

# Practical Congestion Control for NDN: Short-term Design & Implementation

7th NDN Retreat


---

Klaus Schneider

November 4, 2016

The University of Arizona


## Multiple Paths and Endpoints



Mixing RTT measurements from different sources

⇒ Problem: Traditional RTO settings often **too short**

# Multiple Paths and Endpoints




Mixing RTT measurements from different sources

⇒ Problem: Traditional RTO settings often **too short**

1. Use **Route-labels** to know content origin and path [3]
  - Still don't know where next Interest will go! [6]

# Multiple Paths and Endpoints




Mixing RTT measurements from different sources

⇒ Problem: Traditional RTO settings often **too short**

1. Use **Route-labels** to know content origin and path [3]
  - Still don't know where next Interest will go! [6]
2. Predicting location of **future data** [9, 1]
  - Routers mark Data to indicate their content
  - Overhead? Reliable?

# Hop-By-Hop Interest Shaping


**At each hop:** Shape Interests to control returning Data.



Source: Wang et al. - An Improved HBH Interest Shaper for NDN [10]

Much work [2, 10, 8, 4, 12, 11, 7, 5] *based on that principle!*

# Hop-By-Hop Interest Shaping



HBH Interest Shaping assumes that you

- know the **link capacity**
- know the **Data chunk size**

**Estimation errors cost performance!**

# PCON: Design Principles

Schneider et al. – **A Practical Congestion Control Scheme for Named Data Networking** (ACM ICN 2016)


**Remove strong assumptions** about the network:

- Unknown link capacity & Data chunk size
- No route-labels or prediction of data location

**Design Principles:**

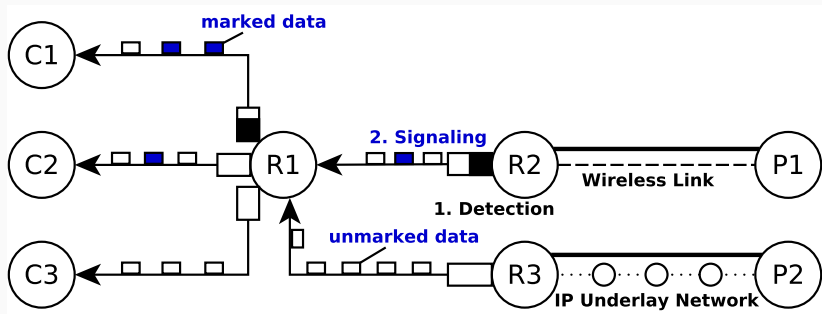
- Detect congestion at the bottleneck!
- Signal it towards consumer

# System Design: Overview

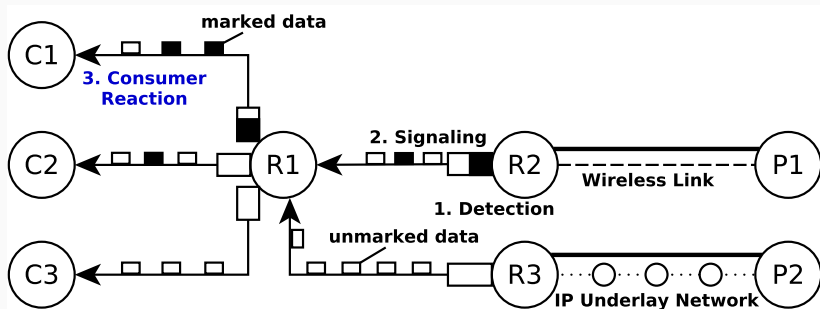




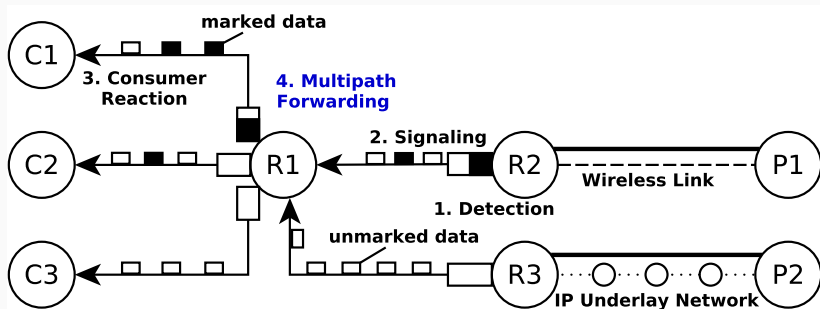
# System Design: Overview



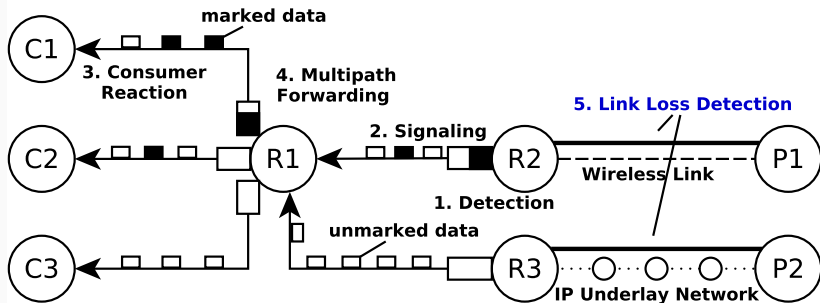
# System Design: Overview



# System Design: Overview



# System Design: Overview



# Implementation Parts:

- ndncatchunks: simple consumer reaction.
    - Conservative window adaptation (SACK)
    - Traditional AIMD with slow start.
    - Other options?
    - Timeout on version discovery?
- ⇒ Consumer/Producer API (Ilya Moiseenko et al.)

# Implementation Parts:

- ndncatchunks: simple consumer reaction.
  - Conservative window adaptation (SACK)
  - Traditional AIMD with slow start.
  - Other options?
  - Timeout on version discovery?

⇒ Consumer/Producer API (Ilya Moiseenko et al.)
- Generic congestion marks
  - For Data and NACKs (and Interests?)
  - Single-bit vs. Multi-bit
  - Read and set at routers
  - Implemented in NDNLP

# Congestion detection on UDP tunnels (Testbed)

Vusirikala et al. – **Hop-By-Hop Best Effort Link Layer Reliability in Named Data Networking** (Tech Report)

Detect local “link” losses and signal to forwarding strategy

# Congestion detection on UDP tunnels (Testbed)

Vusirikala et al. – **Hop-By-Hop Best Effort Link Layer Reliability in Named Data Networking** (Tech Report)

Detect local “link” losses and signal to forwarding strategy

- How many local retransmissions?



# Congestion detection on UDP tunnels (Testbed)

Vusirikala et al. – **Hop-By-Hop Best Effort Link Layer Reliability in Named Data Networking** (Tech Report)

Detect local “link” losses and signal to forwarding strategy

- How many local retransmissions?
- Decouple loss notification from retx.

# Congestion detection on UDP tunnels (Testbed)

Vusirikala et al. – **Hop-By-Hop Best Effort Link Layer Reliability in Named Data Networking** (Tech Report)

Detect local “link” losses and signal to forwarding strategy

- How many local retransmissions?
- Decouple loss notification from retx.
- Positive vs. Negative ACKs?
  - Definitive knowledge about loss vs. overhead.

# Congestion detection on UDP tunnels (Testbed)

Vusirikala et al. – **Hop-By-Hop Best Effort Link Layer Reliability in Named Data Networking** (Tech Report)

Detect local “link” losses and signal to forwarding strategy

- How many local retransmissions?
- Decouple loss notification from retx.
- Positive vs. Negative ACKs?
  - Definitive knowledge about loss vs. overhead.
- ACK duplication?
  - Duplicating ACKs cheaper than spurious retx

# Congestion detection on UDP tunnels (Testbed)

Vusirikala et al. – **Hop-By-Hop Best Effort Link Layer Reliability in Named Data Networking** (Tech Report)

Detect local “link” losses and signal to forwarding strategy

- How many local retransmissions?
- Decouple loss notification from retx.
- Positive vs. Negative ACKs?
  - Definitive knowledge about loss vs. overhead.
- ACK duplication?
  - Duplicating ACKs cheaper than spurious retx
- Selective vs. Cumulative ACKs?

# Congestion detection on UDP tunnels (Testbed)

Vusirikala et al. – **Hop-By-Hop Best Effort Link Layer Reliability in Named Data Networking** (Tech Report)

Detect local “link” losses and signal to forwarding strategy

- How many local retransmissions?
- Decouple loss notification from retx.
- Positive vs. Negative ACKs?
  - Definitive knowledge about loss vs. overhead.
- ACK duplication?
  - Duplicating ACKs cheaper than spurious retx
- Selective vs. Cumulative ACKs?

# Congestion detection on Ethernet links / TCP tunnels?

1. Ethernet: Read out NIC queue size + apply AQM logic
2. TCP Tunnel: Read send buffer occupancy

More work needed:

- How to get the right queue?
- How to get the queue inside NFD?

# More General Questions

- Congestion due to packet processing and memory overhead

## More General Questions

- Congestion due to packet processing and memory overhead
- DDoS in NDN: A congestion problem?



## More General Questions

- Congestion due to packet processing and memory overhead
- DDoS in NDN: A congestion problem?
- BBR's applicability to NDN congestion control

## Any Questions?

Klaus Schneider

`klaus@cs.arizona.edu`

`https://www.cs.arizona.edu/~klaus/`

# References I

- [1] Sebastian Braun, Massimo Monti, Manolis Sifalakis, and Christian Tschudin.  
**An empirical study of receiver-based aimd flow-control for ccn.**  
In *IEEE ICCCN*, 2013.
- [2] Giovanna Carofiglio, Massimo Gallo, and Luca Muscariello.  
**Joint hop-by-hop and receiver-driven interest control protocol for content-centric networks.**  
In *ACM ICN workshop*, 2012.
- [3] Giovanna Carofiglio, Massimo Gallo, Luca Muscariello, Michele Papalini, and Sen Wang.  
**Optimal multipath congestion control and request forwarding in information-centric networks.**  
In *ICNP*, 2013.
- [4] Kai Lei, Chaojun Hou, Lihua Li, and Kuai Xu.  
**A rcp-based congestion control protocol in named data networking.**  
In *CyberC*, 2015.
- [5] Chengcheng Li, Tao Huang, Renchao Xie, Hengyang Zhang, Jiang Liu, and Yunjie Liu.  
**A novel multi-path traffic control mechanism in ndn.**  
In *IEEE ICT*, 2015.

# References II

- [6] Dinh Nguyen, Masaki Fukushima, Kohei Sugiyama, and Atsushi Tagami.  
**Efficient multipath forwarding and congestion control without route-labeling in ccn.**  
In *IEEE ICCW*, 2015.
- [7] Heungsoon Park, Hoseok Jang, and Taewook Kwon.  
**Popularity-based congestion control in named data networking.**  
In *IEEE ICUFN*, 2014.
- [8] Natalya Rozhnova and Serge Fdida.  
**An extended hop-by-hop interest shaping mechanism for ccn.**  
In *IEEE GLOBECOM*, 2014.
- [9] Lorenzo Saino, Cosmin Cocora, and George Pavlou.  
**Cctcp: A scalable receiver-driven congestion control protocol for ccn.**  
In *IEEE ICC*, 2013.
- [10] Yaogong Wang, Natalya Rozhnova, Ashok Narayanan, David Oran, and Injong Rhee.  
**An improved hop-by-hop interest shaper for congestion control in named data networking.**  
*ACM SIGCOMM CCR*, 2013.

# References III

- [11] Feixiong Zhang, Yanyong Zhang, Alex Reznik, Hang Liu, Chen Qian, and Chenren Xu.  
**A transport protocol for content-centric networking with explicit congestion control.**  
In *IEEE ICCCN*, 2014.
- [12] Jianer Zhou, Qinghua Wu, Zhenyu Li, Mohamed Ali Kaafar, and Gaogang Xie.  
**A proactive transport mechanism with explicit congestion notification for ndn.**  
In *IEEE ICC*, 2015.