A Practical Congestion Control Scheme for Named Data Networking

ACM ICN 2016

Klaus Schneider[1], Cheng Yi[2], Beichuan Zhang[1], Lixia Zhang[3]

September 27, 2016

[1]The University of Arizona, [2]Google, [3]UCLA

NDN and IP are different!

NDN and IP networks are different:

1. Traditional Congestion Control doesn't work for NDN
2. Related work often makes **too strong assumptions**

$\Rightarrow$ Developing a more practical solution (**PCON**)

Multiple Paths and Endpoints

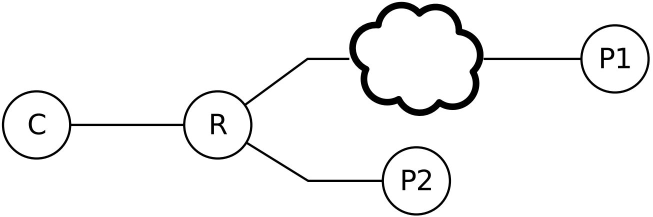

Mixing RTT measurements from different sources

$\Rightarrow$ Problem: Traditional RTO settings often **too short**

Multiple Paths and Endpoints

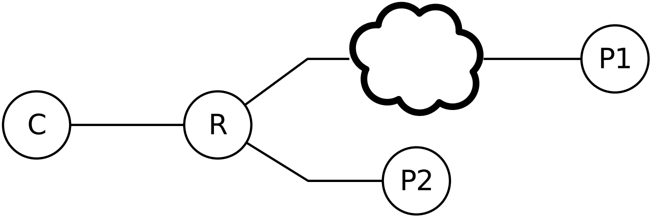

Mixing RTT measurements from different sources
⇒ Problem: Traditional RTO settings often **too short**

1. Use **Route-labels** to know content origin and path [3]
   - Still don't know where next Interest will go! [7]

Multiple Paths and Endpoints

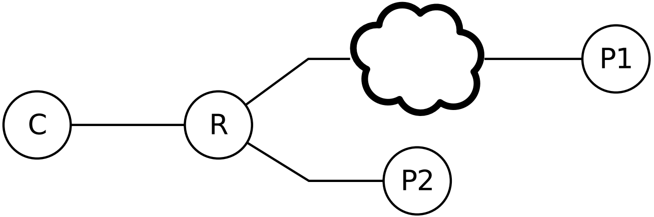

Mixing RTT measurements from different sources
⇒ Problem: Traditional RTO settings often **too short**

1. Use **Route-labels** to know content origin and path [3]
    - Still don't know where next Interest will go! [7]
2. Predicting location of **future data** [12, 1]
    - Routers mark Data to indicate their content
    - Overhead? Reliable?

Hop-By-Hop Interest Shaping

At each hop: Shape Interests to control returning Data.

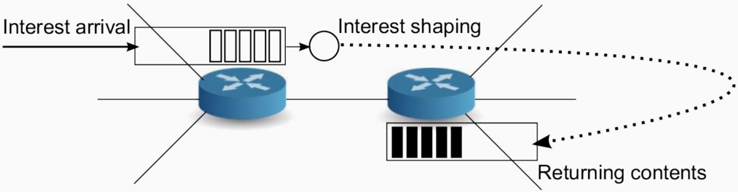

Source: Wang et al. - An Improved HBH Interest Shaper for NDN [14]

Much work [2, 14, 11, 4, 17, 16, 10, 5] *based on that principle!*

Hop-By-Hop Interest Shaping

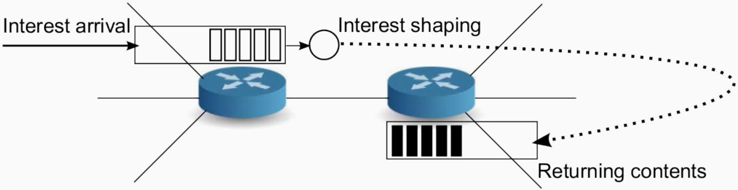

HBH Interest Shaping assumes that you

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

Estimation errors cost performance!

PCON: Design Principles

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

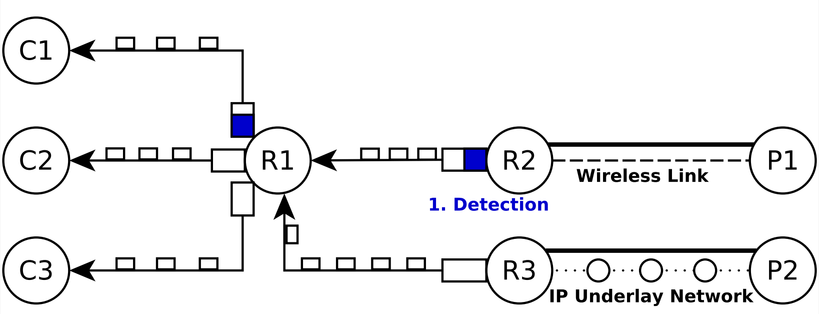

System Design: Overview

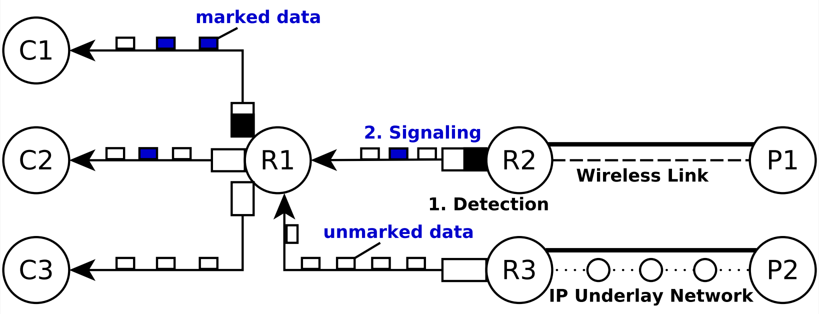

System Design: Overview

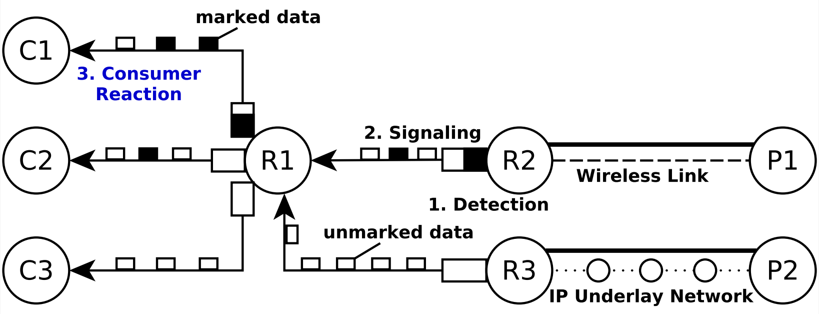

System Design: Overview

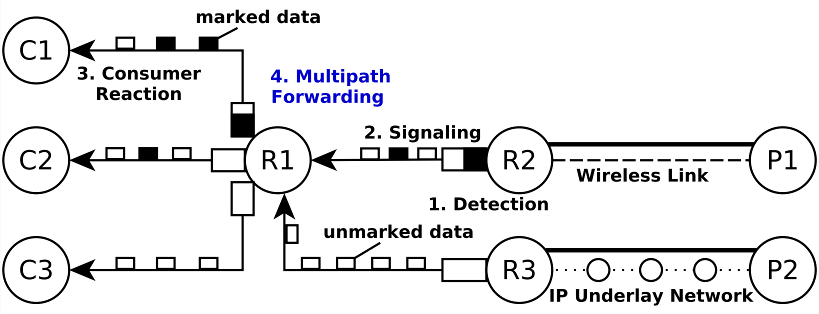

System Design: Overview

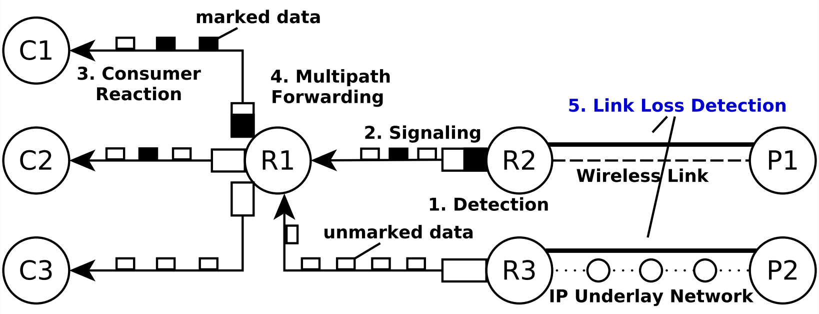

System Design: 1. Congestion Detection

Based on **CoDel AQM** [9, 8] (or any other AQM)

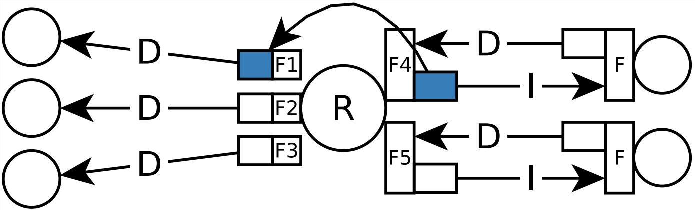

Monitor both **downstream** and **upstream** direction!

System Design: 2. Congestion Signaling

Signaling = Marking NDN **Data packets** with congestion bit.

- Using CoDel's drop spacing logic

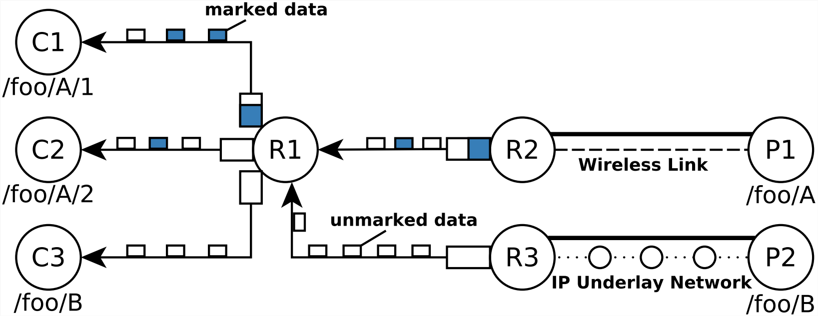

System Design: 3. Consumer Reaction

AIMD window adaptation

- We use TCP BIC [15], but any loss-based TCP algorithm works
- Window decrease on *marked Data*, *NACK*, and timeout.

PCON **removes** traditional sources of packet drops!

- Router buffer overflows
- Drops by AQM mechanism
- Drops by the “link” (UDP tunnel, Wireless)

⇒ Allows to use **higher RTOs!**

System Design: 4. Multipath Forwarding

Adjust the forwarding ratio at each router

- Related work: based on RTT or Pending Interests
- PCON: based on **congestion marks**

Start on **shortest path**; when link congested, divert traffic!

$$fwPerc(F) -= CHANGE\_PER\_MARK$$
$$fwPerc(\bar{F}) += \frac{CHANGE\_PER\_MARK}{NUM\_FACES - 1}$$

When congestion disappears, **shift back to shortest path!**

System Design: 4. Multipath Forwarding Example

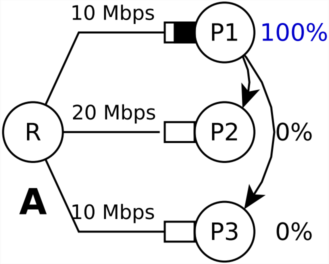

System Design: 4. Multipath Forwarding Example

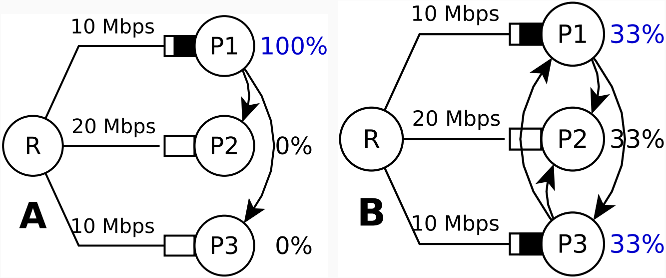

System Design: 4. Multipath Forwarding Example

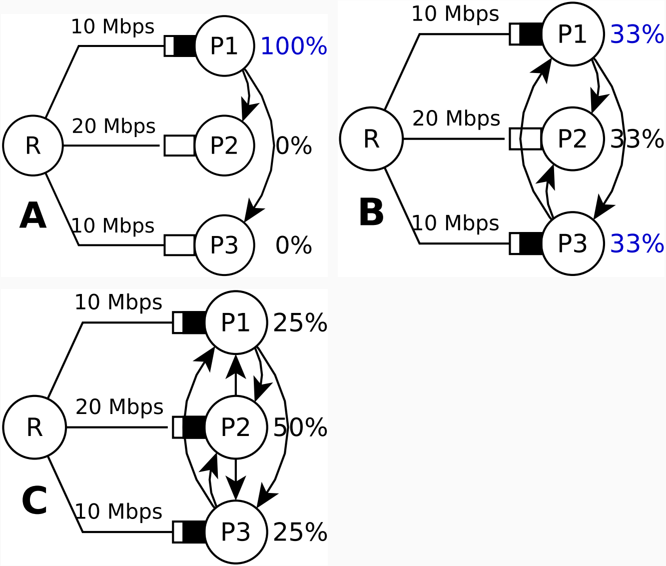

System Design: 4. Multipath Forwarding Example

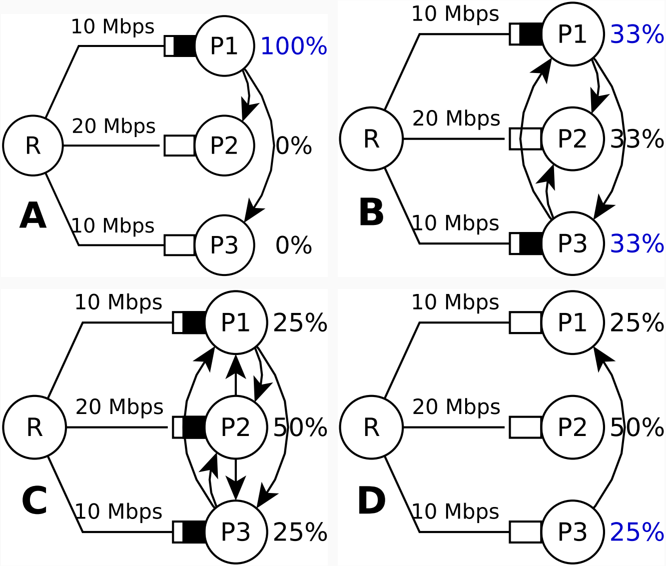

System Design: 5. Local Link Loss Detection

Problems in diverse deployment scenarios:

- **Wireless Links:** Lose packets unrelated to congestion
- **IP Overlays:** UDP tunnels lose packets without notice

Solution: Detect packet loss with a **shim layer** based on positive ACKs [13]; signal consumer with NACK

- **Unmarked NACK:** Only retx, no window adaptation
- **Marked NACK:** Both retx and window adaptation

Evaluation: Caching & Multicast

PCON vs. **CoDel** dropping queues (traditional RTO timer)

Both consumers request same data; C2 starts at 5s.

Evaluation: Caching & Multicast

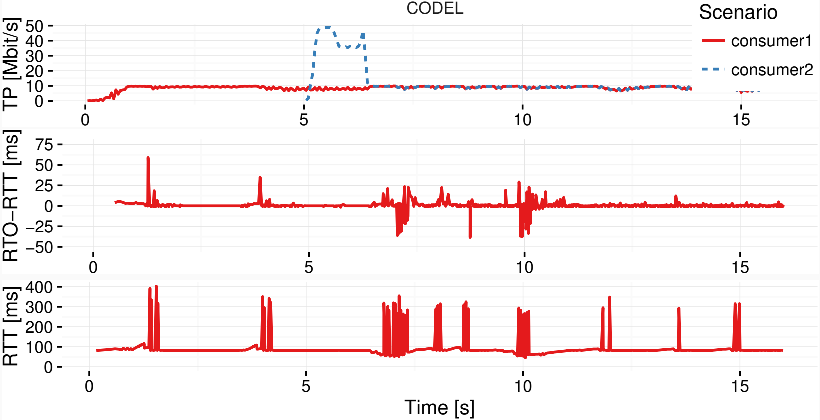

Results: TCP Timers cause spurious retransmissions!

Evaluation: Caching & Multicast

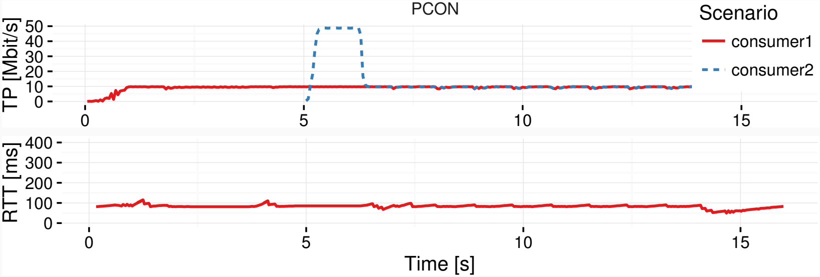

Results: TCP Timers cause spurious retransmissions!

- PCON can use a **fixed higher RTO!**

Evaluation: Multipath Forwarding

Compare PCON against PI-based forward adaptation

- **PI:** Choose face with minimum PI
- **CF** [3]: Weighted round-robin, based on PI
- **PCON:** Adapt to congestion marks

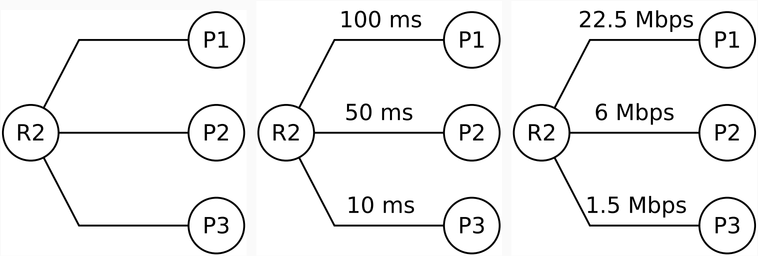

Equal Split **Diff Delay** **Diff BW**

Evaluation: Multipath Forwarding

Split Ratio at R2:

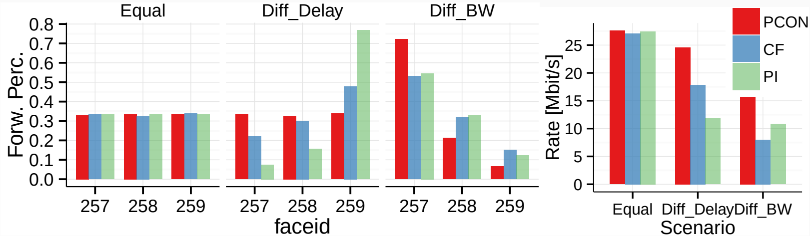

- PI and CF bias against High-Delay and High-BW paths [7]

PCON achieves **optimal split**, i.e, **maximizes throughput!**

Evaluation: PCON vs. ICP

PCON vs. **ICP** with Route-labeling, RAAQM, and CF [3]

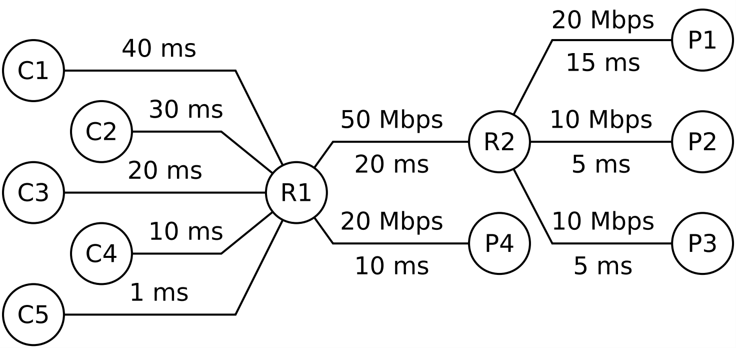

Consumers start with path C–R1–R2–P1.

Evaluation: PCON vs. ICP

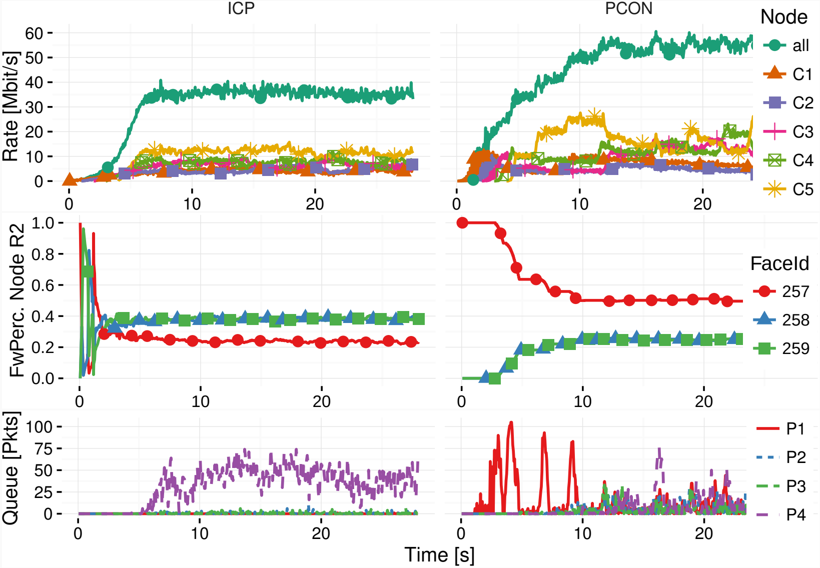

Evaluation: PCON vs. ICP

	ICP		PCON	
	RTT [ms]	Rate [Mbps]	RTT [ms]	Rate [Mbps]
C1	128.31	4.48	132.08	5.53
C2	107.83	5.16	112.37	6.77
C3	88.28	6.34	92.26	9.32
C4	68.01	7.86	72.19	12.69
C5	48.21	12.20	52.46	20.52
All	**78.08**	**7.21**	**79.23**	**10.96**

Table 1: Mean RTT and Rate per Consumer

Trade-off: Throughput vs. Latency

Evaluation: IP Overlay & Wireless Links

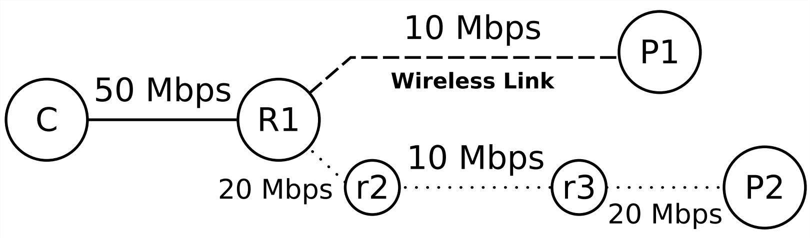

Compare against simplified HBH Interest Shaping [14]

- **Shaper_Ideal:** The shaper at R1 magically knows the link capacity in the underlay network (10 Mbps).
- **Shaper_Overlay:** The shaper uses its local link bandwidth (20 Mbps) as shaping parameter.
- **PCON:** As described earlier.

Evaluation: IP Overlay & Wireless Links

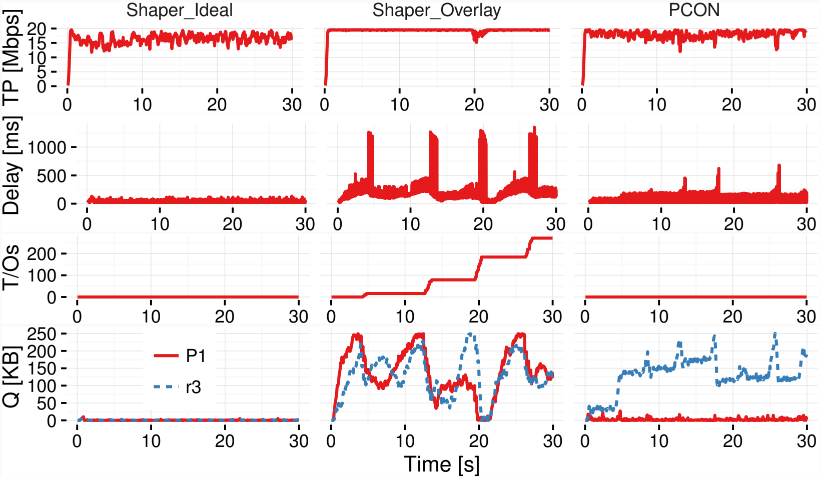

PCON doesn't drain queues, but still **avoids timeouts!**

Conclusions & Future Work

PCON prevents congestion in **diverse scenarios** (WiFi & IP Overlay) **without strong assumptions** about the network.
Novel forwarding adaptation based on congestion marks.

Future Work:

1. Definition of Fairness; handling unresponsive consumers
2. Larger evaluation; parameter setting and dynamics of congestion reaction
3. Implementation in NFD
    - `http://redmine.named-data.net/issues/3636`
    - Consumer/Producer API [6]

Thank you for your attention!

Klaus Schneider
klaus@cs.arizona.edu

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] Ilya Moiseenko, Lijing Wang, and Lixia Zhang.
Consumer/ producer communication with application level framing in named data networking.
In *ACM ICN*, 2015.

[7] Dinh Nguyen, Masaki Fukushima, Kohei Sugiyama, and Atsushi Tagami.
Efficient multipath forwarding and congestion control without route-labeling in ccn.
In *IEEE ICCW*, 2015.

[8] K Nichols, V Jacobson, A McGregor, and J Iyengar.
Controlled delay active queue management: draft-ietf-aqm-codel-03.
RFC draft, 2016.

[9] Kathleen Nichols and Van Jacobson.
Controlling queue delay.
ACM Communications, 2012.

[10] Heungsoon Park, Hoseok Jang, and Taewook Kwon.
Popularity-based congestion control in named data networking.
In *IEEE ICUFN*, 2014.

[11] Natalya Rozhnova and Serge Fdida.
An extended hop-by-hop interest shaping mechanism for ccn.
In *IEEE GLOBECOM*, 2014.

References III

[12] Lorenzo Saino, Cosmin Cocora, and George Pavlou.
Cctcp: A scalable receiver-driven congestion control protocol for ccn.
In *IEEE ICC*, 2013.

[13] Satyanarayana Vusirikala, Spyridon Mastorakis, Alexander Afanasyev, and Lixia Zhang.
A best effort link layer reliability scheme.
Technical report, NDN TR41, 2016.

[14] 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.

[15] Lisong Xu, Khaled Harfoush, and Injong Rhee.
Binary increase congestion control (bic) for fast long-distance networks.
In *IEEE INFOCOM*, 2004.

[16] 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.

References IV

[17] 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.