

Joint Hop-by-hop and Receiver-Driven Interest Control Protocol for Content-Centric Networks

Giovanna Carofiglio

Bell Labs, Alcatel-Lucent

giovanna.carofiglio@alcatel-lucent.com

Massimo Gallo

Orange Labs, France Telecom

massimo.gallo@orange.com

Luca Muscariello

Orange Labs, France Telecom

luca.muscariello@orange.com

ABSTRACT

Content-centric networking (CCN) advocates a new transport model tailored to named-data communication. Three features distinguish CCN transport from the TCP/IP model: unique endpoint at the receiver, pull-based data retrieval in a point to multi-point fashion and in-path caching.

The definition of transport control mechanisms is of fundamental importance within the CCN architectural design and beyond, in the broader scope of information-centric networks. In this work, we propose a joint Hop-by-hop and Receiver-driven Interest Control Protocol (HR-ICP) to regulate user requests (Interests) either at the receiver and at intermediate nodes via Interest shaping. We prove that HR-ICP is stable and converges to an efficient and max-min fair equilibrium. Compared to controlling traffic only at the receiver, HR-ICP accelerates congestion reaction and reduces the loss rate, as we show by means of CCN packet-level simulations. In different network scenarios, we highlight the advantages of our solution in terms of faster convergence to the optimal throughput, robustness against misbehaving receivers and flow protection of delay-sensitive applications.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols

General Terms

Performance

Keywords

Information-Centric Networking; Traffic control; Transport;

1. INTRODUCTION

Information dissemination and retrieval throughout the Internet is defining a novel communication model, in contrast with the founding host-to-host IP architecture principles. Content-centric communication is naturally pull-based,

with a receiver issuing requests for a specific content, in a point-to-(multi)point fashion towards network nodes susceptible of storing a copy of the given item. Today content delivery systems (as CDNs, P2P systems, transparent caching solutions) realize such functionalities by defining overlays on top of the IP network infrastructure. In parallel, novel future Internet proposals are emerging in the research arena to rethink the network primitives around the need of an information-centric communication model [1, 6, 14, 7]. The definition of a suite of traffic control mechanisms adapted to CCN (Content-Centric Networking) still lacks in the literature. In this paper, we focus on CCN flow control and build on our previous work in [3], where we introduce a receiver-based Interest Control Protocol (ICP) to regulate user's requests (Interest) over time. The aim of the present work is to couple ICP with a hop-by-hop Interest shaping mechanism suited to CCN, to realize a faster congestion avoidance, while guaranteeing fair and efficient bandwidth and storage sharing.

Hop-by-hop flow control mechanisms have been extensively studied in the past, as an alternative or in conjunction with end-to-end congestion control, to overcome the difficulties of the sender-driven rate-based flow control in handling bursty traffic. [8], [9] observe that, in presence of bursty traffic, rate adaptation performed at the sender can be imprecise and slow to converge to the optimal value. Authors propose a credit-based flow control at each node, regulating data packet forwarding according to a credit balance. Each time a data packet is received, the credit balance is incremented by one unit, while a data packet forwarding corresponds to a decrement of one unit. The results present the performance improvement w.r.t end-to-end flow control only. In [10], authors develop a hop-by-hop rate-based congestion control mechanism targeted to bursty traffic, as an alternative to its end-to-end version. However, the objective is different: intermediate nodes forward back to the sender information about congestion status and, ultimately, the sender adapts its rate, based on the received feedbacks.

In the context of CCN, or more generally of a pull-based network, such kind of mechanisms are particularly appropriate when applied to the requests' flows, as argued in [5],[11]. Our main contribution is the design of a joint hop-by-hop and Receiver-driven Interest Control Protocol (HR-ICP) suited to CCN. We first analyze HR-ICP stability and steady state equilibrium, in comparison with the case of ICP only, by extending the modeling framework presented in [3]. Further, by means of packet-level simulation we show that

i) the receiver-driven controller is needed to guarantee full

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICN'12, August 17, 2012, Helsinki, Finland.

Copyright 2012 ACM 978-1-4503-1479-4/12/08 ...\$15.00.

bandwidth utilization; ii) HR-ICP provides significant benefits w.r.t. ICP in terms of reduced loss rate, as well as protection against misbehaving receivers; iii) the hop-by-hop mechanism allows to isolate delay-sensitive traffic flows.

The remainder of the paper is organized as follows. In Sec.2 we introduce ICN system description, while Sec.3.1 describes HR-ICP design and implementation. Sec.4 gathers analytical results on HR-ICP stability and steady state equilibrium. The benefits of hop-by-hop Interest shaping are quantified by packet-level simulations in Sec.5, and discussed in Sec.6, in comparison to existing alternatives.

2. SYSTEM DESCRIPTION

In this paper we primarily focus on CCN/NDN proposal [6, 14], though the defined traffic control mechanisms have broader applicability in the context of Information-Centric Networking (ICN[1]) and more generally to cache networks employing a similar transport model (e.g. HTTP based CDNs). Let us briefly describe the main features of CCN.

Content items (data) are split into packets, identified by a unique name, and permanently stored in one (or more) repository(ies). Users retrieve a content item using a receiver driven flow control protocol based on successive packet queries (denoted as *Interests*), triggering data packets delivery in the reverse path (see Fig.1).

A *name-based routing protocol* guarantees that queries are properly routed towards a data repository, following one or multiple paths. Every intermediate node keeps track of outstanding queries in data structures called PITs (Pending Interest Tables), to deliver the requested data back to the receiver on the reverse path. To each entry in the PIT, a timer is associated, so that, during such time interval, all requests for the same Data are not forwarded upstream, as long as the first query is outstanding. In addition, nodes temporarily cache Data packets in a local cache, named *Content Store*.

Upon reception of an Interest packet from an input interface, intermediate nodes perform the following operations : (i) a *Content Store lookup*, to check if the requested Data is locally stored. In case of cache hit, the Data is sent through the interface the Interest is coming from. Otherwise, (ii) a *PIT lookup*, to verify the existence of an entry for the same content name. In this case, the Interest is discarded since a pending request is already outstanding. If not, a new PIT entry is created and (iii) a *FIB lookup* returns the interface where to forward the Interest (selected among the possible ones). As a consequence, Data may come from the repository, or from any intermediate cache along the path storing a temporary copy of the data packet. Packets of the same content item can therefore be retrieved in a multi-path fashion, i.e. from multiple locations with different round trip times, affecting the overall delivery performance.

3. HR-ICP DESIGN

3.1 ICP description

In [3], a receiver-driven Interest Control Protocol for CCN is presented whose description is hereafter summarized. Under ICP, Data are requested via Interests (one Interest per Data packet) in the order decided by the application, according to a window-based Additive Increase Multiplicative Decrease (AIMD) mechanism. The receiver window W , is kept at the receiver and defines the maximum number of

outstanding Interests the receiver is allowed to send. W is increased by η/W upon each Data packet reception. This corresponds to an increment of η ($\eta = 1$ as default) each time a complete W of Interests is acknowledged by Data reception. When an Interest is sent out, a timer is set at the receiver (to manage losses) and its expiration is interpreted by ICP as a *congestion* event, to which the protocol reacts by multiplying W by a decrease factor $\beta < 1$, no more than once in a time interval of the same duration of the timer. Properly setting the Interest timer value τ and the PIT timer (which allows to remove a pending interest at each node) is crucial in CCN.

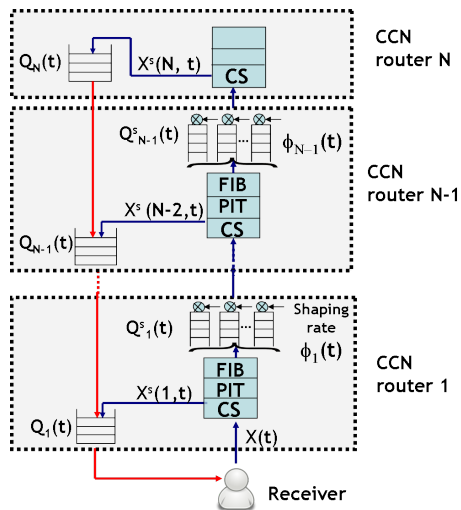


Figure 1: N hops network topology.

Indeed, Interest re-expression after a time-out is necessary to recover losses, even though Data may just be delayed at the bottleneck. At the same time, retransmitted Interests sent just after a time-out can be filtered, if the PIT timer is bigger than the Interest timer at the receiver. The PIT timer is necessary to limit the number of pending Interests, however, it makes loss recovery slower. The larger this value the higher the number of Interests not forwarded upstream towards content repository. On the other hand, a small PIT timer implies a large number of unnecessary retransmissions, since delayed packets arriving after PIT time-outs are discarded. According to the arguments provided in [3], each flow adapts its timer τ with:

$$\tau = RTT_{min} + \delta(RTT_{max} - RTT_{min}) \quad (1)$$

with $\delta < 1$ (0.5 for instance). The protocol maintains round trip delay estimates at every data packet reception, by updating RTT_{min} and RTT_{max} over a history of samples (e.g. 30), excluding retransmitted packets. A complete history of RTT samples is necessary before triggering a window decrease. The stability analysis of delay based congestion control protocol using RTT estimation has been studied in the literature [12]. We leave for future work the proof for ICP under adaptive timer.

3.2 Hop-by-Hop Interest Shaping

This section introduces our hop-by-hop Interest shaping mechanism, which realizes per-flow Interest rate control at the output interfaces of every CCN router (as illustrated in Fig.1). The rationale behind an additional Interest control realized at every network node is (i) to anticipate congestion detection and (ii) to trigger rate reduction via Interest

shaping, before the receiver can detect the congestion via a timer expiration. In every output interface, we maintain one virtual queue per flow, identified by the name of the content (i.e. the flow identifier is a hash of the name). Each virtual queue is associated a credit counter initialized to a maximum value B (in bytes), indicating the number of Data bytes the flow is granted to transmit, with no additional delay. The counter is incremented at the estimated fair rate of the corresponding down-link and decremented by forwarded Interests. The evolution of virtual queue and counter is driven by the algorithm reported in Algo.1. As a note, flow refers to a single content retrieval, and is defined as bottlenecked when at least one Interest packet is queued at the output buffer or the credit counter is null. Notice that keeping one virtual queue per flow implies to maintain a flow counter as long as there are outstanding Interests in the PIT.

Algorithm 1 Virtual queue f management

```

enqueue (interest, f)
  bottlenecked[f] = (length(f)>0 OR counters[f]<= 0)
  if (bottlenecked[f])
    shaping_queue_tail(interest, f)
  else
    link.enqueue(interest,t=0)
    update_rate(rate, interest.data_size)
  endif
  counter[f]=max(0, counter[f]-interest.data_size)

dequeue()
  interest = shaping_queue_head( current_q )
  reduced_rate = downlink_rate - rate
  link.enqueue(interest,t=interest.data_size/reduced_rate)
  current_q = next_q

```

Different implementations of the flow table are possible, e.g. a hash table per content name storing a pointer to an Interest shaper for each output interface. This implementation avoids to have additional name lookups and can be directly included in the PIT. Indeed, HR-ICP exploits the PIT information about ongoing flows, by just introducing a credit counter per active flow. In this way, the scalability with the number of ongoing flows is preserved. When an Interest arrives at the output interface, the algorithm checks if it belongs to a bottlenecked flow or not. If the flow is non-bottlenecked, the Interest is directly sent upstream and the counter decremented by the number of bytes of the corresponding Data packet. For a bottlenecked flow, the Interest is queued in a drop tail FIFO of size Q_{max} served at the shaping rate. The shaping rate, $\phi_i(t)$, of node i at time t , is set to the estimated fair rate, as we detail in the next section, in eq.3, i.e. the link capacity reduced by the rate of non-bottlenecked flows, equally divided among bottlenecked flows. The latter is estimated by the number of non empty waiting queues, while the rate of non-bottlenecked flows is approximated by the total rate of non shaped flows, where we count every non shaped Interest by the size of the corresponding Data packet in bytes. Such size value is deductible from the Interest by using the segmentation information in the name (e.g. in the CCNx prototype), however, a standard component or field would be preferable. The shaping mechanism can be implemented by using virtual times (based on the fair rate) or by round robin based implementation. The latter is detailed in the following pseudo-code.

4. HR-ICP ANALYSIS

In this section, we show that HR-ICP is stable and converges to the same equilibrium of ICP in steady state on average, in absence of packet losses. Such result is important to state that proactive shaping of Interests does not introduce additional delay and does not prevent, in absence of losses, the convergence to the efficient and fair average throughput equilibrium guaranteed by ICP. Let us consider a single flow traversing the network in Fig.1 composed by N hops between the user and the repository. Each intermediate node implements the Interest shaping mechanism described in Sec.3.2.

We describe the instantaneous time evolution of the rate $X(t)$, of a flow controlled by ICP at the receiver, of the output queues occupancy, $Q_i(t)$ and of the shaping queues occupancy, $Q_i^s(t)$. Like in [3], we express the flow rate $X(t)$ as the weighted average, by the stationary miss probabilities p_i , of the rates associated to route i (i.e. sub-path between the user and the i^{th} cache),
 $X(t) = \sum_{i=1}^N X^s(i, t) \prod_{j=1}^{i-1} p_j (1 - p_j)$. $X(i, t)$ is replaced by its "shaped" version, as a result of the Interest shaping performed at node i (see Fig.1). Notice that the receiver does not implement Interest shaping, $X^s(i, t) \equiv X(i, t)$. The temporal evolution of $X(t)$ driven by ICP (with constant τ) and that of the output queues $Q_i(t)$ is a slight modification of eq.(6) in [3], where we replaced $X(i, t)$ with $X^s(i, t)$.

$$\begin{aligned} \dot{X}^s(i, t) &= \frac{\eta}{\text{VRTT}(i, t)^2} - \beta X^s(i, t) \mathbb{1}_{\{\text{VRTT}(i, t) = \tau\}}, \\ \dot{Q}_i(t) &= X^s(i, t) - C_i \mathbb{1}_{\{Q_i(t) > 0\}} \\ &+ \sum_{l=i+1}^N \left(\prod_{j=i+1}^l \mathbb{1}_{\{Q_j(t)=0\}} X^s(l, t) + \mathbb{1}_{\{Q_i(t) > 0\}} C_l \right) \end{aligned} \quad (2)$$

where $\text{VRTT}(i, t) = \sum_{j=1}^i (Q_j(t)/C_j + Q_j^s(t)/\phi_j'(t))$, $i = 1, \dots, N$ denotes the round trip time given by the aggregated queuing delay associated to route i , i.e. all links from the receiver to the i^{th} node. The Interest shaping rate at node i is the max-min fair share $\phi_i(t)$

$$\phi_i(t) = (C_{i+1} - \sum_{j>i, j \in \mathcal{R}_i^{nb}} X^s(j, t)) / (N - \|\mathcal{R}_i^{nb}\|), \quad (3)$$

defined as the link capacity C_{i+1} minus the rate of the routes non bottlenecked at link i , i.e. in $\mathcal{R}_i^{nb} = \{\cup_{j, j \neq \underset{i' \ni i}{\text{argmin}} \gamma_{i'}}\}$ divided by the number of bottlenecked routes. The evolution of shaping queues is described by

$$\dot{Q}_i^s(t) = \sum_{j>i, j \notin \mathcal{R}_i^{nb}} X^s(j, t) - \phi_i(t) \mathbb{1}_{\{Q_i^s(t) > 0\}} \quad (4)$$

where $i = 1, \dots, N - 1$. Indeed, eq.(4) describes the shaping performed at node i on the Interest rate of routes bottlenecked in i , paced to the fair rate $\phi_i(t)$.

PROPOSITION 4.1. *In the regime described by eqq.(2)-(3)-(4), HR-ICP is stable and attains the max-min fair and efficient equilibrium of ICP[3], i.e. the steady state solution is periodical with average value*

$$\tilde{X} = \sum_{i=1}^N (1 - p_i) \prod_{j=1}^{i-1} p_j \tilde{X}^s(i), \quad \tilde{X}^s(i) = \gamma(i) \quad (5)$$

where $\gamma(i)$ is the max-min fair rate for route i (cfr. [2],[3]).

PROOF. The proof consists in showing the equivalence between eqq.(2)-(3)-(4) and the ODEs describing ICP (cfr. [3]), when $Q_i^s(t)$ replace $Q_i(t)$ in the ICP system, $\forall t \geq t^* > 0$, where t^* accounts for the time needed for the shaping queues to fill in, dependently on their initial condition. The proof follows two steps:

- 1) under HR-ICP, $Q_i(t) \leq Q_i(0)$, $\forall t > 0$, $i = 2, \dots, N$;
- 2) $\forall t > t^*$, $Q_1(t) = Q_1^{ICP}(t)$ and $Q_i^s(t) \equiv Q_{i+1}^{ICP}(t)$.

1) Let us focus on the output queues $Q_i(t)$ under HR-ICP. Due to the Interest shaping performed at node N , the input rate of queue Q_N at time t is $X_N^s(t)$ and by definition of the shaping rate $\phi_i(t)$, $X_N^s(t) \leq C_N$. Thus, $\dot{Q}_N < 0$ and $Q_N(t) \leq Q_N(0)$, $\forall t > 0$. Similarly for queue $Q_{N-1}(t)$,

$$\begin{aligned} X_N^s(t) \mathbb{1}_{\{Q_N(t)=0\}} + C_N \mathbb{1}_{\{Q_N(t)>0\}} + X_{N-1}^s(t) \\ \leq C_N + (C_{N-1} - C_N) = C_{N-1} \end{aligned}$$

For $Q_i(t)$, $i > 2$ in general

$$\begin{aligned} \sum_{l=i+1}^N \left(\prod_{j=i+1}^l \mathbb{1}_{\{Q_j(t)=0\}} X^s(j, t) + \mathbb{1}_{\{Q_l(t)>0\}} C_l \right) + X_i^s(t) \\ \leq C_i + (C_{i+1} - C_i) = C_i \end{aligned}$$

Notice that the inequality does not hold for $Q_1(t)$, since $X(1, t)$ is not shaped.

2) Let us focus on a three links scenario like in [3]. The generalization to the case of $N > 3$ is straightforward, but lengthy (we need to distinguish more cases to explicitly compute the max min rate). Under HR-ICP, the queues evolution satisfies the following ODEs according to the configurations listed below.

Case I): $C_1 < C_2, C_3$ - In this case all routes are bottlenecked at link 1, therefore there is no shaping at hop 2 and 3 and

$$\dot{Q}_1(t) = X(t) - C_1 \mathbb{1}_{\{Q_1(t)>0\}}, \quad Q_1^s(t) = Q_2^s(t) = 0$$

like for ICP only. This leads to $\tilde{X}(i) = C_1$, $i = 1, 2, 3$.

Case II): $C_2 < C_1, C_3$ - In this case route 2 and 3 are bottlenecked at link 2, thus the shaping rate for the two at the first hop is $\phi_1(t) = C_2$.

$$\begin{aligned} \dot{Q}_1(t) dt &= X(1, t) + C_2 - C_1 \mathbb{1}_{\{Q_1(t)>0\}}, \\ \dot{Q}_1^s(t) &= X(2, t) + X(3, t) - C_2 \mathbb{1}_{\{Q_1^s(t)>0\}}, \quad Q_2^s(t) = 0. \end{aligned}$$

It suffices to compare the ODE of $Q_1^s(t)$ and that of $Q_2(t)$ in ICP, to conclude about their equivalence and, hence, about the equivalence of their solution.

Case III): $C_3 < C_2 < C_1$ - In this configuration, the ODEs of HR-ICP reduces to:

$$\begin{aligned} \dot{Q}_1(t) &= X(1, t) + C_2 - C_1 \mathbb{1}_{\{Q_1^s(t)>0\}}, \\ \dot{Q}_1^s(t) &= X(2, t) - (C_3 - C_2) \mathbb{1}_{\{Q_1^s(t)>0\}}, \\ \dot{Q}_2^s(t) &= X(3, t) - C_3 \mathbb{1}_{\{Q_2^s(t)>0\}}. \end{aligned}$$

where we used $\phi_1(t) = C_2 - C_3$ and $\phi_2(t) = C_3$ as route 3 is bottlenecked at link 3. This implies $Q_1^s(t) = Q_2^{ICP}(t)$, $Q_2^s(t) = Q_3^{ICP}(t)$.

Case IV): $C_3 < C_1 < C_2$ - Similarly to previous cases, we can specialize eqq.(2)-(3)-(4),

$$\begin{aligned} \dot{Q}_1(t) &= X(1, t) + X(2, t) + C_3 - C_1 \mathbb{1}_{\{Q_1(t)>0\}}, \\ \dot{Q}_2^s(t) &= X(3, t) - C_3 \mathbb{1}_{\{Q_2^s(t)>0\}}. \end{aligned}$$

This implies $Q_1^s(t) = Q_2^{ICP}(t)$, $Q_2^s(t) = Q_3^{ICP}(t)$. The extension to the case $N > 3$ hops is straightforward as in the case of ICP (cfr.[3]). \square

5. EVALUATION

To assess HR-ICP performance, we have implemented the Interest shaping mechanism in CCNPL-Sim (available at <http://perso.rd.francetelecom.fr/muscariello/sim>) a C++ event driven simulator implementing the CCN system: naming, node packet-level caching, name-based forwarding and routing via PIT and FIB data structures: FIB entries are pre-computed according to a shortest path name-based routing protocol (described in [4]). At the receiver we implemented ICP ([3]) and two mechanisms for Interest dispatch at constant rate (CBR) and with a constant window. In this section, we gather a selected set of simulation results illustrating HR-ICP functioning, properties and benefits.

Scenario 1. *Hop-by-Hop Interest shaping is not enough.*

In this scenario we show i) the need for hop-by-hop Interest shaping to be coupled with a receiver-driven Interest control adapting Interest window over time (as ICP), ii) the loss rate reduction realized by our Interest shaping.

To this aim, we consider a two hops network (Fig.3) and focus on a single data retrieval for a content item of 5MB. We set link capacities to $\{C_1, C_2\} = \{100, 40\}$ Mbps, propagation delay $1ms$, 10kB data buffers and Interest shaper with burst size $B = 2kB$, $\Delta T = 100ms$, and infinite queuing capacity (large enough to have no Interest discard). We compare two receivers: one implementing ICP and the other using a constant Interest window W and a constant Interest re-expression timer set to 10ms. Data packet size is 1kB, while Interest packet size is 25B. The results of the comparison are reported in Tab.1. Using ICP at the receiver the delivery time is optimal, equal to 1s and, given the absence of losses, the same average behavior is observed with and without Interest shaping, as proved in Sec.4.

W	Delivery Time[s]		Throughput[Mbps]		Losses[%]	
	w HbH	w/o	w	w/o	w	w/o
2	2.42	2.42	16.30	16.30	0	0
10	1.00	1.00	39.70	39.60	0	0
15	1.00	2.08	39.60	19.20	0	11.20
20	1.00	1.90	39.60	20.90	0	15.30
ICP	1.00	1.00	39.80	39.80	0	0

Table 1: Scenario 1.

On the contrary, in presence of a constant Interest window, the delivery performance depends on W value. A small constant window value W causes inefficient content retrieval due to resources under-utilization (i.e., for $W = 2$, the delivery time is 2.42s). With $W = 10$ the link is efficiently utilized and the delivery time is 1.0s, while increasing W to 15 – 20 induces Data packet losses and leads to higher delivery time, when no hop-by-hop mechanism is used.

In this case, Interest shaping allows to reduce loss rates by queuing Interest packets before the bottleneck. However, we conclude that the shaping mechanism itself is not sufficient to guarantee optimal delivery performance, as the optimal value of W is not known a-priori and should vary according to available network resources.

Scenario 2. *The benefits of HR-ICP over ICP*

We consider a second simulation scenario, in the same network set-up described above. By varying the number of

flows over time as shown in Fig.2, we observe i) how HR-ICP queues Interests before the bottleneck link and ii) how it reacts to a congestion phenomenon arising at a given time instant. A single ICP flow starts at time zero, while a second one begins at 0.5s. In the same simulation run, a greedy CBR flow arrives at $t=1.0s$, with rate equal to link capacity $C_2 = 40Mbps$, sharing the downlink capacity of the bottlenecked link with the two ICP flows.




Figure 3: Network topology 1: nodes 1,2; network topology 2: nodes 1,...,4.

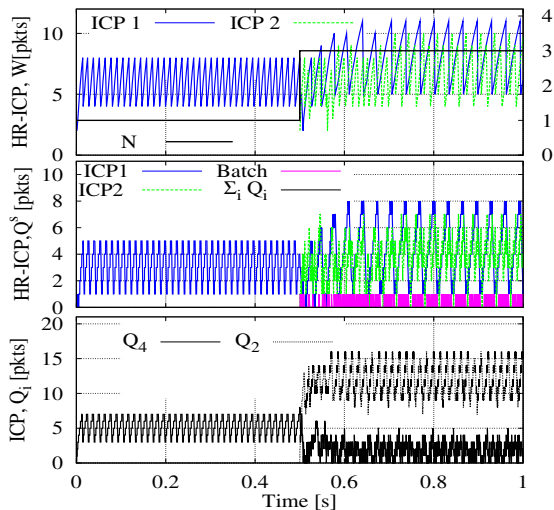


Figure 4: Scenario 3.

Before the arrival of the CBR flow, no losses are observed and, as expected from the analytical predictions in Sec.4, ICP and HR-ICP have the same stationary regime, given by the max-min rate allocation. This can be observed by comparing the time evolution of $Q_2(t)$ and $Q_1^s(t)$, respectively for ICP and HR-ICP, in Fig.2. Under ICP, the output FIFO queue, $Q_2(t)$ is displayed, while under HR-ICP we report the two per-flow queues in the up-link, $Q_1^s(t)$, since $Q_2(t) = 0$. The total queuing is the same in the two cases, though in HR-ICP the buffering is distributed in two per-flow buffers scheduling Interest transmissions at the fair rate.

At $t=1s$, when the CBR flow starts, the two ICP flows observe losses and their windows drastically reduce in absence of Interest shaping. Conversely, with our Interest shaping mechanism, the CBR flow is immediately shaped and gets a penalty of 66% of Interest packets loss rate, achieving a download rate equal to its fair share. In conclusion, HR-ICP is shown to protect conformant ICP traffic flows from greedy misbehaving receivers sending Interests at a rate higher than their fair rate. Indeed, HR-ICP prevents all ICP Data packet losses at the bottleneck, by early reacting to the congestion caused by the greedy CBR flow.

Scenario 3. HR-ICP in a multi-bottleneck scenario.

In this third simulation set, we consider four links including an additional bandwidth bottleneck, link 4, with capacity $C_4 = 20 Mbps$, and $C_3 = 100Mbps$ as in Fig.3. Two ICP flows retrieve Data from the second and fourth nodes, respectively. The first ICP flow, starting at time 0, is bottle-

necked at link 4, whereas the second, starting at $t=0.5s$, is bottlenecked at link 2. As a first observation, HR-ICP creates Interest queuing for each flow before its bottleneck, by leaving empty the shaper queues at other nodes and all Data output queues. Therefore, in Fig.4 we plot $Q_3^s(t)$ for ICP 1 and $Q_1^s(t)$ for ICP 2 under HR-ICP, which correspond to $Q_4(t)$ and $Q_2(t)$ under ICP only. The evolution of the windows is almost the same with and without Interest shaping and coincides with that reported in Fig.4. After the first 0.5s, we introduce a new flow, sending 4 Interests in batch every 10ms, which corresponds to an average Data delivery of 5Mbps, with a peak rate of 100Mbps. The Interest window of the two ICP flows is slightly reduced according to the fair share reduction induced by the new flow, but their rate is almost not affected. In the HR-ICP system, with hop-by-hop Interest shaping, the new flow gets priority along the request path and does not experience queuing in the Data path. Instead, in absence of Interest shaping, the new flow experiences a significant queuing along the Data path given by $Q_2 + Q_4$. Such example shows a desirable property of our shaping mechanism, i.e. prioritization of real-time and delay sensitive traffic, whose rate is lower than the fair rate.

6. DISCUSSION AND CONCLUSIONS

In the paper, we introduced an Interest control mechanism for CCN, HR-ICP, jointly realized by the end-user and within the network by in-path routers. We believe this is an essential building block of the CCN architecture, whose definition needs to exploit the specific features of Content-Centric communication. We have shown that the receiver plays an important role in the rate control loop to guarantee full bandwidth/storage resources utilization and flow fairness. Hop-by-hop Interest shaping enhances rate and congestion control performance and it is particularly suited to CCN for various reasons.

Interest not Data control. Controlling Interests instead of Data packets (as done in [13]), gives the opportunity to prevent congestion onset by delaying Interest forwarding. Interest packets are smaller in size than Data packets, hence they require smaller buffer capacity and with reasonable buffer dimensioning Interest losses may be avoided, except in case of greedy misbehaving receivers.

Early congestion detection. It allows to realize early congestion detection by locally monitoring at each node the per-flow Interest rate in uplink and the corresponding Data fair rate on the downlink. A discrepancy between an Interest rate larger than its associated Data rate signals the beginning of a congestion phenomenon, before the detection of packet losses at the receiver.

Protection from misbehaving receivers. By shaping Interest in a hop-by-hop fashion, the greedy behavior of a non-conformant receiver can be fast detected and controlled in order to protect concurrent flows. HR-ICP and the solution proposed in [5] react to misbehaving receivers by queuing Interests up to a certain threshold, before discarding them. In [11], instead, all Interests exceeding in rate the corresponding Data fair rate are directly discarded. Compared to alternative solutions for CCN Interest control [5], [11] our proposal brings additional benefits due to: (i) the coupling with a rate/fairness optimal receiver control, (ii) Interest shaping at output interfaces, (iii) max-min rate by nly controlling Interests.

Scalability/Feasibility. HR-ICP exploits the PIT in-

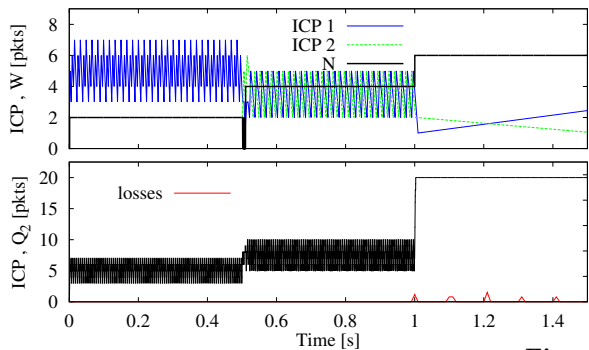



Figure 2: Scenario 2

formation about ongoing flows, by just introducing a credit counter per active flow. Thus, scalability is bounded by the PIT size, differently from the solution presented in [5] where the estimation of the response delay $A(t)$ is required per each flow based on an history of samples.

Delay-sensitive flows protection. HR-ICP brings a considerable reduction of the Data queuing at the bottleneck, by queuing Interests at output buffers (ideally, under the fluid model assumption the queue totally empties). As a consequence, delay-sensitive flows like streaming or voice flows, whose average rate is lower than the fair rate see no Interest queuing and low/zero Data queuing delay due to the Interest shaping of bottlenecked flows. This is a distinctive of our Interest shaping mechanism w.r.t. the solution proposed in [5], where Interests are FIFO queued and every ongoing flow experiences the overall Data queuing delay.

No Interest losses. We have observed that, under a reasonable dimensioning of the output buffer, no Interest is dropped with HR-ICP, except for misbehaving flows. Conversely, in [11], overload control is managed on a per-flow basis through Interest discard. The incoming Interest rate, at every CCN router, is limited to the fair rate realized by the corresponding Data packets at the bottleneck, under the assumption of a per-flow scheduler implemented at each network node. Interests are discarded as long as the Interest rate exceeds the bottleneck fair rate and before any router processing (CS/PIT/FIB). We argue that this kind of Interest drop should be avoided, as it prevents the removal of the associated PIT entry at the current node (PIT not yet accessed) and at previous nodes. If the PIT entries corresponding to the discarded Interest in previous nodes are not removed, subsequent re-expressions of the Interest are filtered, as the first ones were ongoing until their expiration. In [11] authors propose an explicit signaling of the Interest packet loss back to the user, though it would require message prioritization and it would rise security concerns. Similarly, in [5], where a receiver-driven control of the Interest lacks, authors claim that some back-pressure mechanism is needed to avoid Interest losses, leaving its definition for future work.

Additional traffic control opportunities. By implementing Interest shaping at the output interfaces, HR-ICP opens additional traffic control opportunities. As a consequence of an Interest drop at a given node, our mechanism permits to trigger the removal of the corresponding entry in the local PIT or the forwarding of the Interest of another feasible interface, when available in the FIB. In the latter case, a *congestion-aware load balancing of the Interests* may be envisaged. Notice that Interest redirection is not possible when the Interest is discarded after the forwarding, like in



[11]. Also, a *measurement-based admission control* can be implemented by selectively refusing Interests of new flows, while not penalizing already ongoing flows. A natural extension of this work is the definition of traffic control mechanisms for the management of a multi-path communication, either at the receiver and hop-by-hop as coupled with an Interest forwarding policy.

Acknowledgements

This work has been partially funded by the French national research agency (ANR), CONNECT project.

7. REFERENCES

- [1] Information centric network projects. http://www.anr-connect.org/index.php?p=1_5_Related-Projects.
- [2] G. Carofiglio, M. Gallo, and L. Muscariello. Bandwidth and storage sharing performance in information centric networking. In *Proc. of ACM SIGCOMM ICN*, 2011.
- [3] G. Carofiglio, M. Gallo, and L. Muscariello. Icp: Design and evaluation of an interest control protocol for content-centric networking. In *Proc. of IEEE INFOCOM NOMEN Workshop*, 2012. Technical Report available at <http://perso.rd.francetelecom.fr/muscariello>.
- [4] A. Carzaniga, M. Rutherford, and A. Wolf. A routing scheme for content-based networking. In *Proc. of IEEE INFOCOM*, 2004.
- [5] S. Fdida and N. Rozhnova. An effective hop-by-hop interest shaping mechanism for ccn communications. In *Proc. of IEEE INFOCOM NOMEN Workshop*, 2012.
- [6] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard. Networking named content. In *Proc. of ACM CoNEXT*, 2009.
- [7] T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K. Kim, S. Shenker, and I. Stoica. A data-oriented (and beyond) network architecture. In *Proc. of ACM SIGCOMM*, 2007.
- [8] H. T. Kung, T. Blackwell, and A. Chapman. Credit-based flow control for atm networks. *ACM SIGCOMM CCR.*, 24:101–114, October 1994.
- [9] H. T. Kung and K. Chang. Receiver-oriented adaptive buffer allocation in credit-based flow control for atm networks. In *Proc. of IEEE INFOCOM*, 1995.
- [10] P. P. Mishra, H. Kanakia, and S. K. Tripathi. On hop-by-hop rate-based congestion control. *IEEE/ACM Transactions on Networking*, 4:224–239, April 1996.
- [11] S. Oueslati, J. Roberts, and N. Sbihi. Flow-aware traffic control for a content-centric network. In *Proc. of IEEE INFOCOM*, 2012.
- [12] F. Paganini, Z. Wang, J. C. Doyle, and S. H. Low. Congestion control for high performance, stability, and fairness in general networks. *IEEE/ACM Transactions on Networking*, 13(1), Feb 2005.
- [13] P. Viotti. Caching and transport protocols performance in content-centric networks. *Eurecom Master Thesis.*, 2010.
- [14] L. Zhang and al. Named data networking project.