IETF96 RTP Media Congestion Avoidance Techniques (rmcat) working group
Tuesday 19 July, 2016, 10:00-12:30

Update about working documents and way forward to complete them (chairs)

Varun Singh: Is there any need for CC-application interaction doc? Been struggling to write doc since app taxonomy is quite different from CC taxonomy. Asks the room, is it useful?
Varun: We probably need to do it.
Zahed Sarker: working on framework doc, says there's an interaction between CC and app. If there's any more than a paragraph of detail it should go in this separate doc.
Sergei Mena: Cisco is trying to open source a framework where all can be compared. We are not sure of this app/CC interface so this doc would be useful.

Coupled CC
Michael Welzl: What is needed here?
Chairs: Doc needs explaining what would be needed to move from Experimental to Proposed Standard.

Chairs: Same thing needed for Shared bottleneck before WGLC

Eval Criteria
Varun: Need to define what background UDP means. Do we need to specify a model?
Sergio: experimenting with both the initial background UDP and new one. Using background UDP to model trunk capacity. Both drafts should mention loss model and background UDP.
Zahed: loss model needs to be defined in eval criteria
Varun: needs feedback of whether background UDP as specified is working
Stefan Holmer: Is background UDP something that we are going to require? That test cases are run twice, once with background UDP and once with various capacity? I.e. duplicating test cases?
Sergio: Yes, this is what we are doing now, we see some differences between the two.
Zahed: Currently it says you should do with both, and since we see difference, and it sounds like it's useful.
Stefan: This is useful _if_ an algorithm breaks in capacity but not in background.
Varun: Could add this as a separate test case
Mirja Kuehlewind: if you have UDP background traffic you will drop UDP packets, otherwise not. Useful if you want to evaluate how quickly, if you want to evaluate impact on background traffic too. Either use background traffic entirely, or just have a test case for this.
Zahed: If the motivation is what is the impact on the background traffic, we don't need new test cases, just specify it's done like this.
Sergio: Repeat one test case using background traffic.
Mirja: only reason not to have background traffic is it's easier. It's much more realistic to have it. Changing capacity is far less likely.
Zahed: Wirelsss test case is the capacity jump example. What does it give us to have analysis on background traffic?
Mirja: No, it's about how you adapt if the background traffic changes. Do you adapt quickly or not? Will make a very big difference. Two separate test cases will solve this.
Varun: Endorse second way - need separate test case.
Magnus Westerlund: ECN is not in eval criteria, but it's in the charter to be considered. Do we need some ECN cases?

Varun: Send text!
Magnus: Will take a look.

Chairs conclusion: Will have a background model and some new test cases, including those for ECN.

Zahed: there's a "potential test cases" section with ECN mentioned.
Magnus: I looked in criteria!
Zahed: No, don't look in criteria! Look in test cases.
Varun: needs more thought on what this means.
Colin Perkins: It's not as simple as just turning on ECN. Should be separate to loss-based testing, it's sufficiently big and different.
Zahed: What is the behavior you want from ECN? How do you evaluate if it's good or bad?
Michael R: What should we assume the marking means? We can report what happens, but we don't know what should be happening!
Gorry Fairhurst: make sure our feedback feeds back ECN marks so they are useful. Separate out ECN stuff and get it moving. Should be in a centralized, separate doc.
Zahed: Agree with Gorry.

Chairs: Agree kept as separate item for now.

Wireless tests: no updates. Ready for reviewers.
Zahed: LTE cases ready last year.
Sergio: actively working on wireless, no issues with tests themselves.

Chairs: no updates planned, need volunteers for reviews.

eval-test
Zahed: has gone through reviews. Still experimenting with test cases given there. TCP model needs to go to eval-criteria.
Zahed: We should not do over engineering here to make our candidate results better, they need to be realistic.

GCC
Gaetano Carlucci: Paper and slides presented at ANRW 2016 will be sent to list
Zahed: algorithm is implemented in Chrome. Is it true that if draft is implemented it will be the same as Chrome?
Stefan: supposed to be the same. Don't know!
Stefan: We have a startup phase in the code but isn't in the draft. Some work to do on the loss-based controller.

Feedback design team (Zahed)
Jorg Ott: concerned that we don't actually know whether we have everything we need right now, why develop a new concise one when you may need more? Will have to cope with multiple versions in the future of this feedback, it's calling for trouble.

Colin: we could develop a compression mechanism for RTCP at the RTCP level for multiple RTCP blocks. We don't have time to define new stuff now, just use what we have.

Michael Ramalho: we could certainly do more optimal things looking at what we have. But initial versions of NADA covered very similar items so this gives some confidence that we are not missing anything.

Jorg: too many XR blocks already due to what we have defined

Zahed: looking to make all CC algorithms interoperable by defining new feedback.

Colin: take it to AVT for longer discussion

Varun: thinking we can go back to 1.6s for feedback. If we say 0.8s is the max we can stick with Option 1.

Mo: Would you never NACK as a receiver?

Zahed: No.

Mo Zanaty: Expect to interact with current feedback?

Zahed: Yes. Should interact with TMMBR.

Magnus: Makes sense to separate feedback vs actual explicit requests.

Varun: sender should not infer retransmits from this signal - it could, but its input to its algorithm and not a request.

Jon Lennox: sender and receiver are different.

Jon: clock needs to be synced across all sources being reported on

Chairs: feeling of design team as to whether to continue?

Colin: yes we need to continue, e.g. timing resolution need to be decided

Michael R: interesting to know what DSCP it was received on? May be rewritten.

Zahed: not considered

Jon Lennox: As AVTEXT chair, if it went to AVT they'd have another design team with the same members!

Randell Jesup: yes should continue, more to cover. Any benefit to negotiating resolution of timestamps etc for extensibility, and what the CC algorithm needs? To avoid wasting bandwidth on stuff that algorithms don't need.

Chairs: conclusion is clear, continue Design Team

CC feedback draft (Colin)

Varun: the 5% mean you can't send anything else, e.g. NACKs etc (?)

Colin: NACKs would be on AVPF so probably fine

Varun: FURs too -

Colin: and it depends how often people want to send them

Colin: What else are people sending? Need to find out and re-run.

Jon Lennox: Many operating modes lower than 2.5Mbps, e.g. want to find out whether you can go from Audio to Audio + Video. Audio only would imply feedback every 3s which is completely implausible.

Colin: provide more numbers

Mo: Need much better reporting for low bitrates. Feedback assumes contiguous loss.

Colin: no, not true. Loss does not affect feedback format as defined, it reports loss.

Framework document (Zahed)

Varun: document is not following RTP Taxonomy, leads to confusion.
Zahed: have not run thought this yet.
Varun: Does CC run on specific 5-tuple? What about unbundled? What about rtcpmux?
Zahed: description should say it runs on 5-tuples but include information on what happens if you multiplex.
Varun: this is a good start but needs to be more rigorous.
Varun: not sure "streams" means RTP stream or ICE transport
Zahed: if we continue we need to go through taxonomy

Varun: this is a simplification, if the encoders are on different transport, should CC run on one or both?

Varun: should the output of the document be used to realign architectures and terminology?
Zahed: yes, that is the mission
Jorg: want to leave open what final goal will be. Modelling is an intellectual exercise, not sure what end result will be, don't want to specify end goal.

Chairs: Adopt now or after more updates? No hums for adopting now, weak hums for waiting for updates. Will wait for updates to revisit adoption.

Video Traffic Model (Sergio)

Sergio: Unable to find good sources of test sequences.
Mo: NETVC testing draft provides pointers to video conferencing sequences.

NADA Performance Evaluation (Sergio)

Charis: Expecting WGLC after next update.

Scream Update (Ingemar Johansson)

Zahed: mismatch is in implementation not spec?
Ingemar: yes

Chairs: Send comments on the list due to time.

Ingemar: asked for WGLC for draft-ietf-rmcat-scream-cc
Chairs: please add a section explaining what experiments would be needed to move from Experimental to Proposed Standard and submit a new update, after that the document should be ready for WGLC.

Adaptive FEC (Varun)

Mo: note well, two IPR declarations, no licensing terms

Varun: group not concerned
Mo: terms required to be given (Nokia, Polycom)

Zahed: did you see any difference between use of this on different algorithms? Would like to see as a generic mechanism, not tied to CC.
Varun: it was proven to be useful
Zahed: What do we need to do on the CC to implement?
Varun: depends on the CC algorithm
Jorg: up to the respective algorithm to define the implementation, e.g. at what point to send FEC or not. It is algorithm-specific
Michael R: no problem with this doc being an individual doc. Need IPR specificity before any formal adoption. What kind of FEC to be used?
Varun: up to individual implementations as to what kind of FEC algorithm is used (XOR, Reed Solomon, etc)