

Framework for Loss Detection, Retransmission, and Congestion Control in ICN

Anders Eriksson, Börje Ohlman

2017-03-26

Packet Loss Rate in Real Networks

Measurement and Analysis of Internet Interconnection
and Congestion

David Clark
Steven Bauer
William Lehr
CSAIL/MIT
ddc,bauer,wlehr@mit.edu

kc claffy
Amogh Dhamdhare
Bradley Huffaker
Matthew Luckie
CAIDA
UC, San Diego
kc,amogh,brad,mjl@caida.org

September 9, 2014


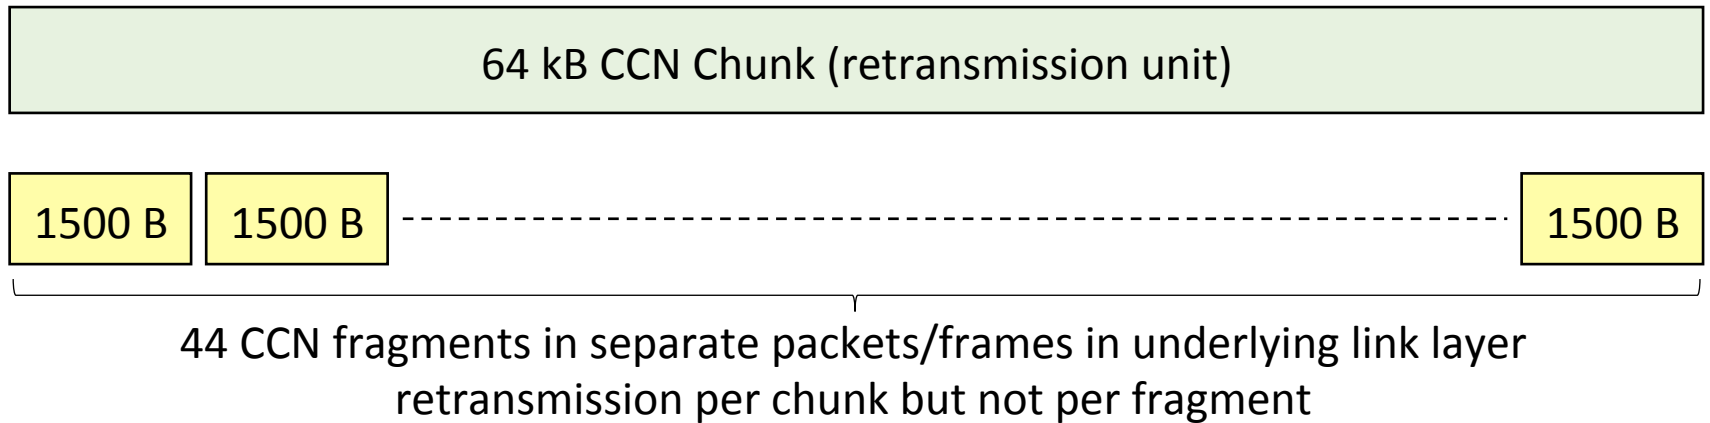



Figure 2: Congestion on an interdomain link between Comcast and Cogent, measured from a VP within Comcast. The round trip time to the Cogent (far) router increases from 20ms to 70ms while the round trip time to the Comcast (near) router is stable at 20ms. The square wave indicates the queue is always close to full when the round trip time increases to 70ms. The loss rate from the Cogent router increases after this level shift occurs.

Example 1: Fragmentation of a 64 kB CCN Chunk



Example 2: 1500 B CCN Chunks

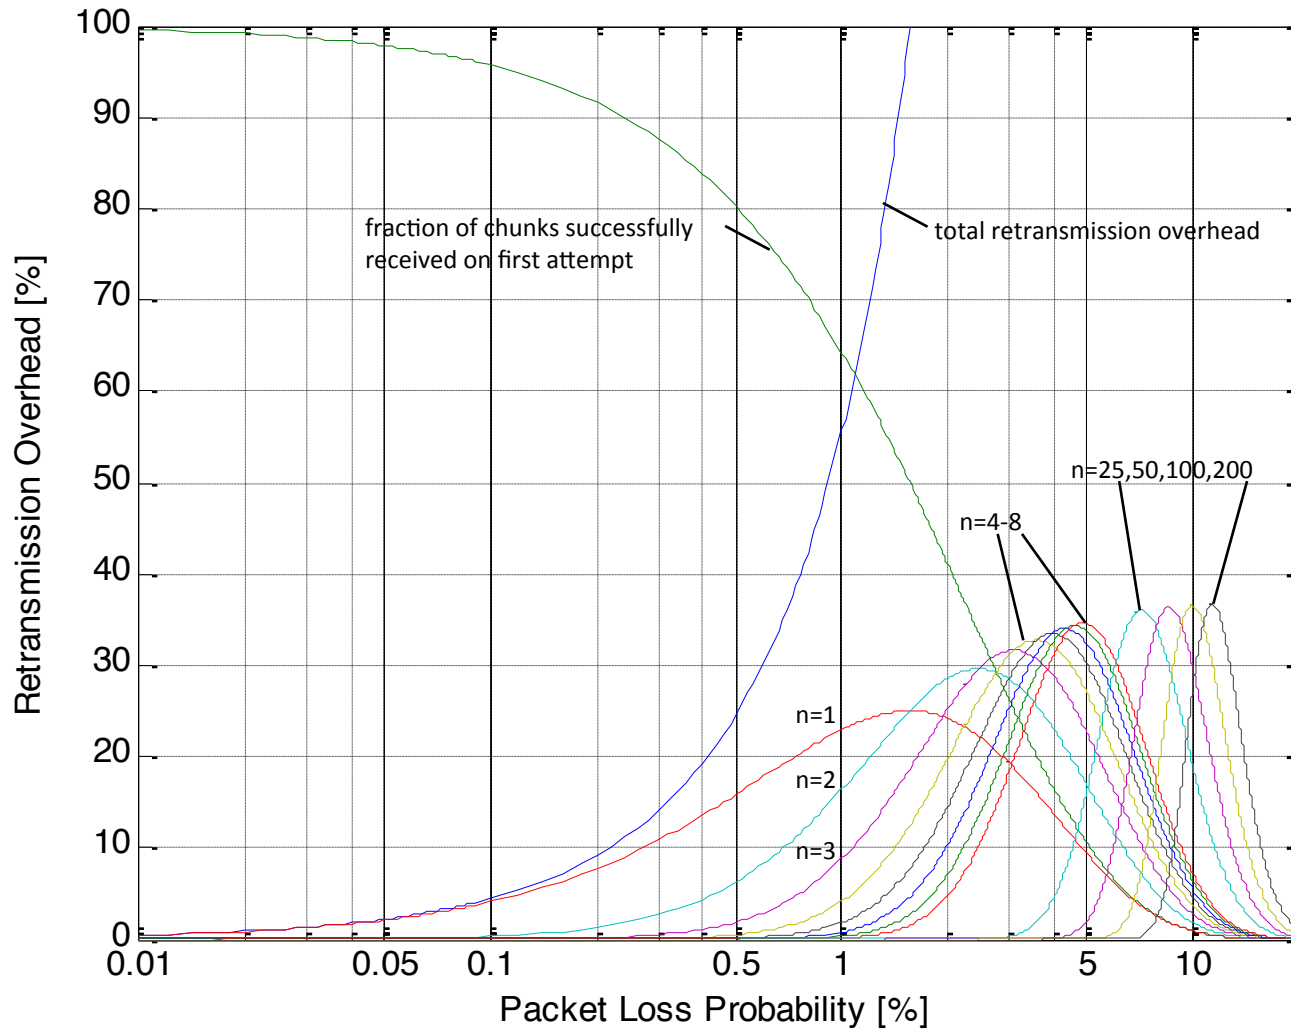


Overhead issue: One Interest message per 1500 B Content Object

retransmission unit

CCN Fragmentation with 44 Packets per Chunk

Overhead from CHUNKS Retransmitted n Times



The sum of all traces $n=1 \rightarrow \infty$ = total retransmission overhead (blue line)

TCP/UDP/IP vs. CCN/NDN

- **The IP network layer** is best-effort and IP packets may be lost, reordered, or corrupted. Also, there may be congestion in IP routers, or in e.g. Ethernet switches.
- **For IP**, there is a layering approach with **TCP** (loss detection, retransmission, and congestion control) or **UDP** (none of this, lightweight).
- **TCP and UDP** offer socket **APIs** to the application.
- **CCN/NDN packets** may be lost or corrupted. If they are carried over IP and/or Ethernet, congestion in IP routers and Ethernet switches is an issue.
- Some **link layer mechanisms** detect packet loss and perform retransmission.
- **Question:** Is an API and layering framework needed for CCN/NDN which takes packet loss, retransmission, and congestion control into account?

Approaches

- TCP-like approach: Hide CCN/NDN packet loss and congestion from the application
- UDP-like approach: Let the application deal with CCN/NDN packet loss and congestion
- Two possible API approaches:
 - Basic API, only Interest and Content Object messages
 - Feature-rich API which exposes information relevant for packet losses, retransmissions, and congestion control
- Underlying layer mechanisms that could support some of this:
 - CCN/NDN over TCP/IP per hop
 - WLDR (Wireless Loss Detection and Recovery)*
 - Loss detection and retransmissions in cellular networks

* G. Carofiglio, L. Muscariello, M. Papalini, N. Rozhnova, X. Zeng: Leveraging ICN In-network Control for Loss Detection and Recovery in Wireless Mobile Networks, ACM ICN 2016

CCN over Lossy Link Layer

Network API 3

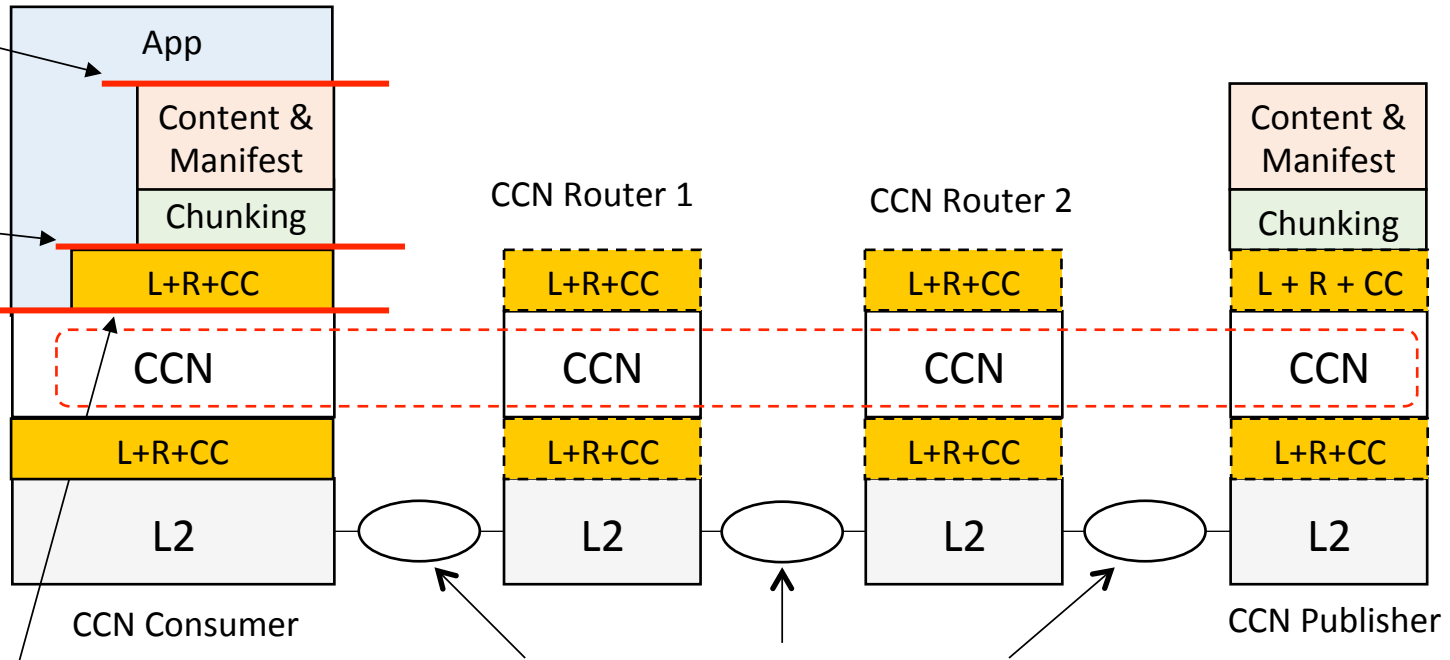
Reliable GET
content-name
(e.g. FLIC)

Network API 2

Reliable GET
chunk-name

Network API 1

Unreliable GET
chunk-name



Chunk-centric
Network Layer
(chunk transfer only)

Layer 2 carrying both CCN traffic and legacy TCP IP traffic.

CCN (L+R+CC) must then:

- be TCP-friendly
- deal with packet loss
- packets may be lost both in L2 and in the CCN layer

L Packet Loss Detection
R Retransmission
CC Congestion Control

Loss Detection, Retransmission, and Congestion Control at Different Layers

Above the CCN layer

- API with Interest and Content Object messages, but no awareness of topology
- Timeout-based Retransmission of Interests
- With a richer API, more can be done at this layer

In the CCN layer

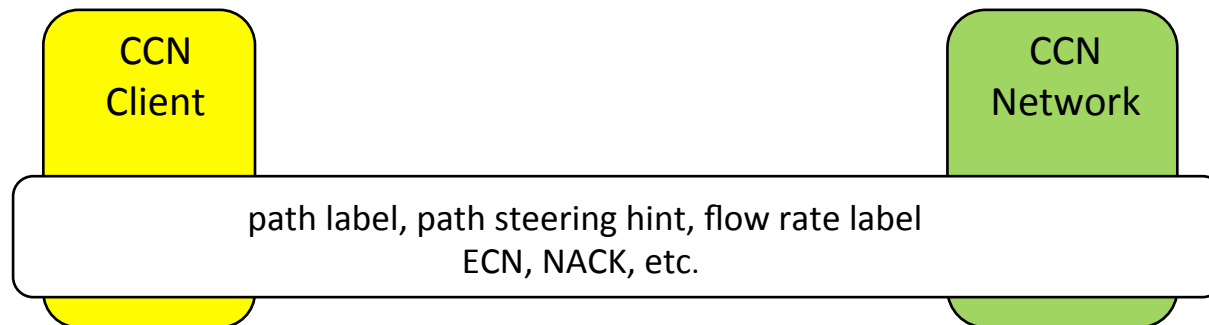
- FIB & PIT awareness
- Multi-path congestion control using for example:
 - path labels (data), path steering hints (Interest), number of pending Interests per interface or prefix, ECN, flow rate marking
- Retransmission of expired Interests, NACK

Below the CCN layer

- A CCN hop may include a wireless link, an Ethernet network, an IP network, etc.
- TCP-like loss detection, retransmission, and congestion control per CCN hop
- TCP-friendliness is an issue

Research & Standardization Challenge

- TCP and QUIC run in clients, but not in IP routers.
- TCP and QUIC can therefore be developed independently of IP router technology.
- By contrast, CCN/NDN congestion control, loss detection, and retransmission run **both in clients and in routers**.
- Research & standardization challenge:
Design **forward and backward compatible** client and router mechanisms for reliable packet delivery and congestion control.



Summary

- TCP/UDP/IP has a layered framework with APIs and transport protocols handling different needs with regard to packet loss, retransmission and congestion control.
- Is something similar needed for ICN?

Backup Slides

Why is Retransmission Part of Congestion Control?

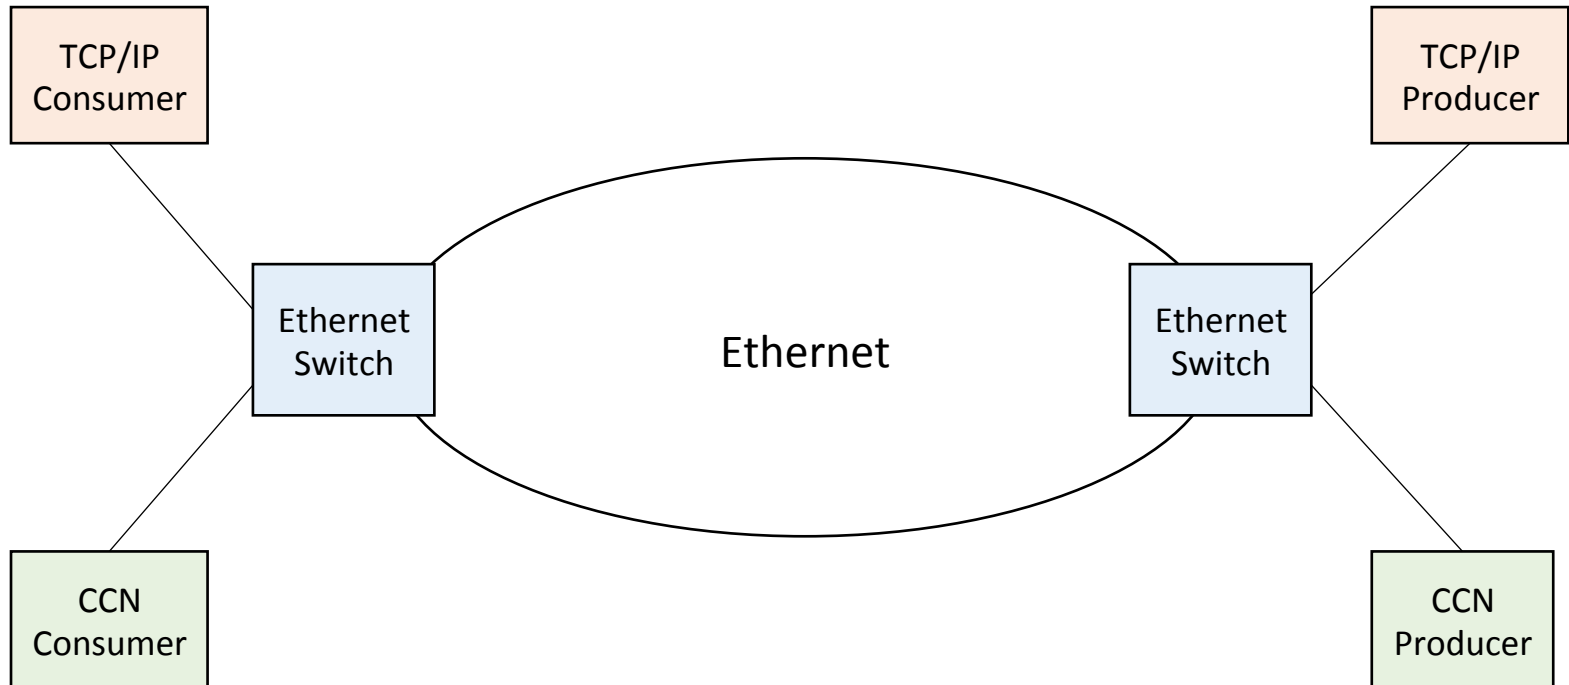
“In October of '86, the Internet had the first of what became a series of ‘congestion collapses’. We were fascinated by this sudden factor-of-thousand drop in bandwidth and embarked on an investigation of why things had gotten so bad.

Since that time, we have put seven new algorithms into the 4BSD TCP:

- (i) round-trip-time variance estimation
- (ii) exponential retransmit timer backoff
- (iii) slow-start
- (iv) more aggressive receiver ack policy
- (v) dynamic window sizing on congestion
- (vi) Karn’s clamped retransmit backoff
- (vii) fast retransmit”

Van Jacobson and J. Karels, November 1988

ICN Packet Loss Scenario

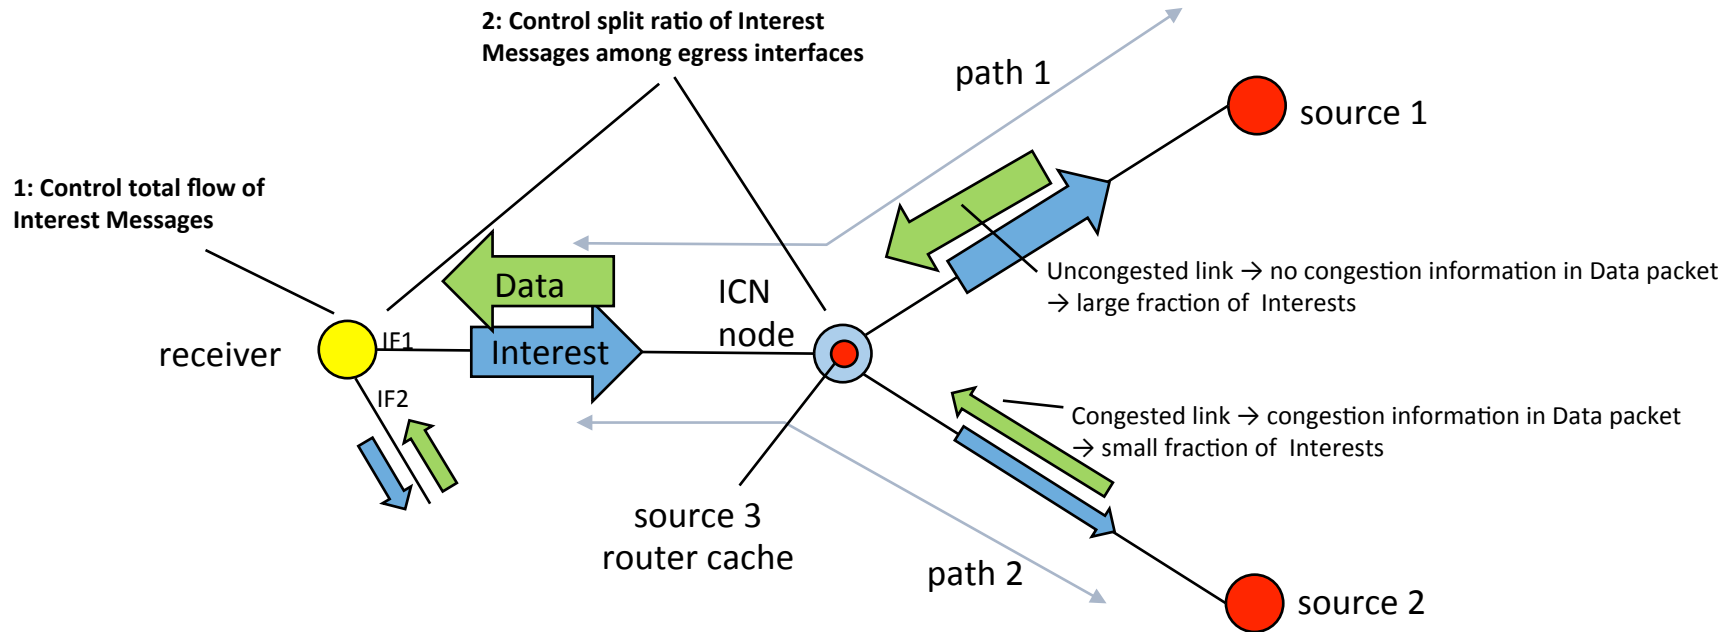


- TCP will ramp up traffic until packet loss occurs, and this will hit both TCP/IP and ICN packets.
- Separation of TCP/IP and ICN traffic using VLANs is an alternative, but this adds management complexity.

CCN Terminology

- One **Interest Message** is used to request one **Content Object**.
- A **Content Object** is the data message sent in response to an Interest Message. A Content Object has a maximum size of 64 kB. Larger amounts of application data than 64 kB can be **chunked** into Content Objects. A Content Object may have a metadata object that describes the original pre-chunked object.
- In this presentation we use the term **chunk** for a Content Object, since it is typically a chunk of a larger application data object.
- **End-to-end fragmentation:** CCN protocol capability to fragment and reassemble Interest Messages and Content Objects in end systems. Fragmentation headers in the CCN protocol are used for hop-by-hop forwarding and reassembly of fragments. The Interest Message is used for path MTU discovery. The minimal fragment size is 1280 bytes. No retransmission of lost fragments.
- **Hop-by-hop fragmentation:** CCN protocol capability to fragment and reassemble Interest and Content Messages on a hop-by-hop basis. Retransmission of lost fragments is not supported.
- A **Manifest** is a Content Object with a well-known payload format. It contains hash-based names of a collection of Content Objects along with metadata for the collection. A hierarchy of manifests can be used to represent a very large application data object.
- **Improve on this:** A **Flow** is the collection of all Interest Messages and Content Objects needed to retrieve an application data object, possibly from multiple sources. If Content Objects are retrieved along multiple paths, the flow consists of one sub-flow per path.
- Streams???

CCN Congestion Control Problem Statement



- congestion control per flow (data object) with chunks from different sources
- some chunks of a data object may be retrieved from a local cache, and others from a source on another continent -> retransmission timeout issue
- lost packets should be retransmitted quickly
- fairness among flows
- Interest messages are so big that they can contribute to congestion

Packet Loss Recovery in icnrg RFCs

- "Information-Centric Networking (ICN) Research Challenges" RFC 7927:

" In case of disruption (message not delivered), a node can resend the request, and it could be answered by an on-path cache, i.e., on the other side of the disrupted link. The network itself would be able to send local retransmissions, which enables shorter round-trip times and the offloading of origin servers and other parts of the network."

- RFC 7933 "Adaptive Video Streaming over Information-Centric Networking (ICN)"

Packet loss recovery and forward error correction mentioned in the context of IPTV (§7.1)

ICN Conference Papers on this Topic

- Multitude of papers on ICN congestion control
- Commonality: CCN/NDN congestion control mechanisms are needed both in clients **and** in CCN/NDN routers.
- Co-existence between TCP/IP and ICN traffic → ICN packets will be lost (unless resource separated).
- “CCN & TCP co-existence in the future Internet: Should CCN be compatible to TCP?”
(S. Braun, M. Monti, M. Sifalakis, C. Tschudin; IFIP 2013)
Highlights the TCP/IP co-existence issue and the RTT estimation issue with multiple content sources.
- “Leveraging ICN In-network Control for Loss Detection and Recovery in Wireless Mobile networks” (ACM ICN 2016, G. Carofiglio, L. Muscariello, M. Papalini, N. Rozhnova, X. Zeng)
- “Transport Layer Design for Named Data Wireless Networking” (IEEE Infocom 2014 Workshops, M. Amadeo, A. Molinaro, C. Campolo, M. Sifalakis, C. Tschudin)

ICNRG drafts on this topic

- CCNx Semantics draft-irtf-icnrg-ccnxsemantics-00
If a node sends an Interest and receives a
Congestion InterestReturn,
then try a different forwarding path, if one exists.