CoRE Working Group Internet-Draft Intended status: Informational Expires: September 10, 2015

# Qualifying Questions for a CoAP Advanced Congestion Control Scheme draft-bormann-core-cc-qq-01

## Abstract

CoAP ( $\underline{\mathsf{RFC7252}}$ ) comes with a conservative base congestion control scheme. Advanced congestion control schemes can be defined where better performance is desired for a certain area of application.

This document is a strawman for a set of questions that could be used in qualifying a CoAP advanced congestion control scheme.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <u>http://datatracker.ietf.org/drafts/current/</u>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 10, 2015.

#### Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>http://trustee.ietf.org/license-info</u>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in <u>Section 4</u>.e of Internet-Draft

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

<u>1</u> .	Introduction	2
<u>2</u> .	Area of application	2
<u>3</u> .	Protection	<u>3</u>
<u>4</u> .	Stability	<u>3</u>
<u>5</u> .	Scalability	<u>4</u>
<u>6</u> .	Range	<u>4</u>
<u>7</u> .	Scope	<u>4</u>
<u>8</u> .	Performance	<u>5</u>
<u>9</u> .	Concurrent traffic	<u>5</u>
<u>10</u> .	Evaluation quality	<u>5</u>
<u>11</u> .	IANA Considerations	<u>5</u>
<u>12</u> .	Security considerations	<u>5</u>
<u>13</u> .	Acknowledgments	<u>5</u>
<u>14</u> .	References	<u>6</u>
<u>1</u>	4.1. Normative References	<u>6</u>
1	<u>4.2</u> . Informative References	<u>6</u>
Aut	nor's Address	<u>6</u>

# **<u>1</u>**. Introduction

(See abstract.)

This document should be read in conjunction with more fundamental documents such as [<u>RFC2914</u>], [<u>RFC5405</u>].

The set of questions posed here cannot be deemed to be a set of acceptance criteria. The questions are broad enough that it is unlikely good research will be available to answer each and every single facet of them.

(The set of questions in the current version of the document is clearly just a start; this version is being published to elicit contributions.)

## **2**. Area of application

Q(1) Is the algorithm meant for general use? If not, can the scope/ area of application be defined in an unambiguous way? This is particularly important if some of the below questions only can be answered in a positive way for that area of application.

## 3. Protection

Q(2) Does this scheme really protect the network?

Answering this question requires realistic simulations (see <u>Section 10</u>). Generally, a single set of simulations should vary a parameter such as offered load, number of clients etc.

"Protecting the network" is not easily defined. Comparing the behavior to that of base CoAP ([RFC7252], Section 4.7) is an acceptable proxy. Indicators that might be evaluated include:

- Number of retries (or the related metric: energy per delivered bit)
- \* Number of spurious retransmissions
- \* Goodput/throughput ratio (average, burst)
- \* Settling time (e.g., reaction time to and recovery after a burst)
- Q(3) Does the protection rely on the self-protection of the underlying network? If that can be switched off, does the scheme still protect the network?

(Anecdote: An early version of CoCoA turned out to work well only as long as it was run over a MAC with exponential backoff.)

## 4. Stability

Beyond congestion collapse, there are other situations that a congestion control algorithm should try to avoid.

- Q(4) Are synchronization effects expected in CORE environments (also see granularity statement below); i.e., if an application tries to deliver an exchange at a predetermined point of time (i.e. temperature reading every 5 min), all stacks might "synchronize" into colliding with each other, and backing off in a lock-step fashion. That might result in unreasonably large RTOs in significant parts of the population of sensors; It might be good to have the binary backoff combined with some kind of dithering/randomization, in order to break such sync early on?
- Q(5) What is the expected (required, desirable) granularity of the RTT measurements? For an algorithm that has all the intervals

specified in seconds, an implementer not aware of this issue might choose a granularity of a second. If that is not intended, or RTT timer granularity should be a certain resolution (i.e. in the same order of magnitude of the lowest expected RTT), a hint on that might be good.

- Q(6) What is the expectation of the algorithm on the stability of the parameters of the network? How long is a history of measured RTTs expected to be useful in predicting the future?
- Q(7) If any mechanisms are adopted from other congestion control algorithms, what analysis has been undertaken to avoid known problems of those mechanisms (e.g., [<u>RFC6298</u>] will increase RTO when RTT decreases).

### 5. Scalability

Q(8) Do we have numbers for larger networks?

CoAP applications are expected to be run in networks with thousands of nodes (and even many more). At least some of the qualifying questions (and in particular protection) should be examined up to such a scale.

## 6. Range

Q(9) What is the range of parameters the scheme is supposed to cover?

Congestion control schemes need to adapt to a large range in each of the governing parameters such as latency, loss, and offered load. What do we know about the range actually being covered? (Note that it is quite acceptable for a scheme not to "use" the full range, e.g., not to be able to exploit very short latencies for improved performance.)

## 7. Scope

- Q(10) What is the scope of a single instance of the algorithm? (E.g., a five-tuple, a host pair, a single host and an IP address prefix with many peers?)
- Q(11) What is done to control the aggregate congestion behavior (cf. [RFC5405] section 3.1)?

### 8. Performance

Q(12) Is it worth it?

While improved performance certainly is not part of the acceptance criteria, deployers are unlikely to switch on a scheme that is worse than the default one.

Metrics might include goodput, latency (average, median, 95th percentile, etc.), goodput/throughput ratio, ... Again, these are best presented over a scale varying some input parameter.

## <u>9</u>. Concurrent traffic

While TCP fairness is both overrated and almost trivially achieved for what is basically a lockstep protocol, some information is desirable on how the scheme fares with concurrent traffic (such as base CoAP, TCP, or even inelastic UDP flows).

- Q(13) Does the scheme starve?
- Q(14) Does it do significant damage to the control algorithms of the concurrent traffic?

#### **10**. Evaluation quality

Of course, for all simulations and experiments, we need to know more about the models and environments used. Ideally, the evaluation would not fail the criteria in [incredibles].

## **<u>11</u>**. IANA Considerations

This document makes no requirements on IANA. (This section to be removed by RFC editor.)

#### **<u>12</u>**. Security considerations

The security considerations of [RFC2914] apply.

Q(15) Does the scheme have any special security considerations beyond those intrinsic to congestion control?

### 13. Acknowledgments

The development of this document was spurred by the questions asked at the IETF90 CoRE WG session on congestion control, in particular those by Lars Eggert and Richard Scheffenegger (who also supplied most of the text for section <u>Section 4</u>). It is also based on the

Expires September 10, 2015 [Page 5]

experience in CoCoA evaluation by August Betzler, Carles Gomez, Ilker Demirkol, Josep Paradells, Matthias Kovatsch.

## **<u>14</u>**. References

## <u>**14.1</u>**. Normative References</u>

[RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", <u>RFC 7252</u>, June 2014.

## 14.2. Informative References

- [RFC2914] Floyd, S., "Congestion Control Principles", <u>BCP 41</u>, <u>RFC</u> 2914, September 2000.
- [RFC5405] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", <u>BCP 145</u>, <u>RFC 5405</u>, November 2008.
- [RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", <u>RFC 6298</u>, June 2011.
- [incredibles]

Kurkowski, S., Camp, T., and M. Colagrosso, "MANET simulation studies: the incredibles", SIGMOBILE Mob. Comput. Commun. Rev. 9(4) p. 50-61, DOI: 10.1145/1096166.1096174, 2005.

Author's Address

Carsten Bormann Universitaet Bremen TZI Postfach 330440 Bremen D-28359 Germany

Phone: +49-421-218-63921 Email: cabo@tzi.org

Expires September 10, 2015 [Page 6]