

Some Congestion Experienced

Some Congestion Experienced (SCE)

An RFC3168 backward compatible approach to high fidelity ECN signaling and congestion control

<https://tools.ietf.org/html/draft-morton-taht-tsvwg-sce-00>

Pete Heist
Network Analytics
Rodney W. Grimes
FreeBSD

Jonathan Morton
Principal Architect: sch_cake and Cobalt AQM systems
TekLibre, LLC
IETF tsvwg, March 25, 2019

The problem

- Transports spend most of their time either:
 - Building a bottleneck queue → high delay
 - Under-utilising the path capacity → low goodput
- The network doesn't give them enough information!
 - Binary signal per RTT → oscillation
 - Bang-bang control theory


The solution

- Congestion should be explicitly signalled earlier than loss
- ECT(1) is unused, treated identically to ECT(0)
 - RFC 3168, actual measurements in the wild (RFC 8311)
- Rename ECT(1) as SCE, ECT(0) as ECT
 - Some Congestion Experienced
- Other ECN codepoints retain current meaning & semantics

Two Congestion Signals

- SCE (some congestion experienced)
 - High fidelity congestion information
 - 100% marking when CE threshold reached
 - Less marking means less congestion
- CE (congestion experienced)
 - Coarse congestion information
 - Continue using as defined now
 - May be encountered without SCE (existing middleboxes)

State Diagram



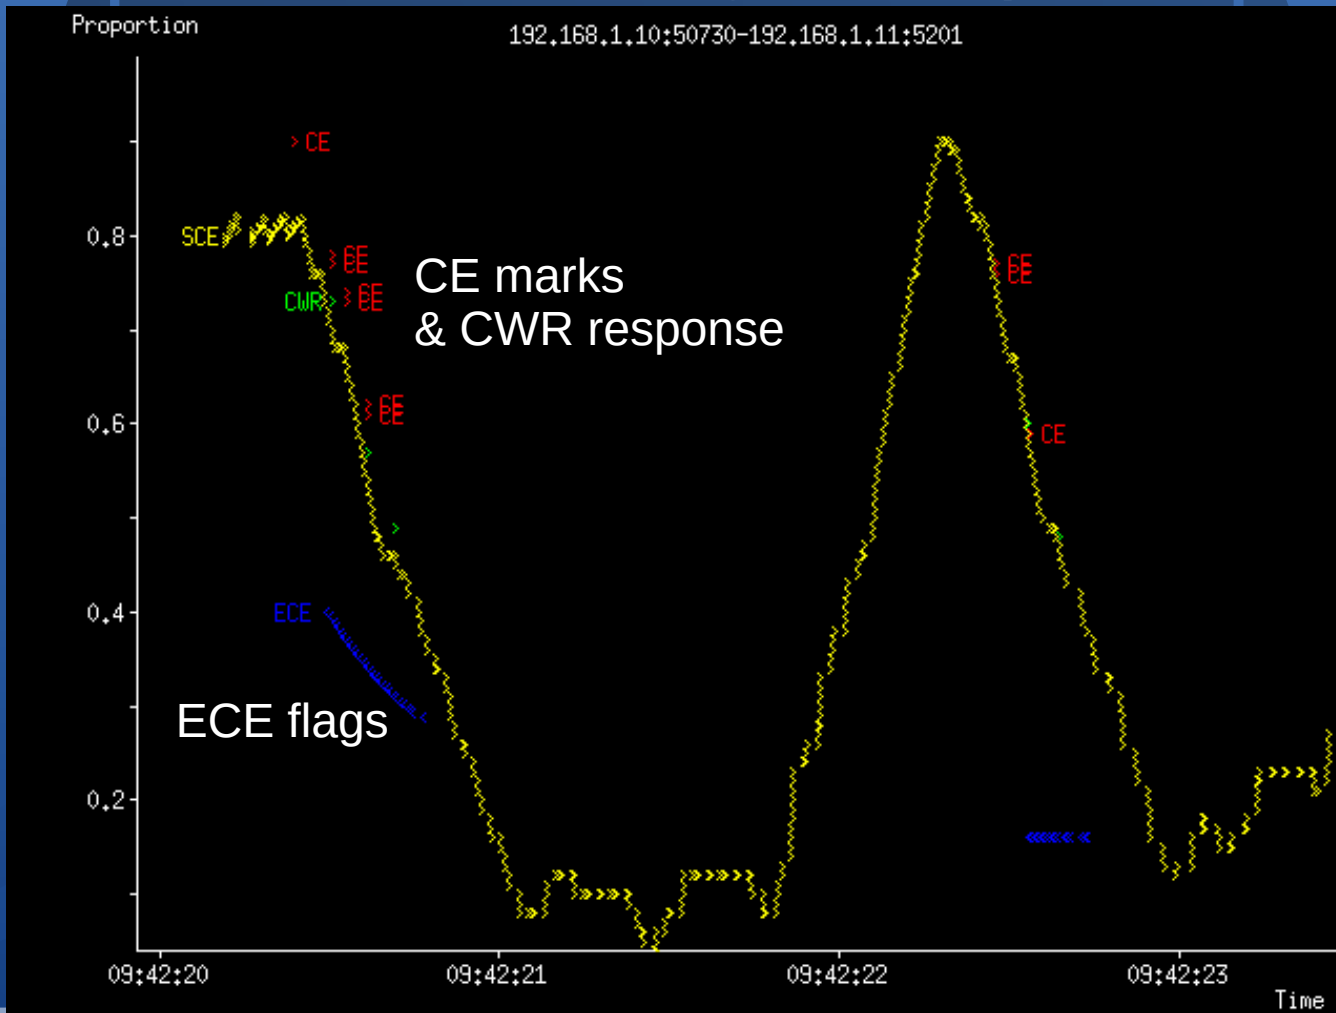
Not-ECT: Not ECN Capable Transport

ECT: ECN Capable Transport

SCE: Some Congestion Experienced

CE: Congestion Experienced

SCE marking example



Experimental Plan

- Experiment #1 – SCE marking basics – success!
- Experiment #2 – middlebox AQM tests - ongoing
- Experiment #3 – check our maths
- Experiment #4 – integrate with transport protocols
 - Requires a feedback path receiver → sender

More details on our web site at:

<http://www.bufferbloat.net/projects/ecn-sane/wiki/>