

# TCP Alternative Backoff with ECN (ABE)

[Naeem Khademi](#), Michael Welzl, Grenville Armitage, Chamil Kulatunga, David Ros, Gorry Fairhurst, Stein Gjessing and Sebastian Zander

TCPM WG - IETF 93, Prague – July 22, 2015

[naeemk@ifi.uio.no](mailto:naeemk@ifi.uio.no)



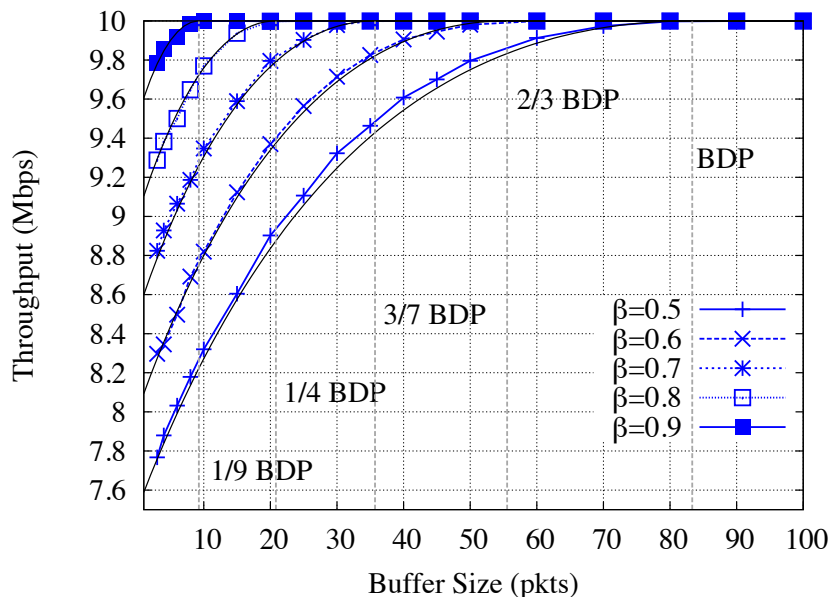
REDUCING INTERNET TRANSPORT LATENCY



# Internet's Latency Problem

- **With standard TCP CC's multiplicative decrease factor ( $\beta = 0.5$ ):** more throughput means more buffering (more latency)
  - 1 BDP for full utilization
- **AQM solutions:** aiming to keep the average latency low: (FQ\_)CoDel, PIE, ARED, ...
- AQMs' default thresholds translate into a tiny average buffer (e.g. 5ms~20ms for CoDel and PIE)

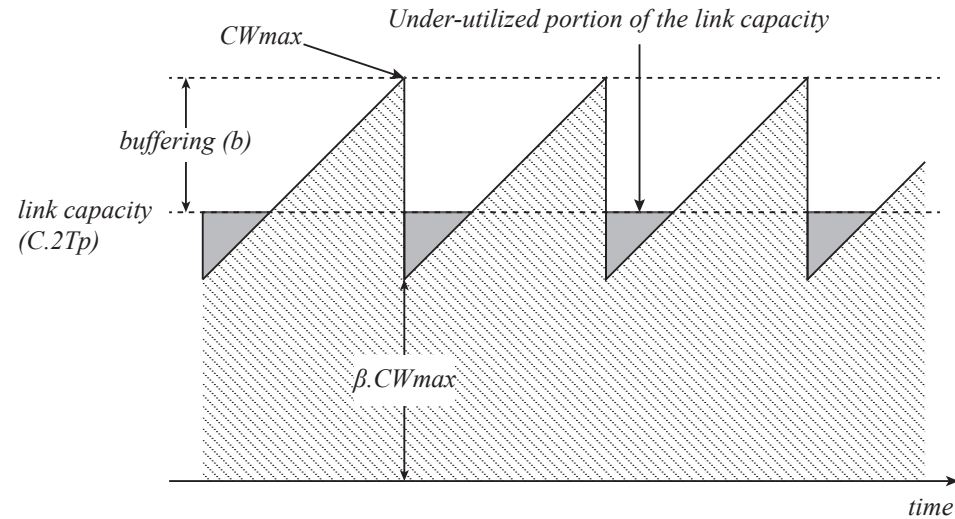
Simulation vs. model  
One NewReno flow @ 10 Mbps, RTT 100ms, DropTail



# The Problem with using AQMs

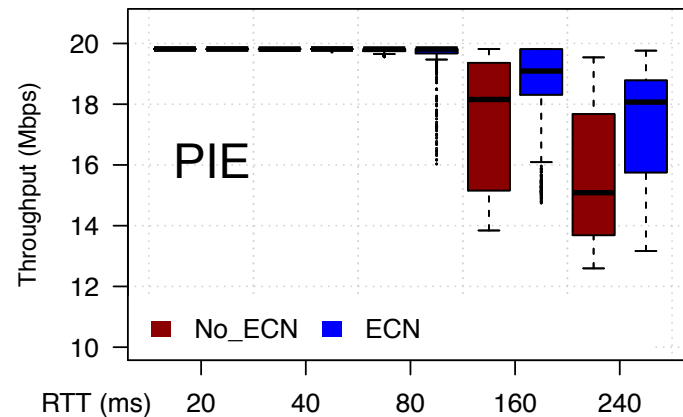
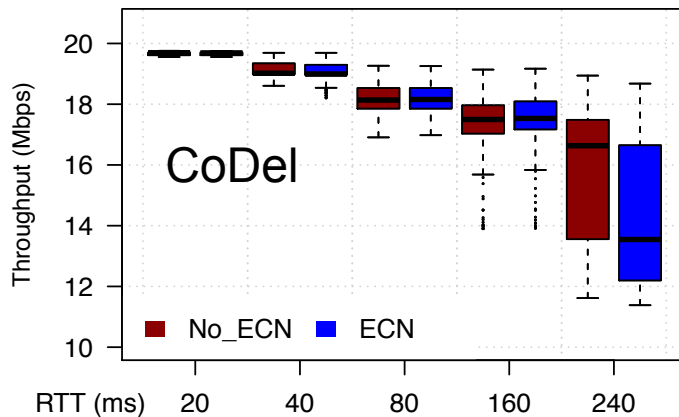
## Problem with long-RTT paths

- Link under-utilization with AQMs' default thresholds when RTTs head above 60ms~80ms

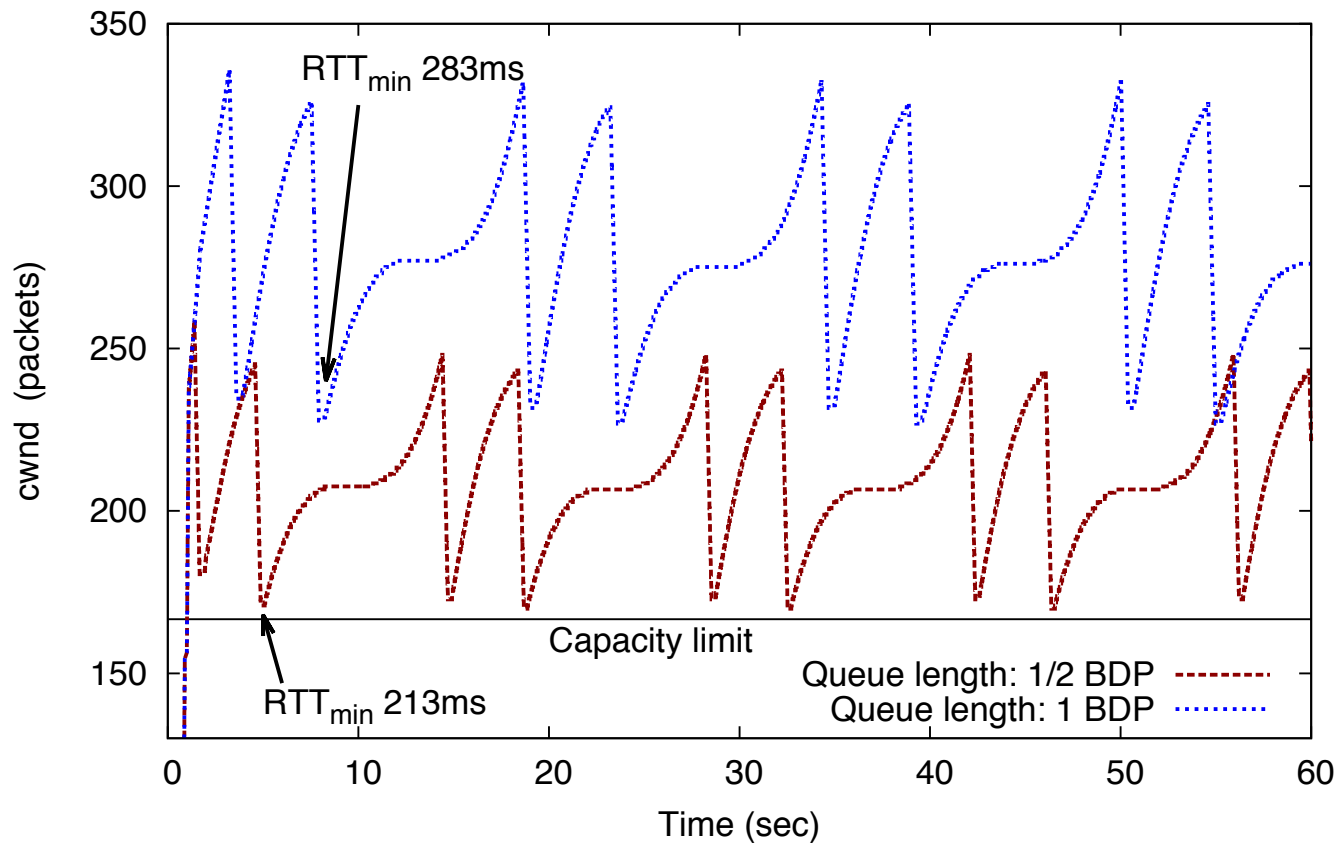


- Under-utilization still a problem when RTT increases

## One CUBIC flow @ 20 Mbps (real-life tests)



# Can we use a larger TCP $\beta$ ?


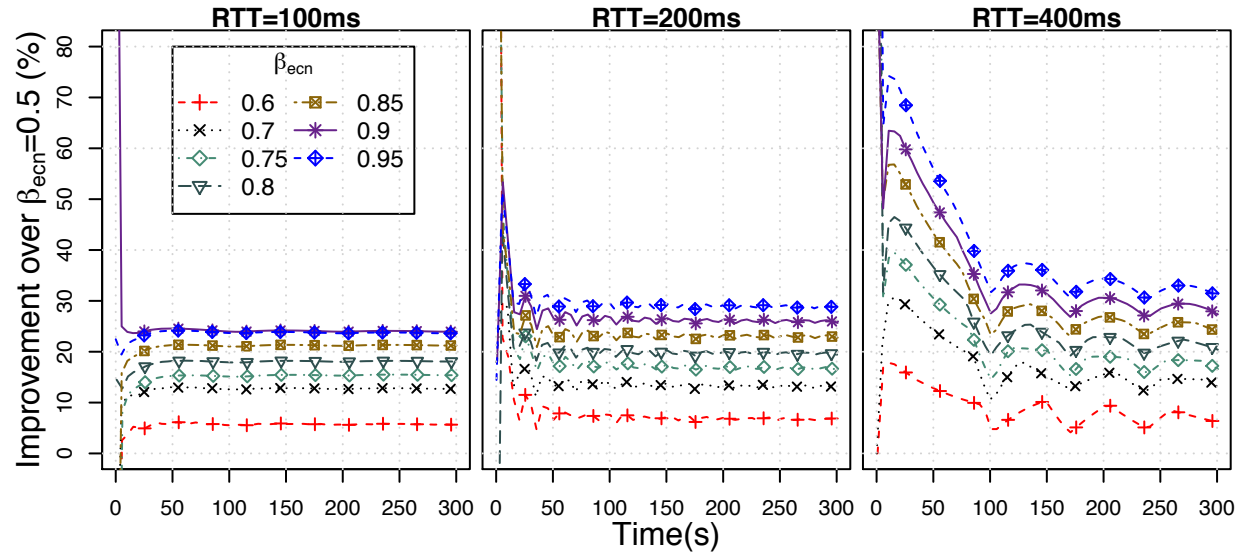


# Alternative Backoff with ECN (ABE)

- **Packet loss:** tail-loss (and other sources of loss) versus AQM-loss
  - Not a good indicator of the onset of congestion
- An **ECN-CE** signal *explicitly* indicates an AQM's presence (most likely to be (FQ\_)CoDel or PIE) with a low marking threshold
- **Alternative Backoff with ECN (ABE):** only a simple TCP parameter change
  - => use a  $\beta$  higher than 0.5 in response to ECN-CE ( $\beta_{\text{ECN}}$ ) and react with  $\beta=0.5$  in response to loss
    - Few lines of code in kernel (in the sender *only*)



# Latency vs. throughput tradeoff (#1)

One NewReno flow, CoDel @ 10 Mbps, RTT=100

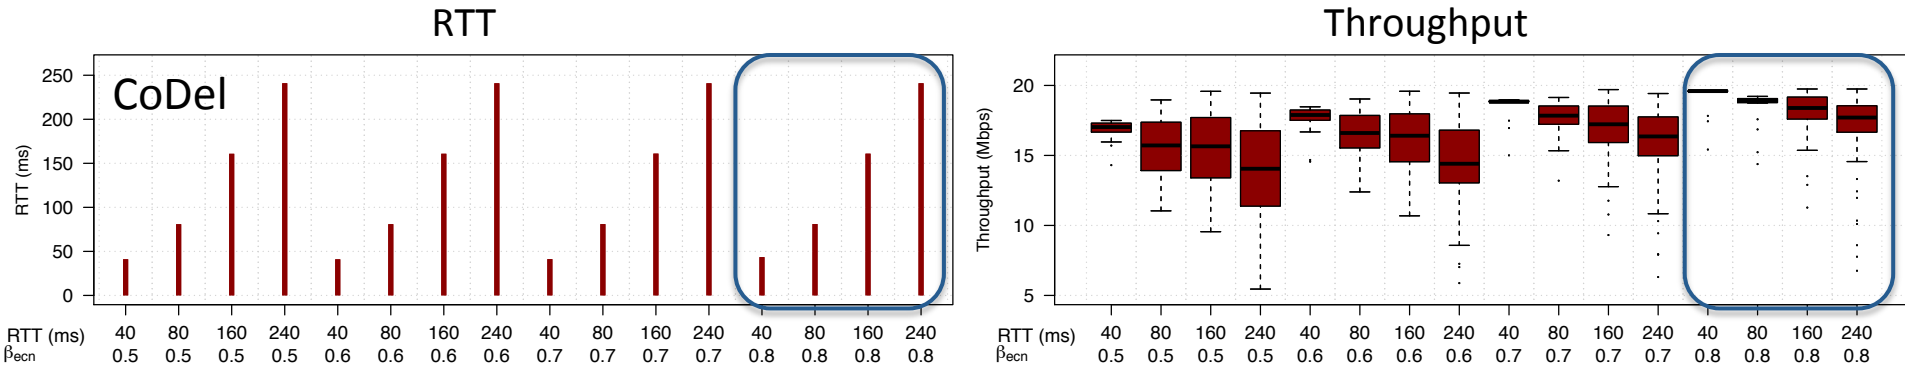


# Latency vs. throughput tradeoff (#2)

One NewReno flow, **PIE** @ 10 Mbps, RTT=100



# Latency vs. throughput tradeoff (#3)



Increased  $\beta_{ECN}$  causes insignificant change in median RTT

## ABE Performance


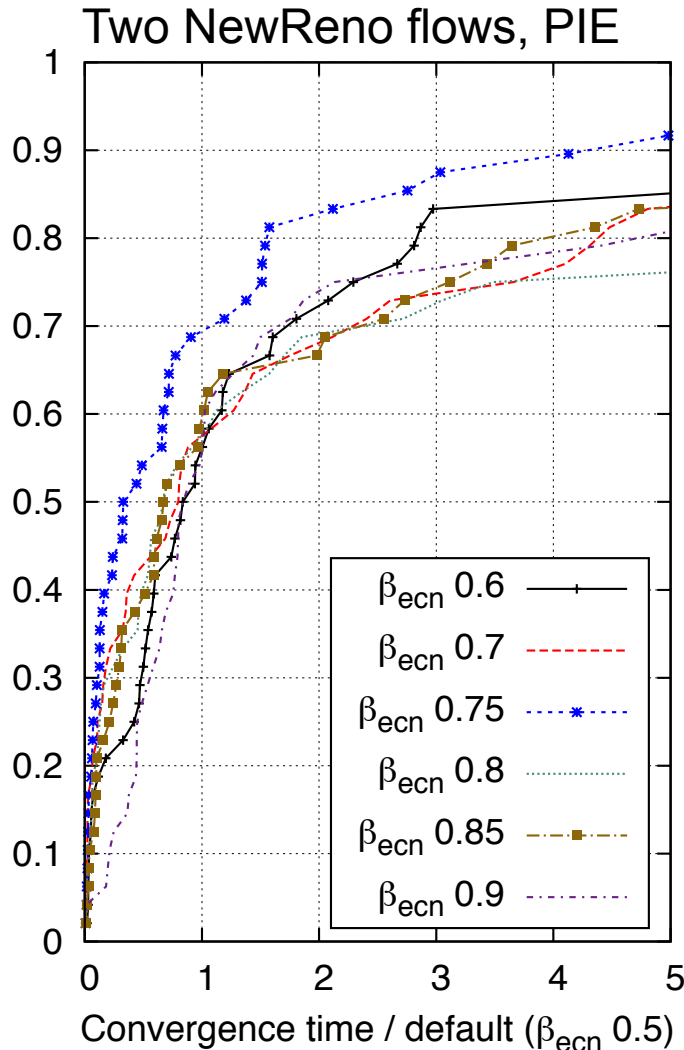
- Low latency and high utilization



Source: <http://frenchyme.blogspot.no>



# TCP Convergence Time




@ {1, 5, 10, 20, 40, 100} Mbps and RTT={10, 20, 40, 60, 80, 160, 240, 320} ms

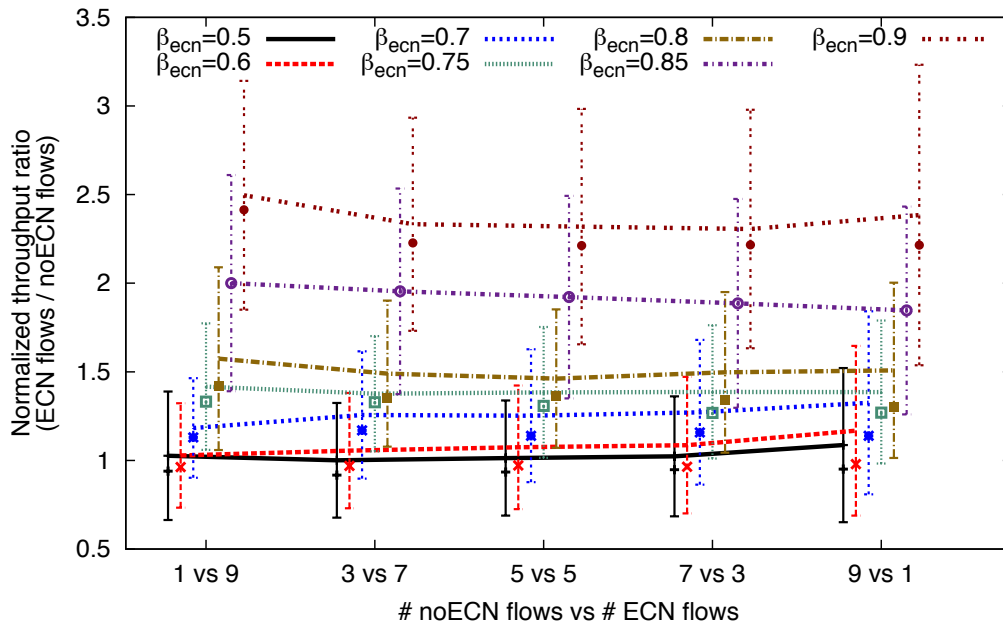
JFI of the cumulative number of bytes transferred per flow from the time a new flow entered, (after 30 sec during a 90 sec experiment) and noting the time it took until this index reached a threshold (0.95)

# Friendliness with standard (loss-based) TCP

## CoDel




## PIE



# ABE and slow-start (#1)

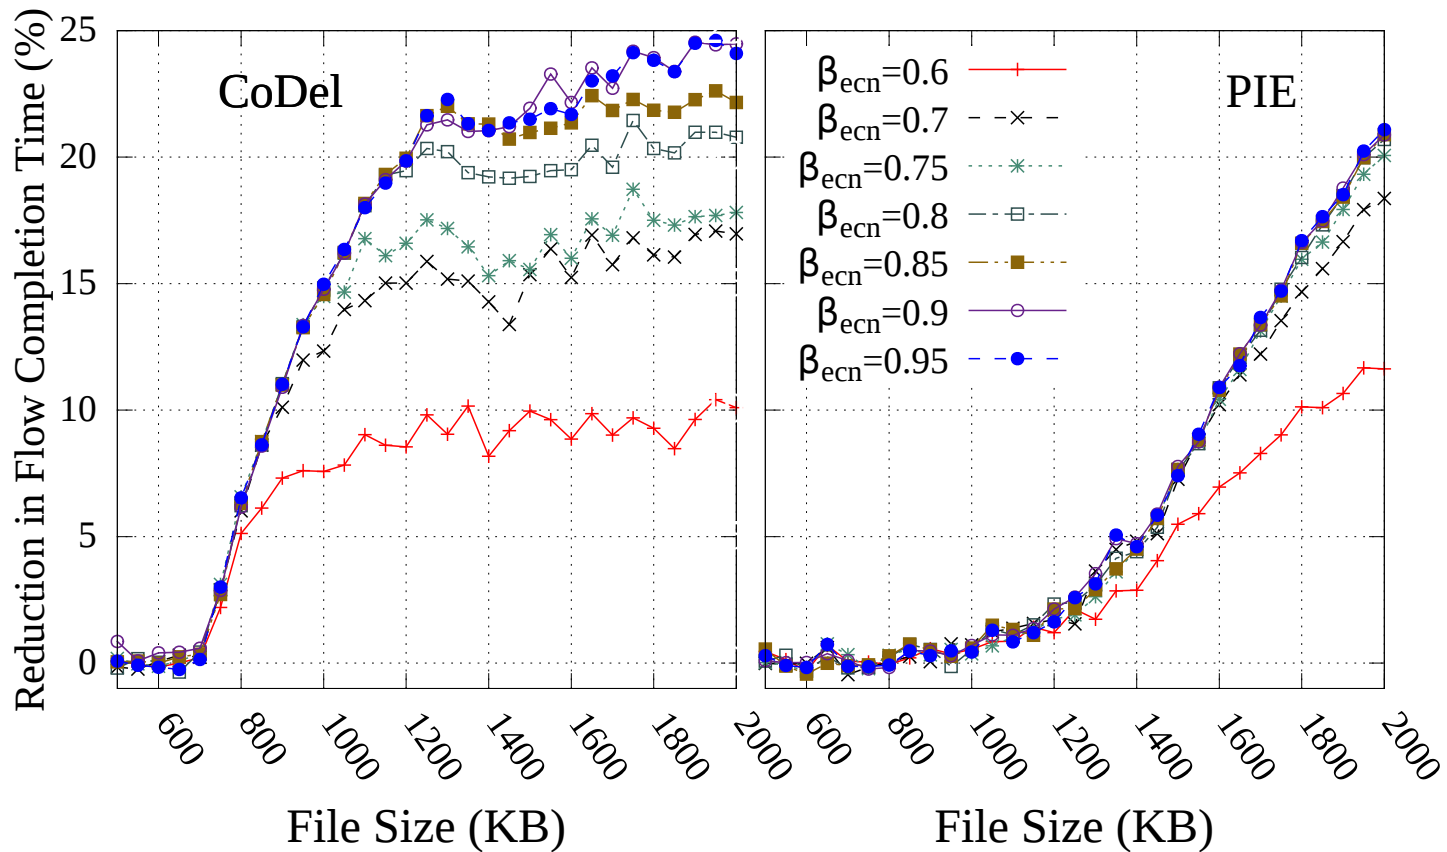
- Benefits short flows terminating right after SS
- E.g. reduction in PLT of large web-pages
- Number of reduction steps bound to  $\log_{\beta_{ECN}}(0.5)$

The effect of overshoot at the end of slow-start  
@ 20Mbps (experiment)



# ABE and slow-start (#2)

Percentage reduction in FCT (compared to using  $\beta_{ECN} = 0.5$ ) with different flow sizes @10 Mbps link with RTT =100ms



# Conclusions

- **Alternative Backoff with ECN (ABE):**

Naeem Khademi, Michael Welzl, Grenville Armitage, Chamil Kulatunga, David Ros, Gorry Fairhurst, Stein Gjessing, Sebastian Zander, "**Alternative Backoff: Achieving Low Latency and High Throughput with ECN and AQM**", Technical report 150710A, 10 July 2015, Swinburne University of Technology, Australia

<http://caia.swin.edu.au/reports/150710A/CAIA-TR-150710A.pdf>

## draft-khademi-alternativebackoff-ecn

- A minor sender-side modification (changes  $\beta_{\text{ECN}}$ )
- Complies with RFC 3168
- Incremental deployment with no flag-day!

### ABE Performance

- Significant throughput gain with lightly-multiplexed traffic
- Low latency (using CoDel or PIE)
- No starvation of standard TCP
- Reasonable convergence and fairness with recommended  $\beta$ s

# Q&A