

NWCRG
Internet-Draft
Intended status: Informational
Expires: August 9, 2020

N. Kuhn
CNES
E. Lochin
ISAE-SUPAERO
F. Michel
UCLouvain
M. Welzl
University of Oslo
February 6, 2020

**Coding and congestion control in transport
draft-irtf-nwcrg-coding-and-congestion-01**

Abstract

Coding is a reliability mechanism that is distinct and separated from the loss detection of congestion controls. Using coding can be a useful way to better deal with tail losses or with networks with non-congestion losses. However, coding mechanisms should not hide congestion signals. This memo offers a discussion of how coding and congestion control can coexist. This document can help the comparison of FEC schemes by identifying at which level they are operating with respect to the transport congestion control.

This document is the product of the Coding for Efficient Network Communications Research Group (NWCRG). The scope of the document is end-to-end communications: coding for tunnels is out-of-the scope of the document.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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 <https://datatracker.ietf.org/drafts/current/>.

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 August 9, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) 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 Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1.](#) Introduction [2](#)
- [2.](#) Separate channels, separate entities [3](#)
- [3.](#) Scope [4](#)
 - [3.1.](#) Type of application [4](#)
 - [3.2.](#) End-to-end [4](#)
 - [3.3.](#) Objective of the document [5](#)
- [4.](#) FEC above the transport [5](#)
- [5.](#) FEC within the transport [6](#)
- [6.](#) FEC below the transport [6](#)
- [7.](#) Acknowledgements [7](#)
- [8.](#) IANA Considerations [8](#)
- [9.](#) Security Considerations [8](#)
- [10.](#) Informative References [8](#)
- Authors' Addresses [8](#)

[1.](#) Introduction

There are cases where deploying coding improves the quality of the transmission. As an example, it may take time for the server to detect tail losses and this would impact the experience of applications using short flows. Another example are networks where non-congestion losses are persistent and prevent a sender from exploiting the link capacity.

Coding is a reliability mechanism that is distinct and separated from the loss detection of congestion controls. [[RFC5681](#)] defines TCP as a loss-based congestion control; because coding repairs such losses, blindly applying it may easily lead to an implementation that also hides a congestion signal to the sender. It is important to ensure that such information hiding does not occur.

This memo offers a discussion of how coding and congestion control can coexist. One objective is to help the research community consider congestion control aspects when proposing and comparing coding solutions.

This document represents the collaborative work and consensus of the Coding for Efficient Network Communications Research Group (NWCRCG); it is not an IETF product and is not a standard. The proposed recommendations apply for coding at the transport or application layer. Coding for tunnels is out of scope for the document.

2. Separate channels, separate entities

Figure 1 presents the notations that will be used in this document and introduce the Congestion Control (CC) and Forward Erasure Correction (FEC) channels. The Congestion Control channel carries data packets (from a server to a client) and potential information signaling the packets that have been received (from the client to the server). The Forward Erasure Correction channel carries coded packets (from the server to the client) and a potential information signaling the packets that have been repaired (from the client to the server). It is worth pointing out that there are cases where these channels are not separated.




Figure 1: Notations and separate channels

Inside a host, the CC and FEC entities can be regarded as conceptually separate:

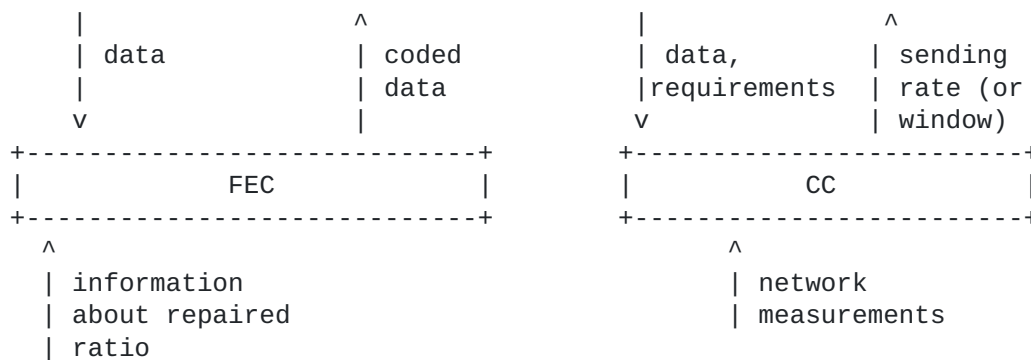


Figure 2: Separate entities (server side)

Figure 2 provides more details than Figure 1 by focusing on the server side. The inputs to FEC (data to work upon, and signaling from the receiver about losses and/or repaired blocks) are distinct from the inputs to CC. The latter calculates a sending rate or window from network measurements, and it takes the data to send as input, sometimes along with application requirements such as upper/lower rate bounds, periods of quiescence, or a priority.

It is not clear that the ACK signals feeding into a congestion control algorithm are useful to FEC in their raw form, and vice versa - information about repaired blocks may be quite irrelevant to a CC algorithm. However, there can be meaningful other interactions (indicated by the horizontal double arrow) between the two entities, usually as a result of their operation rather than by relaying their own raw inputs.

3. Scope

This section describes the scope of the document.

3.1. Type of application

The document focuses on reliable transmissions.

3.2. End-to-end

The document focuses on end-to-end coding, i.e. cases where coding is added at the server and client end points. The discussions should then consider fairness with non-coding solutions.

3.3. Objective of the document

[Section 2](#) presents that FEC and CC can be seen as two separate channels. In practice, implementations may mix the signals that are exchanged on these channels. This document discusses how FEC and CC relate. The discussion can help in avoiding comparing FEC schemes that do not operate at the same level (with respect to the CC). [NK: We may want to later add that we provide recommendations - but for the moment, we may just list the server side coding solutions in the layering approach without providing recommendations for each case]

4. FEC above the transport

Figure 3 illustrates the FEC above the transport case.

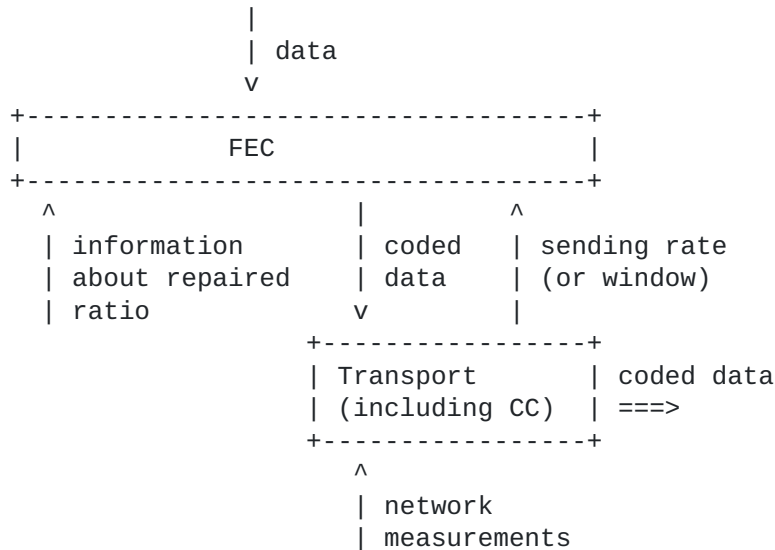


Figure 3: FEC above the transport (server side)

Figure 3 presents an architecture where FEC is on top of the transport. The coded packets are sent within what the congestion window allows.

The advantage of this approach is that the FEC does not contribute to adding congestion in the network.

The drawback of this approach is that it is useless if the transport guarantees data delivery by retransmitting packets.

5. FEC within the transport

Figure 4 illustrates the FEC within the transport case.

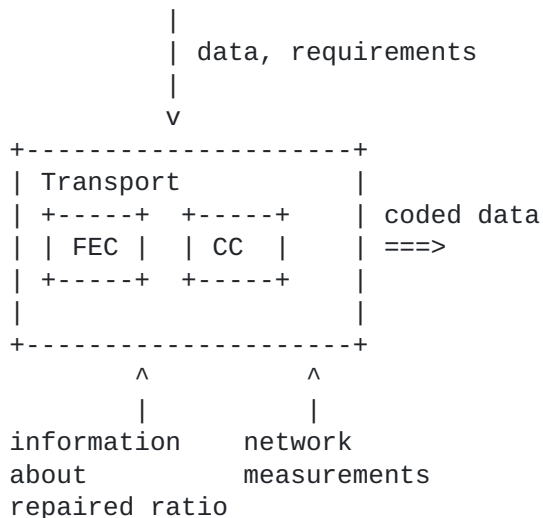


Figure 4: FEC in the transport (server side)

Figure 4 presents an architecture where FEC is within the transport. The coded packets are sent within what the congestion window allows, such as in [CTCP]. Examples of the solution would be sending coded packets when there is no more data to transmit or preferably send coded packets instead of the following packets in the send buffer.

The advantage of this approach is that it can enable conjoint optimization between the CC and the FEC. Moreover, the transmission of coded packets does not add congestion in potentially congested networks but help repair lost packets (such as tail losses).

The drawback of this approach is that it may not result in much gains as opposed to classical CC retransmissions mechanisms and it costs bandwidth that could have been exploited to transmit non coded data packets. The coding ratio needs to be carefully designed.

6. FEC below the transport

Figure 5 illustrates the FEC below the transport case.

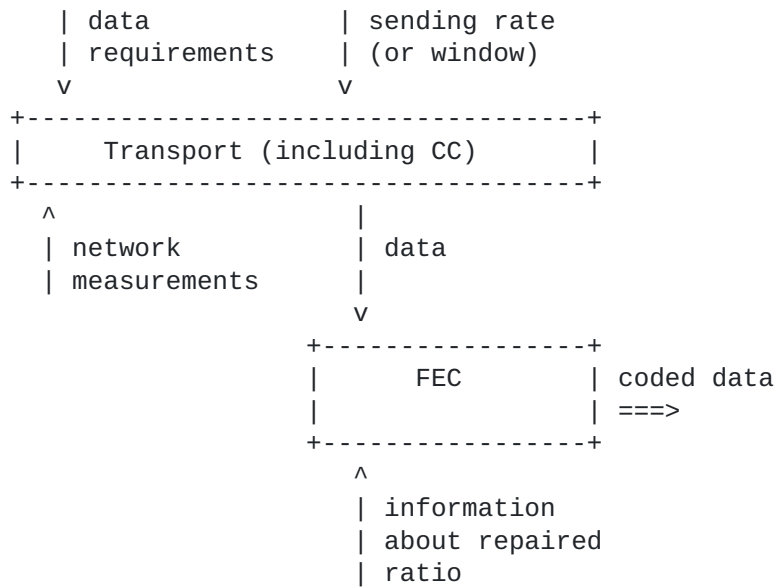


Figure 5: FEC below the transport (server side)

Figure 5 presents an architecture where FEC is below the transport. The coded packets are sent on top of what is allowed by the congestion control. Examples of the solution could be adding a given percentage of the congestion window as supplementary packets or sending a given amount of coded packets at a given rate. The redundancy flow can be decorrelated from the congestion control that manages source packets: a secondary congestion control could be introduced underneath the FEC layer, such as in coupled congestion control for RTP media [[I-D.ietf-rmcat-coupled-cc](#)]. An example would be to exploit a lower than best-effort congestion control [[RFC6297](#)].

The advantage of this approach is that it can result in performance gains when there are persistent transmission losses along the path.

The drawback of this approach is that it can induce congestion in already congested networks. The coding ratio needs to be carefully designed.

7. Acknowledgements

Many thanks to Spencer Dawkins, Dave Oran, Carsten Bormann, Vincent Roca and Marie-Jose Montpetit for their useful comments that helped improve the document.

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

There is no security considerations required for this document.

10. Informative References

- [CTCP] Kim (et al.), M., "Network Coded TCP (CTCP)", arXiv 1212.2291v3, 2013.
- [I-D.ietf-rmcat-coupled-cc] Islam, S., Welzl, M., and S. Gjessing, "Coupled congestion control for RTP media", [draft-ietf-rmcat-coupled-cc-09](#) (work in progress), August 2019.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", [RFC 5681](#), DOI 10.17487/RFC5681, September 2009, <<https://www.rfc-editor.org/info/rfc5681>>.
- [RFC6297] Welzl, M. and D. Ros, "A Survey of Lower-than-Best-Effort Transport Protocols", [RFC 6297](#), DOI 10.17487/RFC6297, June 2011, <<https://www.rfc-editor.org/info/rfc6297>>.

Authors' Addresses

Nicolas Kuhn
CNES

Email: nicolas.kuhn@cnes.fr

Emmanuel Lochin
ISAE-SUPAERO

Email: emmanuel.lochin@isae-supaero.fr

Francois Michel
UCLouvain

Email: francois.michel@uclouvain.be

Michael Welzl
University of Oslo

Email: michawe@ifi.uio.no