



Summary of draft-zhang-pcn-performance-evaluation and draft-charny-single-marking

Anna Charny
acharny@cisco.com

Joy Zhang
joyzhang@cisco.com

Francois Le Faucheur
flefauch@cisco.com

Vassilis Liatsos
vliatsos@cisco.com

Outline and Scope

- **Performance comparison of admission control**
 - **Single-marking vs. draft-briscioe-style virtual-queue based admission**
 - **New results since last meeting**
- **Summary of Termination performance**
 - **The same algorithm for draft-brsicoe and single-marking**
- **Other single marking tradeoffs**
- **Impact of single marking on other drafts**
- **Note: recent proposals (draft-babiarz- and draft-westberg) need separate comparative evaluation**
 - **Work on alignment of performance criteria and simulation setups in progress**

High-level Results: Admission Control (Previously reported)

	Virtual Queue	Single Marking Admission
Marking Parameters	<ul style="list-style-type: none"> • Ramp vs. Step: no difference • Upper/Lower Threshold: relatively insensitive 	<ul style="list-style-type: none"> • Token Bucket Depth: Relatively insensitive
Absolute or Relative RTT Difference	No Effect on bottleneck utilization	No Effect on bottleneck utilization
EWMA weight and CLE	Insensitive	Relatively Insensitive

High-level Results: Admission Control (New in these drafts)


	Virtual Queue	Single Marking Admission
Ingress-Egress Aggregation	No effect	<ul style="list-style-type: none">Over-admission at low aggregation (synchronization effect)Sensitive to bursty flow arrivals at low aggregations
Multi-Bottleneck	<ul style="list-style-type: none">No effect at bottleneckUnfair to long-haul aggregates	<ul style="list-style-type: none">No effect at bottleneckUnfair to long-haul aggregates

High-level Results: Termination Control

Sanity check on SingleLink topology	Worked as Expected
RTT Difference	<ul style="list-style-type: none">• Absolute Difference: No effect• Relative Difference: Visible over-termination, though not significant
Multi-Bottlenecks	Worked as Expected, long-haul aggregates are more affected
Ingress-Egress Aggregation	Visible over-termination at low aggregation (synchronization)



Marking Synchronization

- **Cause: for periodic traffic and certain parameter combinations marking is not well distributed among flows sharing the bottleneck**
 - some flows are always marked and some are never marked
 - most relevant for CBR, but visible for near-CBR portions of other traffic types
- **Relevant only to excess-rate token bucket marking/metering when ingress-egress aggregation is low**
 - Detrimental to excess-rate admission: overadmission
 - Beneficial to termination: less over-termination than theoretical worst case


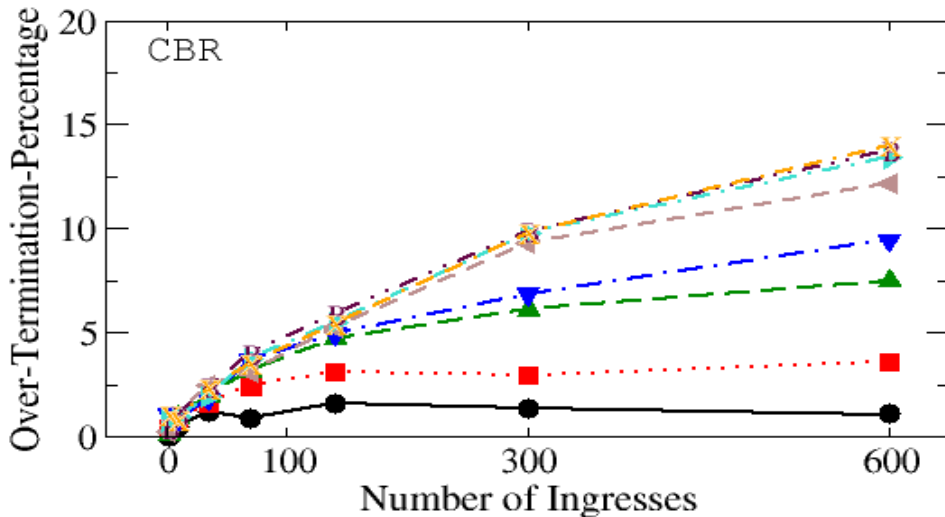


Is it a simulation artifact?

- **Probably Not!**




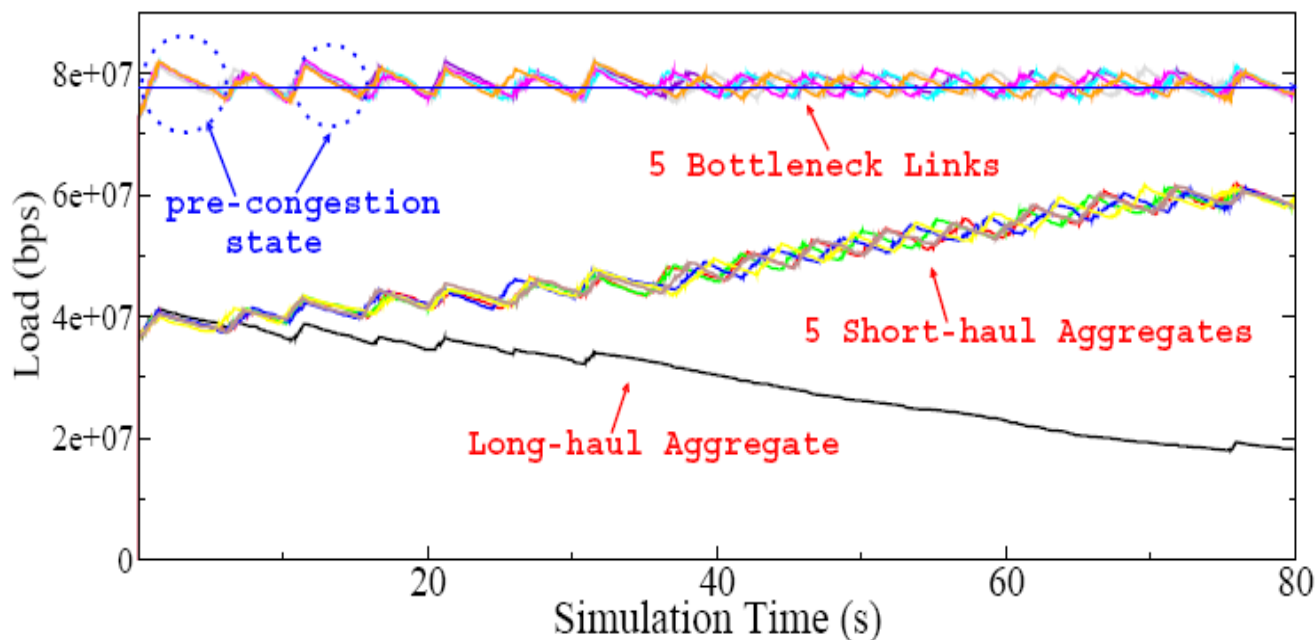
Single Marking Admission



Termination

Multi-Bottleneck Admission

Unfairness to long-haul flows (Beat-Down Effect)



- **Common problem for Single-Marking and VQ admission**
 - is a known property of many MBACs and likely to be a problem for all other admission control proposals
- **Is probably of limited practical worry**
 - for its effect to be significant needs large demand overload of long duration
 - overload is not large under "normal" conditions
 - in exceptional condition utilization control is more important than fairness

Conclusions of Performance Evaluation

- **Admission**
 - **At reasonable ingress-egress aggregations (~10 flows or more) performance of both schemes is comparable**
 - **At low ingress-egress aggregation single marking performance degrades**
 - **over-admission**
 - **sensitivity to call arrival assumptions**
 - **Both schemes unfair to long-haul flows for Multi-Bottleneck**
- **Termination**
 - **Performs as advertised in most scenarios**
 - **Occasional over-termination – but typically tolerable**

Single Marking Tradeoffs

- **Pros (or why we may want single marking)**

- Single codepoint

- One metering/marking scheme in the forwarding path of core equipment

- Easy (easier?) deployment path

- Can be viewed as an intermediate step for dual-marking

- **Cons (or why we should not do just single marking):**

- some configuration restrictions

- unclear how to do anti-cheating in multi-domain case

- lower accuracy in the multipath case

- some traffic engineering tradeoffs

- some performance tradeoffs at very low ingress-egress aggregation levels

Impact on PCN Architecture and Encoding Documents

- **Architecture**

- No changes to architecture required in the core**

- Addition of a single configuration parameter at the edges**

- already planned for inclusion in the architecture draft**

- Proposed optional renaming of marking-related terms**

- **Encoding**

- Any encoding choice suitable for dual marking works for single-marking**

- Allows additional encoding options**

- already in the current encoding draft**

What Next?

- **WG needs to decide whether single-marking needs to be allowed/accomodated in various WG documents.**
We argue – YES!
- **This is part of more general decisions: which/how many of the proposed approaches should be chosen?**
Need definition/performance criteria to aid the decision process (work in progress)
- **Smaller decision: should naming of the marking be associated with function (admission/termination) or semantics (excess rate/queue threshold/excess-rate-proportional/other).** *We argue naming should reflect semantics to help interoperability if more than one marking algorithm can support a given function*

CISCO SYSTEMS



Relationship to Other PCN Proposals

- **Core functionality**

 - a subset of core functionality needed for draft-briscoe**

 - a special case of core functionality required for termination of draft-babiarz**

 - coexistence with draft-westberg to be understood**

- **Edge functionality**

 - a trivial superset of the edge functionality of draft-briscoe**

 - differs substantially from edge functionality of draft-babiarz and draft-westberg**