draft-kuhn-aqm-eval-guidelines-01

Nicolas Kuhn (ed) [1] Preethi Natarajan (ed) [2] David Ros [3]
Naeem Khademi [4]

[1]IMT Telecom Bretagne, France

[2]Cisco Systems, United States

[3]Simula Research Laboratory, Norway

[4]University of Oslo, Norway

AQM Evaluation Guidelines – consensus on purpose and objectives

- Purpose of this draft
  - Provide a set of recommended guidelines for evaluating AQM schemes – similar to RMCAT
  - how does this draft affect standardizing algorithms in the WG – should these guidelines be defined before WG takes up algos ? (if after, what is the purpose of this draft ?)
  - how does any draft discussing detailed test suites affect standardization ?
- Objectives
  - high level guidelines on what should be general considerations for evaluating AQM schemes
  - performance
  - stability/deployability
- this presentation : updates of the *.00 version presented at IETF89

Comments on draft-kuhn-aqm-eval-guidelines-00

IETF89's (+ emails') feedbacks

- the 00 version was a mix of *evaluation guidelines* and *evaluation suite*
- agree first on a general set of principles for evaluating AQM
- before attempting to specify the details of evaluation tests

Simplify the guidelines - which are not an *evaluation suite*

- in short, the guidelines will detail :
  - classes of traffic to consider (Web traffic, Video traffic, etc.)
  - various scenarios (set of principles to evaluate)
  - congestion control protocols
- but have less details in the experimental approach :
  - do not specify how to generate the traffic (Web traffic modeling, encoded video, type of games, etc.)
  - do not focus on a specific context : network characteristics (link capacities, RTTs) are not precised (except for RTT fairness scenarios)

Towards draft-kuhn-aqm-eval-guidelines-01 : remove/add

Remove

- queue-level metrics :
  - hard/impossible to obtain
  - E2E metrics show the impact of the introduction of AQM schemes
- QoE metrics
- scenarios related to the context (Wi-Fi, rural broadband, data-centers) : the guidelines will define set of aspects that will not depend on a specific context

Add

- discussions around the methodology for :
  - running experiments
  - comparing AQM schemes

New ToC

New ToC :

- Introduction
- End-to-end metric
- Generic set up
- Various TCP variants (Aggressive, Friendly)
- RTT fairness
- Burst absorption
- Stability (congestion levels, varying bandwidth)
- Implementation cost
- Control Knobs and auto-tuning
- Interactions with ECN
- Interactions with scheduling
- Methodology (AQM comparison, packet sizes)

End-to-end metrics

- flow completion time
- packet loss
- packet loss synchronization : degree of synchronization of loss events between two flows on the same path
- goodput : end-to-end appreciation of how the AQM improves transport and application performance - compulsory to evaluate the flow starvation when scheduling comes into play
- latency and jitter : AQM's capacity to reduce the queuing delay
- trade-off : these metrics help to evaluate the trade-off between reducing the goodput and maximizing the goodput
- these metrics should be considered for each scenario in the document

Generic set up

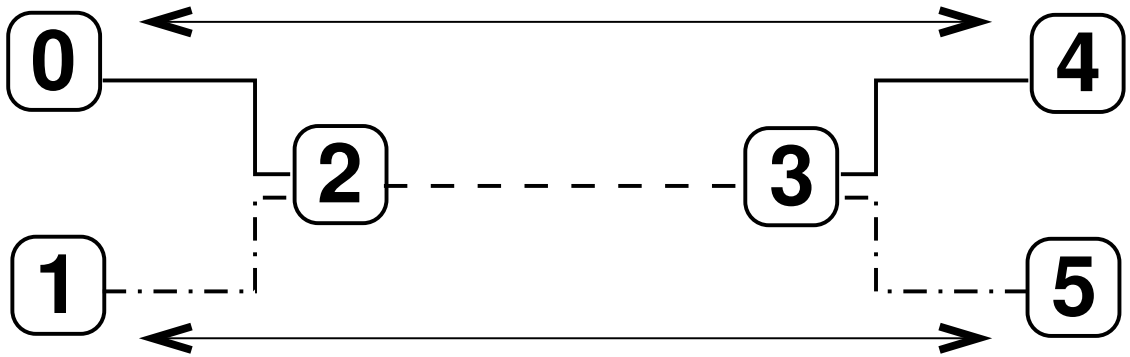

- classical dumbbell network - such as in "Common TCP Evaluation Suite"
- AQM in the router before the bottleneck (node 2)
- the (RTT,Capacity) of each link is independent from the others
- the links are supposed to be asymmetric
- size of the buffer should be carefully set, considering the bandwidth-delay product
- three classes of congestion controls : TCP-friendly (example : NewReno), agressive (example : Cubic) and Less-than Best Effort

Various TCP variants

Why various TCP variants scenario ?

- applications run over different flavors of TCP (unresponsive flows or aggressive flows)
- AQM should ensure queuing delay is under control with traffic profiles

Details on the various TCP variants scenario

TCP-friendly Sender	TCP New Reno or others
aggressive Transport Sender	TCP Cubic or others
unresponsive Transport Sender	UDP flows
TCP inigial congestion window	mix of TCP New Reno, TCP Cubic with IW3 and IW10
traffic mix	TCP transfer, HTTP traffic, VoIP, Gaming, CBR, adaptive video streaming

RTT fairness

Why RTT fairness scenario ?

- AQM must be evaluated against a set of RTT as the capability of AQM schemes to control the queuing delay depends on the way end-to-end protocols react to congestion signals
- asymmetry (various RTT) SHOULD be considered : fairness between the flows might be improved with AQM

Details on the RTT fairness scenario

- To evaluate intra-protocol fairness :
  - flows with RTT in [5ms ;200ms]
- To evaluate inter-RTT fairness :
  - flows with RTT of 5ms
  - flows with RTT in [5ms ;200ms]

Burst absorption

Why burst absorption scenario ?

- bursty packet arrivals
- AQM scheme may bring bursts under control by balancing between (1) minimizing the queuing delay spikes and (2) minimizing the performance penalties (losses) for ongoing flows

Details on the burst absorption scenario

The traffic should be a combination of :

- one CBR (UDP) traffic
- repeating TCP transfer (IW10)
- HTTP web traffic
- bursty video frames

Stability

Why stability scenario ?

- varying operating conditions (time of day or deployment scenario)
- stability of the AQM's parameters over time is challenged

Details on the stability scenario

- mild congestion
- medium congestion
- heavy congestion
- varying available bandwidth

Parameters sensibility and stability analysis

- AQM proposals SHOULD provide background material to ease the understanding of the AQM control law
- or could use other ways to discuss its stability

Implementation cost

Context

- AQM's successful deployment is directly related to its ease of implementation

Requirements

that help identify costs associated with implementing the AQM on a particular hardware or software platform

- AQM proposals SHOULD provide pseudo-code for the complete AQM scheme, highlighting generic implementation
- AQM proposals SHOULD highlight parts of AQM logic that are platform dependent and discuss if and how AQM behavior could be impacted by the platform

Control knobs and auto-tuning

Context

- AQM scheme's safety is directly related to its stability under varying operating conditions (varying traffic profile, fluctuating network conditions)

Requirements

- an AQM scheme should be stable in varying conditions without the need for external tuning (employing auto-tuning if needed)
- an AQM scheme should minimize the control knobs exposed for operator tuning, to be more user-friendly and easier to deploy and debug

Interactions with ECN

Requirements

- An AQM scheme SHOULD support ECN
- An AQM SHOULD leverage ECN as an initial means to control queuing delay before resorting to packet drops
- An AQM scheme SHOULD self-adapt and remain stable even with faulty and/or unresponsive ECN implementations

Interactions with scheduling

Context

- there may be confusions whether a scheduling scheme is added on top of an AQM or is a part of the AQM - the guidelines consider that AQM is an hybrid scheduling/dropping policy algorithm
- a router may schedule the transmission of packets in a specific manner by introducing a scheduling scheme

Requirements

- the tester MUST discuss the feasibility to add scheduling on top of its algorithm
- this discussion MAY detail if dropping policy is placed while packets are enqued and dequed

Methodology

Replicating tests

- provide test setup, software and hardware versions
- no particular evaluation toolset
- data available

Metrics measurement

- no details on the way the metrics must be obtained
- consider queue-level metrics when possible

Comparing AQM schemes

- consider performance AND deployment for a fair comparison

Packet sizes and congestion notification

- an AQM should adhere to recommendations outlined in RFC7141 (no advantage to flows with smaller packets)

Discussion

- any more end-to-end metrics ?
- are we missing any performance evaluation ?
- are we missing any deployment evaluation ?
- we tried to remove all the "recommendation" aspect of this draft - are there any particular thing that we mention in the draft that should not be included in this draft (ECN, scheduling, etc.) ?
- next step : other drafts should detail the scenarios proposed in these guidelines for environment specific evaluations