

NWCRG N.
Kuhn
Internet-Draft
CNES
Intended status: Informational E.
Lochin
Expires: September 6, 2020 ISAE-
SUPAERO F.
Michel
UCLouvain
Welzl M.
Oslo University of
2020 March 5,

**Coding and congestion control in transport
draft-irtf-nwcrg-coding-and-congestion-02**

Abstract

FEC coding is a reliability mechanism that is distinct and separate from the loss detection of congestion controls. Using FEC coding can be a useful way to deal with transfer tail losses or with networks having non-congestion losses. However, FEC coding mechanisms should not hide congestion signals. This memo offers a discussion of how FEC coding and congestion control can coexist. Another objective is to encourage the research community to also consider congestion control aspects when proposing and comparing FEC coding solutions in communication systems.

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: FEC 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 September 6, 2020.

Kuhn, et al.
1]

Expires September 6, 2020

[Page

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](#). FEC and CC layering
 - [3.1](#). FEC above the transport
 - [5](#)
 - [3.2](#). FEC within the transport
 - [6](#)
 - [3.3](#). FEC below the transport
 - [7](#)
- [4](#). Acknowledgements
- [8](#)
- [5](#). IANA Considerations
- [9](#)
- [6](#). Security Considerations
- [9](#)
- [7](#). Informative References
- [9](#)
- Authors' Addresses
- [9](#)

[1](#). Introduction

There are cases where deploying FEC coding improves the performance of a transmission. As an example, it may take time for the sender to

detect transfer tail losses (losses that occur at the end of a transfer, where e.g. TCP obtains no more ACKs to quickly repair the loss via retransmission). This would improve 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 separate from the loss detection of congestion controls. [[RFC5681](#)] defines TCP as a loss-based congestion control; because FEC 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.

FEC coding and congestion control can be seen as two separate channels. In practice, implementations may mix the signals that are exchanged on these channels. This memo offers a discussion of how FEC coding and congestion control can coexist. Another objective is to encourage the research community to also consider congestion control aspects when proposing and comparing FEC coding solutions in communication systems. This being said, this document does not aim at proposing guidelines for characterizing FEC coding solutions: comparing FEC Schemes without considering congestion control can be relevant if the goal is to compare those schemes only.

The proposed document considers FEC coding at the transport or application layer for an end-to-end unicast data transfer. The typical application scenario that is considered in the current version of the document is a client browsing the web. This memo may be extended to cases with multiple paths.

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 document follows the terminology proposed in the taxonomy document [[RFC8406](#)].

2. Separate channels, separate entities

Figure 1 presents the notations that will be used in this document and introduces the Congestion Control (CC) and Forward Erasure Correction (FEC) channels. The Congestion Control channel carries source packets from a sender to a receiver, and packets signaling information about the network (number of packets received vs. lost, ECN marks, etc.) from the receiver to the sender. The Forward Erasure Correction channel carries repair packets (from the sender to the receiver) and potential information signaling which packets have been repaired (from the receiver to the sender). It is worth pointing out that there are cases where these channels are not separated.

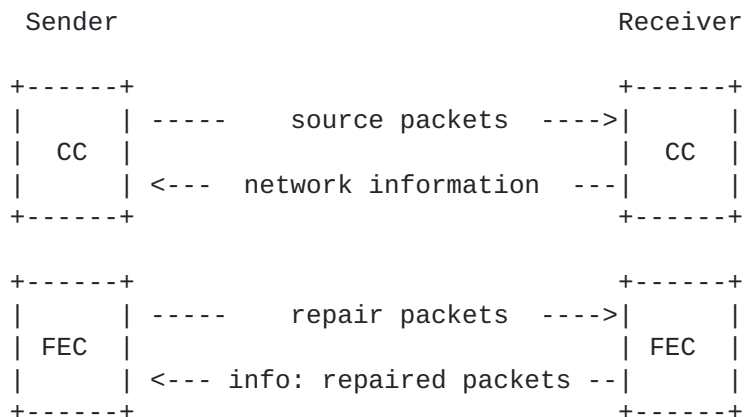


Figure 1: Notations and separate channels

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

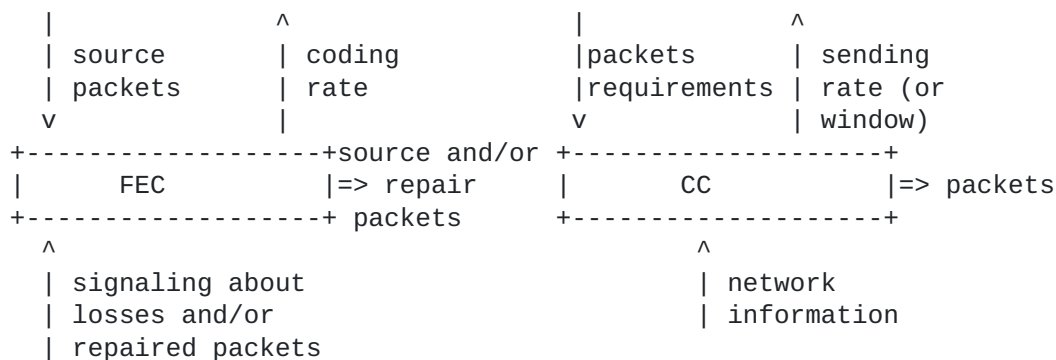


Figure 2: Separate entities (sender-side)

Figure 2 provides more details than Figure 1 by focusing on the sender-side. Some elements are introduced:

- o 'network information' (input control plane for the transport including CC): refers not only to the network information that is explicitly signaled from the receiver, but all the information a congestion control obtains from a network (e.g. TCP can estimate the latency and the available capacity at the bottleneck).
- o 'requirements' (input control plane for the transport including CC): refers to application requirements such as upper/lower rate bounds, periods of quiescence, or a priority.

- o 'sending rate (or window)' (output control plane for the transport including CC): refers to the rate at which a congestion control decides to transmit packets, based on 'network information'.
- o 'signaling about losses and/or repaired packets' (input control plane for the FEC): refers to the information a FEC sender can obtain from a FEC receiver about the performance of the FEC solution as seen from the receiver.
- o 'coding rate' (output control plane for the FEC): refers to the coding rate that is used by the FEC solution.
- o 'source and/or repair packets' (data plane for both the FEC and the CC): refers to the packets that are transmitted. There can only be source packets (if the coding rate is 0), repair packets (if the solution decides not to send the original source packets) or a mix of both.

The inputs to FEC (packets to work upon, and signaling from the receiver about losses and/or repaired packets) are distinct from the inputs to CC. The latter calculates a sending rate or window from network information, and it takes the packet 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.

3. FEC and CC layering

This section discusses how FEC and CC can relate in different cases (FEC above the transport, FEC within the transport, FEC below the transport).

3.1. FEC above the transport

Figure 3 illustrates the FEC above the transport case.

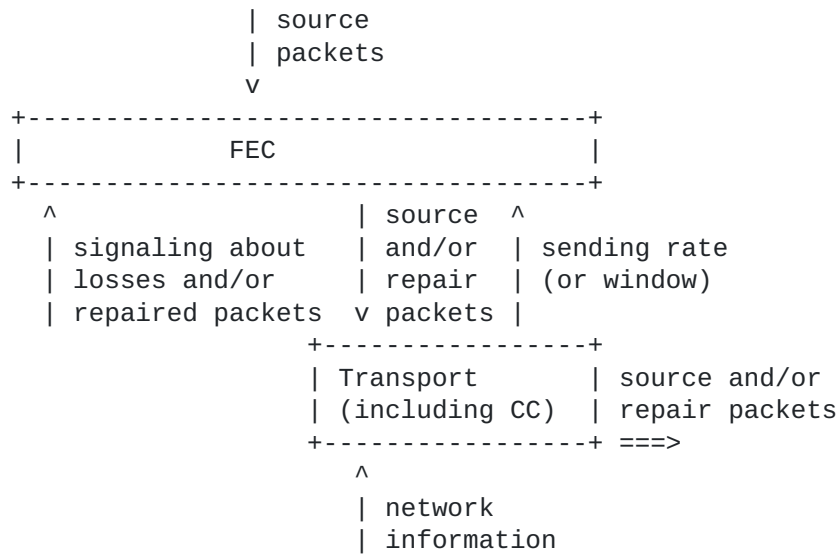


Figure 3: FEC above the transport (sender-side)

Figure 3 presents an architecture where FEC is on top of the transport. The repair 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.

While CC is in principle independent of other transport functions such as reliability, we note that CC is often embedded in reliable transfer protocols (e.g. TCP). This approach requires that the transport protocol does not implement a fully reliable data transfer service (e.g., based on lost packet retransmission). UDP is an example of a protocol for which this approach is relevant.

3.2. FEC within the transport

Figure 4 illustrates the FEC within the transport case.

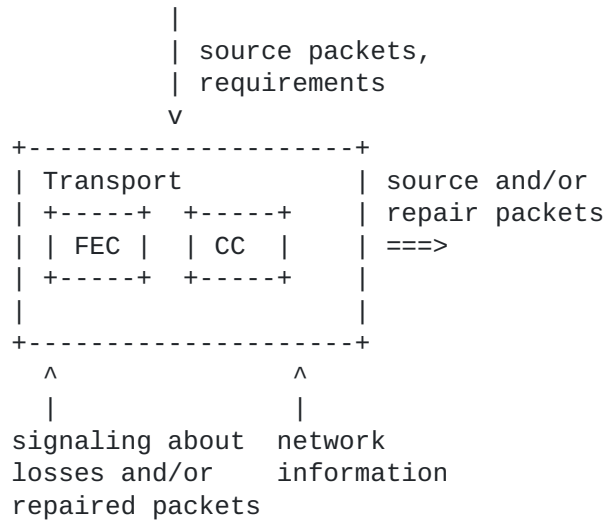


Figure 4: FEC in the transport (sender-side)

Figure 4 presents an architecture where FEC is within the transport. The repair packets are sent within what the congestion window allows, such as in [CTCP]. Examples of the solution would be sending repair packets when there is no more data to transmit or preferably send repair 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 repair packets does not add congestion in potentially congested networks but helps 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 retransmission mechanisms and it costs bandwidth that could have been exploited to transmit source packets. The coding ratio needs to be carefully designed.

3.3. FEC below the transport

Figure 5 illustrates the FEC below the transport case.

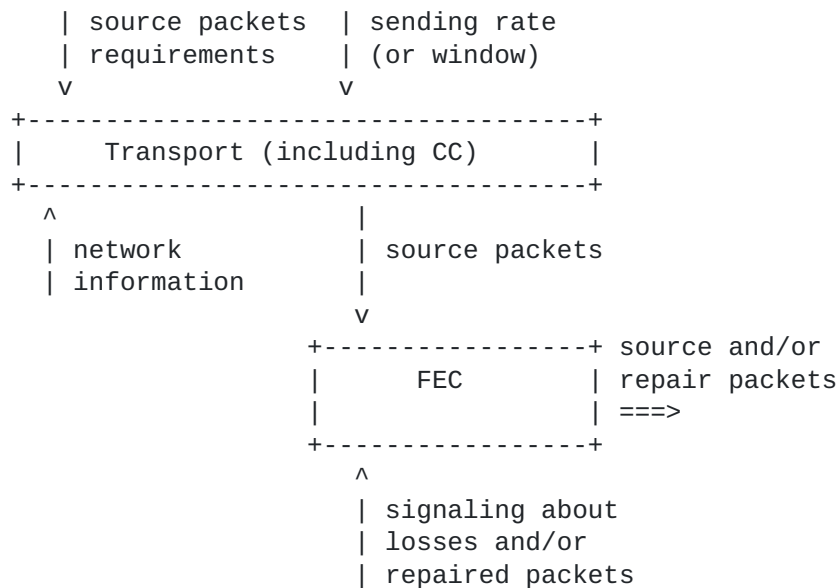


Figure 5: FEC below the transport (sender-side)

Figure 5 presents an architecture where FEC is below the transport. The repair 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 repair 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. The separate congestion control mechanisms could be made to work together while adhering to priorities, as in coupled congestion control for RTP media [[I-D.ietf-rmcat-coupled-cc](#)]. Another possibility would be to exploit a lower than best-effort congestion control [[RFC6297](#)] for repair packets.

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.

4. 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.

5. IANA Considerations

This memo includes no request to IANA.

6. Security Considerations

FEC and CC schemes can contribute to DoS attacks. This is not specific to this document.

In case of FEC below the transport, the aggregate rate of source and repair packets may exceed the rate at which a congestion control mechanism allows an application to send. This could result in an application obtaining more than its fair share of the network capacity.

7. Informative References

[CTCP] Kim (et al.), M., "Network Coded TCP (CTCP)", arXiv 1212.2291v3, 2013.

[I-D.ietf-rmcat-coupled-cc]
congestion Islam, S., Welzl, M., and S. Gjessing, "Coupled 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>>.

[RFC8406] Adamson, B., Adjih, C., Bilbao, J., Firoiu, V., Fitzek, F., Ghanem, S., Lochin, E., Masucci, A., Montpetit, M-J., Pedersen, M., Peralta, G., Roca, V., Ed., Saxena, P., and S. Sivakumar, "Taxonomy of Coding Techniques for Efficient Network Communications", [RFC 8406](#), DOI 10.17487/RFC8406, June 2018, <<https://www.rfc-editor.org/info/rfc8406>>.

Authors' Addresses

Nicolas Kuhn
CNES

Email: nicolas.kuhn@cnes.fr

Internet-Draft
2020

Coding and congestion

March

Emmanuel Lochin
ISAE-SUPAERO

Email: emmanuel.lochin@isae-supero.fr

Francois Michel
UCLouvain

Email: francois.michel@uclouvain.be

Michael Welzl
University of Oslo

Email: michawe@ifi.uio.no

