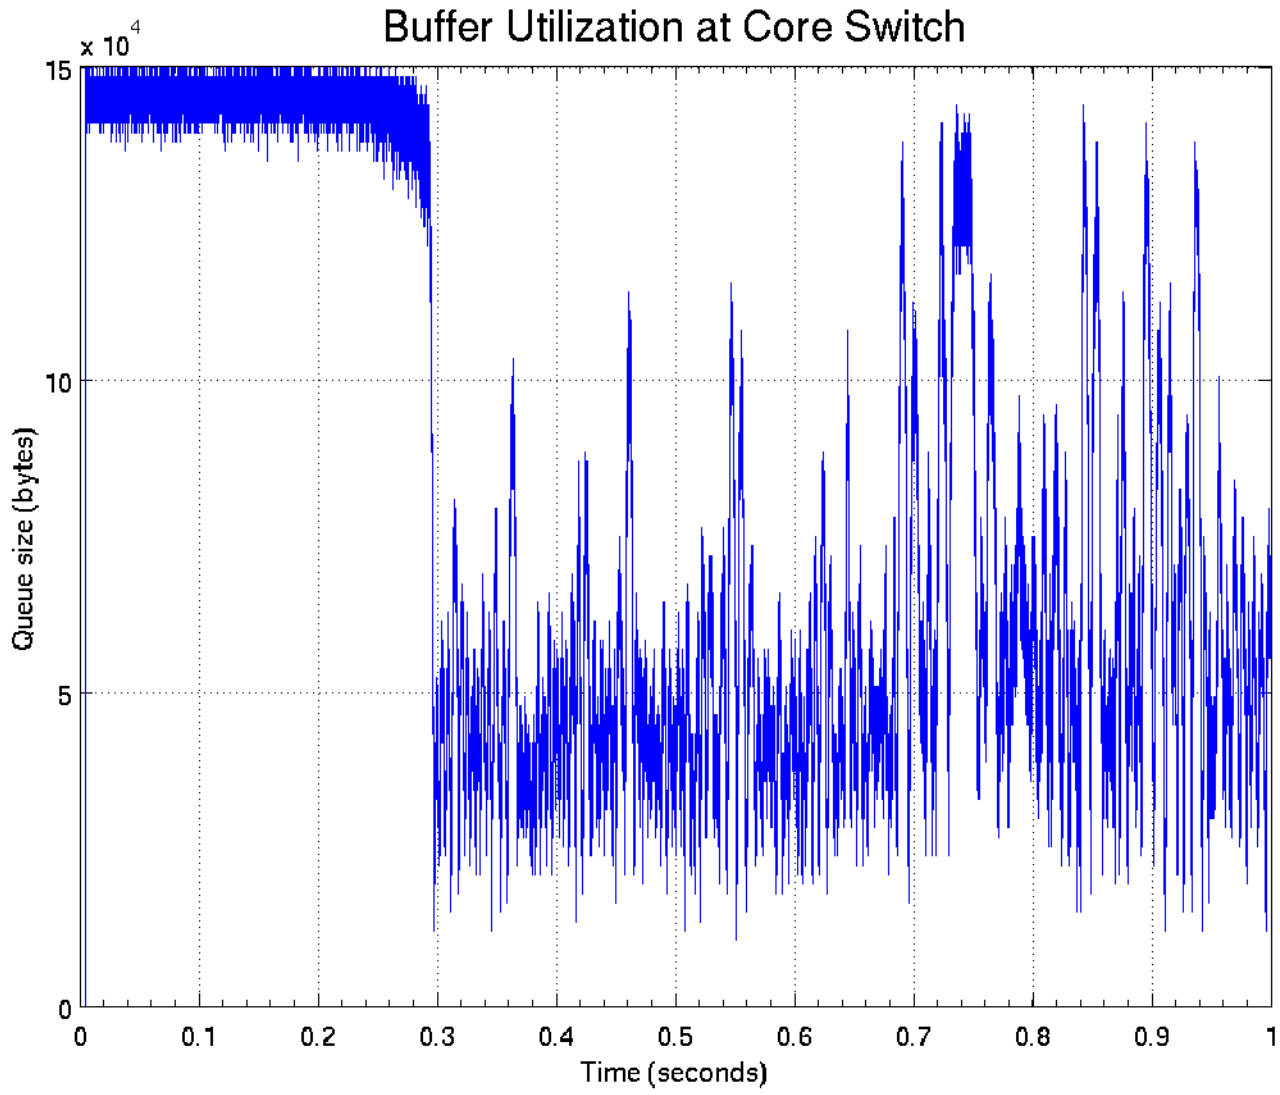

80 Flows.

System is on the verge of becoming unstable.

FULCRUM microsystems

Conclusion

- **For the baseline topology**
  - Time To Stability(TTS) largely linear with hotspot degree
  - BCN starts to become unstable with large number of flows
    - Perhaps this indicates a need for a heuristic to estimate number of flows?
  - HOL Blocking has subtle advantage in helping to reduce TTS although the cost to the innocent flows may be undesirable.

Future Work

- **Larger number of flows**
- **Different Topology**
- **Repeat with BCN variants, BCN(0,0), BCN-Max, BCN+-Max**