Pursuing a Pub/Sub Internet

PURSUIT
Publish Subscribe Internet Technology
FP7-INFSO-ICT-257217

TR 13-0001

On Congestion Control in Information-Centric Networks

Document Properties:

Title of Contract	Publish Subscribe Internet Technology
Acronym	PURSUIT
Contract Number	FP7-INFSO-ICT 257217
Start date of the project	1.9.2010
Duration	30 months, until 28.2.2013
Document Title:	On Congestion Control in Information-Centric Networks
Date of preparation	2012-12-18
Author(s)	Somaya Arianfar (AALTO-HIIT), Pasi Sarolahti, and Jörg Ott
Responsible of the deliverable	Somaya Arianfar Email: somaya.arianfar@aalto.fi
Target Dissemination Level:	PU
Status of the Document:	Completed
Version	1.0
Document location	http://www.fp7-pursuit.eu/
Project web site	http://www.fp7-pursuit.eu/

SEVENTH FRAMEWORK PROGRAMME

European Commission
Information Society and Media

On Congestion Control in Information-Centric Networks

Somaya Arianfar, Pasi Sarolahti, and Jörg Ott

December 2012

Contents

Many different Information-Centric Networking (ICN) proposals have been introduced recently; mainly targeting better information dissemination by the means of different naming and routing modules. Many open questions regarding various other aspects of ICN designs, however, remain unanswered. For instance, the need for having a specific resource management module and its defining principles has yet to be clarified in most ICN designs. In this work we take the early steps to discuss the requirements and options for having congestion control functionality in ICN networks. We first introduce a set of principles combining the design goals of many ICN proposals and their congestion control requirements. We then discuss different possible options for defining a congestion control module based on the expressed principles.

Chapter 1

Introduction

Information-Centric Networking (ICN) has been introduced as a potential future networking architecture with a primary design principle of location-independent dissemination of information. There are a number of ICN proposals that try to achieve this goal with variable design choices for naming and API, routing, forwarding, and other functionalities. For example, the CCN architecture [1] uses hierarchical names and a dynamic routing and forwarding model capable of adapting to real-time situations such as link/node load and failure. In CCN the functions related to name resolution, routing and forwarding are fairly strongly coupled, and implemented as part of the basic router functionality. PURSUIT [2] on the other hand separates the naming infrastructure from the forwarding and routing infrastructure.

One of the ultimate key goals for the location-independent information networking in ICN is to achieve more efficient resource management and resource availability for different networked parties. However, different ICN solutions have somewhat different core designs that has implications on thinking the resource management and congestion control in those systems. For example, CCN is fundamentally a store-and-forward system with name resolution happening at each hop. In CCN neighboring hosts exchange interest messages and multiple interests can be sent at a time, hence CCN can implement a form of window-based flow control between neighboring hosts. However, the report does not discuss what happens when the demand for data from the edges increases beyond the capacity that the network can serve and how would this affect the applications. More recently other flow-control algorithms for CCN have been developed [3] with the direct involvement of the application level receiving end.

The forwarding model advocated by the PURSUIT project separates the forwarding functionality from the caching functionality and name-resolution, and is more similar to multicast forwarding, for example based on Bloom filters [4]. In such model the effects of congestion can be expected to be more similar to the traditional (multicast) networks, and the need for an congestion control algorithm based on network feedback is more obvious. A receiver-based algorithm has been

proposed for PURSUIT architecture [5].

Despite the earlier individual attempts such as those discussed above, we think that a discussion on different possible options for congestion control in information-centric networks is still missing from most ICN proposals. Therefore, here we aim to clarify the required models for signaling the resource availability and avoiding congestion. We also aim to discuss various congestion control options that could be applied to different flavors of ICN design. We believe that although conceptually congestion control in an ICN and traditional end-to-end IP networks are similar, there are some fundamental differences between the two approaches that deserve some special attention, as will be discussed in this report.

Chapter 2

Background

In the following we discuss few basic fundamental properties of information-centric networks we are assuming when building and discussing the congestion control framework in the remainder of this report.

2.1 Pull-based logic

In ICN networks, information items are named. Information names and range are used to identify each part of an information item that is of interest to a receiving application or in other words a subscriber. The subscriber retrieves information through a interest/ response style of interaction with the network. The interaction usually follows a data-oriented pull-based mechanism, where the subscriber controls the information reception based on the information name and interest range. The range of the interests could possibly be defined through the adjustments of the congestion window size, e.g. [1, 5]. The receiver does not need to necessarily know the identity of the end-point(s) responding to its interests.

2.2 Resolution

While most ICN proposals do not use host addresses, they still need to be able to route and forward the packets to an ultimate destination that is able to serve the application requests. As described previously, application requests are represented to the the network in the form of information names. Therefore, some method is needed to resolve the name into forwarding actions applied at the ICN routers.

Two main types of resolution can be applied. In architectures such as the one proposed in the context of PSIRP project [4] the content name is resolved into a forwarding identifier in the beginning of transmission. We call this *out-of-band resolution*. The forwarding identifier implicitly determines the host(s) that are going to receive the subscriptions for content, and the nodes participating in the forwarding process are simply forwarding packets based on the identifier. Because the resolution is done in advance, the requesting end-host protocol implementation is

able to see the resolved identifier and interpret it as much as it can, e.g. get an idea about the possible host, or group of hosts that are going to participate in transmission. The resolved identifier could simply be an IP anycast address, or it could be a flat, topology-independent label not revealing anything about the location of the possible host(s) participating in transmission.

Another way to resolve the name into forwarding actions is to do it in-band with the data transmission. In such case an explicit resolution phase is not needed, but the network routers choose the next hop directly based on information name. CCN follows this approach, and we call it *in-band resolution*. With in-band resolution, the requesting end host does not have any means to determine the resolved identifiers beforehand. This is a significant difference when designing congestion control, as will be discussed later in this report.

2.3 Packet caching

In ICN there is a possibility for caching the information items as they traverse the network. In particular, in addition to forwarding packets, forwarding nodes may also cache them. A network node may use such cached copies to satisfy requests, thereby becoming an unpredictable source for the subscriber [6]. While this helps efficiently repairing packet losses, caching is especially beneficial and could be counted on when serving the same content to many receivers (within a short period of time).

Chapter 3

Congestion Control Principles

The traditional Internet congestion control is based on end-to-end model between two communicating end hosts that treats the network as a non-transparent "black box", with all congestion control intelligence located at the end hosts, mainly at the data sender. In ICN environments, as described in previous section, some of the fundamental assumptions of the original congestion control environment are different. A communication session cannot be assumed to be a one-to-one relationship between two end hosts, but within the same session data can originate from multiple locations. Therefore the congestion control approach cannot be based on a feedback loop between one data sender and receiver.

This section discusses a general set of principles necessary for a congestion control module in information-centric networks. The principles are classified into two different categories: *1) Exclusion principles* that help drawing a borderline for elements *not* to be covered by a congestion control module in an ICN network, i.e., principles that define a separation of concerns of congestion control and other functions. *2) Inclusion principles* that define elements to be reflected in a congestion control module in the sense of positive requirements for such module.

3.1 Exclusion Principles

Consider the congestion control separately from the other protocol and stack logic

Congestion control is vital to the network and that is fundamentally independent of other protocol and network functionality such as reliability and global reachability. Therefore, the congestion control logic should not be intertwined with other protocol logic mainly concerning applications, e.g., in terms of what to send next, or which address to bind to. Nevertheless, this does not preclude providing visibility of other protocol state to congestion control if necessary, nor sometimes (re)using a protocol's own signaling and feedback mechanisms as an input method for congestion control.

Differentiate between the fairness state and the congestion avoidance state

Logically the outcome of a fairness mechanism should affect the outcome of congestion control mechanism, but we emphasize that they are not necessarily the same [7]. Therefore, we would like to differentiate between fairness and congestion control concepts and states.

3.2 Inclusion Principles

Maintain at least one control loop that involves the ultimate source/destination of the data but not necessarily both of them.

As mentioned earlier, some ICN proposals such as CCN claim to avoid the possibility of congestion collapse purely based on using in-network adaptive forwarding and indirect involvement of the source/receiver. However, they do not discuss how a subscribing application residing on the receiving side of a communication should be notified and respond to the resource (un)availability in the network. On the other hand, in most other ICN proposals without embedded resource management at the forwarding layer it is the sole responsibility of the receiver to impact the number of packets in the network and prevent the congestion collaps, e.g., ConTug [5]. Consequently, for both mentioned cases above, a congestion control module is required that supports control-loops involving either the ultimate source or receiver end-points.

Account for different paths

Being compatible with different styles of forwarding and resolution in ICN networks, we consider the case for different paths as part of the support required by a congestion control module; e.g., working congestion control module when middle points forward the packets on different routes for load balancing. Multiple paths could be fed to the end-points either before a communication starts or they could be created at the time of routing [1].

Account for different sources

In the context of defining a connection, source and receiver are the entities that have application implementation. In other word, in ICN a source could be interpreted as the application-level publisher, while a receiver is the subscriber. A resolved forwarding identifier can end in many different sources at the same time. Additionally, packet caches in the middle of a path can also count as sources. Considering this, the above case for supporting multipath can be extended to supporting the model where different chunks of data can arrive from different sources. Therefore, the packets might follow different paths towards different sources that may have different congestion control characteristics.

In summary the principles above state that a congestion control module for ICN should be flexible enough to allow both congestion collapse prevention in the network and minimal reporting to the applications. In next section, we introduce different options that cover some of our inclusion principles assuming the exclusion principles are already well understood and applicable. In other words, putting our thinking in the context of a layered protocol stack, we are going to discusse the

congestion control module at the same level as networking module and not on top of it.

Chapter 4

A congestion control framework

The current Internet congestion control is based on a couple of fundamental principles that have been effective in avoiding major congestion collapses in the internet. A sender implementation applies flow balance, i.e. controlling the number of outstanding packets in the network by triggering transmissions only when acknowledgments of successfully delivered packets arrive. In addition, the responsibility to reacting to congestion is taken at the end of transmission, while network has remained – at least in principle – fairly simple.

Adhering to the traditional principles is difficult in an information-centric network, where content and packets transmitting it may be replicated for multiple subscribers. Nevertheless, we aim to apply these principles in an information-centric fashion. First, we think that closed-loop feedback model should be kept, where the amount of traffic in the network is adjusted based on the feedback that a system receives. The rate of outgoing packets will be adjusted based on the congestion control rules being applied. Second, the end host that is the ultimate source/receiver for the information item is responsible in adapting its rate and reacting to congestion in response to the feedback from the network. Although, other components in the network path can still set up the incentives such that the end host operates according to agreed rules, for example by applying approaches similar to Re-Feedback [7]).

Inline with our arguments in previous sections, we first introduce the separate roles that different elements of a closed-loop feedback model have. We then go further on to explore different congestion control options for ICN, based on the elements of a closed-loop feedback model. However, before going into more details we have to clarify that in the context of this report we use the term **resource** quite often but that is independent from ICN specific terminology. Here, a resource always means a network resource relevant for the congestion control decisions, for instance an end-point, link, or buffer.

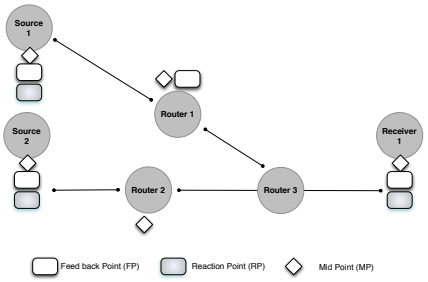

Figure 4.1: Path elements

4.1 The Control Elements

In the context of a closed-loop feedback model we define the *Reaction Point (RP)* as the entity responsible for keeping the required congestion control *state* that helps managing the rate. The measurement data that is used by the RP is collected along the path aside with data communication, and can take different forms depending on the congestion control mechanism: it could be packet losses, explicit congestion marks made by routers, or fine grained information about bandwidth. We call a node that reports the gathered congestion information back to the Reaction Point, a *Feedback Point (FP)*.

Additionally in a control loop there might exist different types of intermediaries that are significant for congestion control. This could be a router that applies multipath routing based on some load balancing rules [1], it is useful to identify these intermediaries, or *Midpoints (MP)*. Although, MPs might not play a logical role in the congestion control algorithm, knowledge of their very existence could become vital for the congestion control operations; e.g. when MPs duplicate the packets.

A feedback loop comprises a Reaction Point and a Feedback Point at its ends, and spans a sequence of Midpoints in the middle if they exist. A valid connection path interpreted through a resolved identifier may also only have an RP and FP, without any additional Midpoints.

Figure 4.1 shows these basic control-loop elements in an example scenario. It is important to note that RP, FP and MP constitute logical roles defined in a control loop and are not necessarily located in different nodes. For example, as can be seen in figure 4.1, the end-point that has the role of a RP or a FP can act as a Midpoint as well.

4.2 Discussing Congestion Control Options

Typically, congestion control algorithm aims to estimate the network conditions on the communication path and adjust the sending rate accordingly. One challenge in estimating conditions on the communication path in ICN is that the path may not be stable, and the source of information may change during the session. Therefore, we are going to discuss different options for designing congestion control

in information-centric networks based on the control loop elements that we have defined above.

1. *RP at the source, separate state for each destination.* The first congestion control approach that comes to mind is a model similar to TCP, where a source controls the sending rate of data. This approach requires the source to keep a congestion control state for every receiver. Consequently, this requirement results to a need for clients or paths towards the client to be identifiable to the RP in the source. In other words, RP either needs to be aware of the resolved identifier that implicitly or explicitly identifies the client-host(s) beforehand or this information has to be communicated to it afterwards.

More importantly what makes this model difficult to use in ICN networks is the fact that in these networks clients are not always bindable to specific sources for the whole duration of an information-retrieval operation. Thus, the source should keep a congestion control state for every possible client that has been served sometime in the past. This problem becomes more serious when packet caching nodes in the network also start to act as sources. In this case, any caching node that might act as a probable source of a probable packet would need to keep some congestion control states for that probable client. Resulting to an enormous number of states inside the network, this does not seem to be the most efficient way to implement a congestion control module.

2. *RP in the receiver, FP in the source*

Based on argument above it is sensible to put the RP somewhere else in the network rather than in the source, while FP remains in the source. Putting the RP in the receiver is one potential approach that matches the pull-based nature of the ICN proposals. However, the requirements on this model may change based on the information availability regarding the FPs and the MPs. Here are the possible cases:

(a)*Shared congestion control states for all possible FPs.* In case there are no specific information regarding the possible FPs and MPs that could be fed to the RP, the RP has to react to the resource availability based on implicit hints. The receiver has to able to make all the congestion control related decisions solely based on the information items it receive in response to its requests. Independent of the number and characteristics of different sources, paths, and MPs, one congestion control state could be used for every information item that is being retrieved from the network. This model is useful for ICN proposals with both in-band and out-of-band resolution, pull-based methods, and support for packet caching.

Being similar to the TCP's idea of a "black box" network, unless some additional intelligence is embedded to the RP, this model is only capable of adapting to the worst case scenario of network resource availability. Therefore, it is not clear how such model would benefit from the improved availability

promised by ICN networks. Examples of applying this model to ICNs could be found in [3, 5].

(b)*Separate congestion control states for all possible FPs (and MPs).* Instead of making complicated intelligent designs for the receiver-based RP to adopt to the implicit resource availability hints, there should exist simpler ways of improving the congestion control related reactions. Adding implicit or explicit FP/MP visibility for the RP could result to a more practical congestion control model that enables separate congestion control states for different FPs. This model does not seem to have been yet explored in ICN networks. based on this model, the receiver could adapt its request rate based on the resource availability regarding a specific FP. However, this requires an out-of-band resolution mechanism with hints about the FPs and MPs, or this information should be given to the RP afterwards. The RP can reflects the received information in adjusting the size of its interest window. Every source that sees the window of requests will allocate a send buffer based on the availability of the data in its local memory and the interest range. A major problem, however, is that it is difficult to see how the RP knows beforehand which source is actually going to respond to its window of requests, so that it can adjust the window size based on the congestion information regarding that source. There are two possible scenarios to resolve this issue:

i.*Bind information retrieval to specific sources* If the receiver already knows which source would respond to its specific request, it can adjust its requests rate based on that relevant congestion control information for that specific source. In other words, if the resolved identifier already enforces explicit binding of information to a specific source then this model would work. However, binding information retrieval to a specific source would make the whole design of an ICN proposal very similar to the TCP/IP model and would reduce the chances of improved availability through the usage of random packet caches or multiple co-existing sources.

ii.*Dynamic set of sources* Not binding the information retrieval to a source improves the chances of receiver to get the information from the best possible sources at any moment. However, it makes it difficult to decide which congestion control state is relevant for every outgoing request. Knowing the simple form of a resolved forwarding identifier in the form supported by most ICN proposals does not help here, because the set of the host that could reply back to an interest might vary within the set of the hosts supported by the resolved forwarding identifier. Details of a possible solution for this problem are discussed in next section.

3. *Many RPs and FPs throughout the network* Another option for locating the RPs and FPS is spreading them throughout the network. In this way, similar to the hop-by-hop model of congestion control, all the nodes involved in a routing process can act either as a FP or RP and a similar mode of back pressure would be appli-

cable. This is similar to the original model suggested by CCN [1]. The strength of this model comes from the information that it can provide to the RP about the FP. Even with in-band resolution mechanisms the RP is always located in the node that does the actual resolution, therefore the RP always knows the resolved identifier, which ends at a FP. Thus, at first glance it seems RP can react properly to the FP-specific resource availability without putting any extra requirements on ICN. However, adapting to the back-pressure mode, in practice every RP in this model needs to keep specific states about all other congested FP/RP throughout the end-to-end path and the subscriber who are actually using that FP/RP. Therefore, additional information regarding resources and publisher/subscriber identification seem necessary for proper reactions in each RP. Besides, depending on the implementation details the number of states in the network could easily become unmanageable.

Different options above show the possibility of having a wide range of congestion control models for ICN networks, compromising one thing or another. In our thinking, the most reasonable compromise is allowing the RP know about the network resources; the option that does not seem to be explored in the context of ICN networks. Our reasoning for preferring this compromise is the fact that topology independent resource identifiers are already feasible and in use in many ICN proposals including PURSUIT [2] and CCN's custodian-based information sharing [8]. Therefore, network resource identification on its own does not add any significant assumption to most of the existing ICN proposals. Communicating this information to the RP is the only remaining issue that depending on the resolution model might not be that difficult. In the next section we elaborate on communicating resource identifiers to the RP, the identification granularity, and RP's possible reaction in the context of option 2(b)ii above.

Chapter 5

ICN congestion control in context

Most of the feasible congestion control options described in previous section require network resource identification and communicating them to the RP. This is a reasonable requirement. To clarify, ICN is about making applications independent from knowing the location of information items and therefore independent from requiring global reachability and permanent application bindings. Resource management in ICN networks becomes unnecessarily difficult when such abstraction with no requirement for global addressing is confused with the existence of no network resource identifier whatsoever in the network. However, many ICN proposals already us different forms of network resource identifiers, and the fact that this identifier with fine enough granularity could be communicated to the RP is not a significant change to the way applications are supposed work in an ICN environment. In this section, we explain a specific form of identifying and communication network resources for the RP in the context of option 2(b)ii in previous section. We also explain a possible model of reaction to this information at the RP.

5.1 Resource identification and communication : pathlets

Assuming an out-of-band resolution mechanism is in place resulting to a resolved forwarding identifier without binding to any specific source, the type and granularity of the information that is passed to the RP afterwards is of most importance. Considering the suggested congestion control model the subscriber plays the role of a RP and controls the rate on top of a resolved networking identifier that it is aware of it. The RP already knows the resolved forwarding identifier but the more it knows about the different parts of the feedback loop that it would act on, the better it can decide on how to react. To have this additional information with finer granularity, we conceptually divide the resolved forwarding identifier between a subscriber and a source into segments, we call pathlets, as inspired by earlier work by Godfrey et al. [9]. The concept of pathlet gives us a way to model multiple sources and participating in a communication session, and allows identifying the shared portions of communication path as well as the bottlenecks.

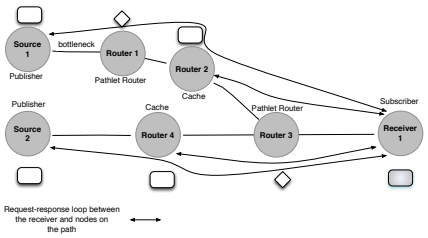

Figure 5.1: Information-centric networking scenario

We call a fragment of path separated by consecutive routers, or the source or the subscriber, a pathlet. A valid node in this case becomes a *pathlet router* and needs to identify itself to the subscriber as a pathlet representative. The pathlet identifier should be included along with each piece of data (feedback) that traverses that pathlet. The pathlet identifier does not need to have any topological context, such as IP address, but it can be any identifier that is statistically unique within its local domain. For example, a random number selected at deployment time would suffice. According to our terminology in previous section a pathlet router could be an example of a Midpoint (MP).

By monitoring the pathlet Identifiers, the subscriber is able to detect what pathlets are involved in communication. If a data (feedback) packets also indicate the pathlet on which congestion occurred or the bottleneck pathlet (for example by a mechanism similar to ECN), the subscriber can make informed decisions about how to adjust the rate at which data is being requested. The subscriber then can keep different state for each identified pathlet involved in communication.It is possible that a communication session involves several sources and alternative feedback loops that can be differentiated through sequences of pathlet identifiers. Therefore, pathlets could be used as one mechanism to identify FPs or at least the most important network resources used by each FP. For instance the pathlet routers could be inserted at 2 ends of the links with least capacity in the network and then used by the subscriber to identify which sources use that link. Clearly an FP itself could be one end of a pathlet that ends in a router or in the subscriber. The pathlet information is fed back to the RP in addition to the resolved identifier known to it. In other words, while an out-of-band resolved identifier gives implicit information about possible host(s) and explicit information about the forwarding path, the pathlet information that is fed back to the RP gives explicit information about the active part of the path and active sources that use that path and congest it.

5.2 Reaction

As mentioned earlier, the subscriber maintains the corresponding congestion control states per source or in this specific case per pathlet. Feedback messages are piggy-backed on interest and response packets, thus implicitly capturing congestion state of the pathlets that an interest-response pair traverses. The bottlenecked pathlets are identified through explicit marking of the congestion, e.g. ECN, for

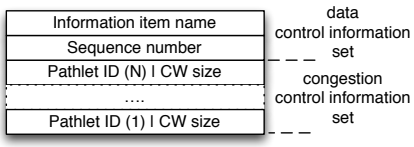

Figure 5.2: Interest packets control information set

that specific pathlet. Figure 5.1 shows the feedback loops and the congestion control point located at the data receiver (subscriber). For simplicity we assume that every MP in this context acts as a pathlet router as well. Since any node in the path can respond to an interest, not all the packets and feedback in the system will traverse the whole path: an interest may be satisfied from routers 2 or 4 if they have stored the requested packet. Then, the routers also becomes an FP for one of the control loops ended at the subscriber. All feedback messages carry identifiers showing the pathlets they have traversed.

The data (feedback) packets in this model would include pathlet identifiers of the path they have traversed. In Figure 5.1 for example, the data packet that is originated from source 2, would contain information regarding traversing 2 pathlets one between source 2 and router 3 and another one between router 3 and receiver 1. While a data packet originating from the router 2 would only include the pathlet identifier between router 3 and receiver 1.

The information regarding the overall pathlet identifiers and the bottlenecked pathlets help the subscriber to decide which feedback is relevant for its congestion control decisions. The subscriber can then change the per pathlet congestion control's window size based on the congestion control algorithm it is using. The congestion control information, however, needs to be interacted to the sources if it is not clear which source would respond to a request.

One simple mechanism to interact the congestion control window size with all possible sources is as following: the subscriber keeps the congestion information per active pathlets that it sees used during the retrieval of an information item. The subscriber can then piggy-back the window size for all the expected pathlets to the interest packet that goes through the network. Therefore, a interest packet would look like Figure. 5.2. In addition to the name of the information item, and a sequence number to identify the beginning of the interested byte range for that item, the subscriber also includes the pathlet specific congestion window in each request.

Every, source (publisher/cache) checks its memory/disk and if there is no matching data it it would forward the request upstream. Once the request packet reaches a pathlet router that needs to forward it upstream the pathlet router then pops off its own congestion and identification information from the request control information set and forwards the request upstream. If there is a source upstream that is able to serve the request, it can use the rest of the congestion window information located in the header to adjust its rate. Note that in this way, if possible pathlets congestion information could be shared among all the sources that use it.

Chapter 6

Conclusions

In this report, we have discussed different congestion control models that could be implemented and used in ICN networks. We have defined a set of congestion control principles to capture the required logic for an explicit resource management functionality in ICN networks. Our main observation has been for congestion control it would be useful and practical to avoid treating the ICN network as a black box.

In our suggested example of a preferable model of congestion control, we chose to be explicit about the availability of network resources to the reaction points in congestion control loop. The main consequence of our model is that (some) network resources need to be identifiable and the identity of active resources should be fed back to the reaction point. We claim that is a reasonable requirement to put on congestion control modules designed for ICN networks, as long as it is not used for topological binding, or enforce global address uniqueness and reachability as in IP. This comes within the limits of many ICN proposals without binding the subscribers to a specific source for the whole duration of an information-retrieval operation. In-depth investigation of this model, however, remains as part of our future work.

Finally, it is worth mentioning that this work only touches the surface of issues regarding the congestion control problem for ICN networks. More work is required to address problems that can raise regarding various aspects of congestion control modules and algorithms in an information-centric network.

Bibliography

[1] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking Named Content," in *Proc. ACM CoNEXT*, 2009.

[2] D. Trossen (ed.), "Architecture Definition, Components Descriptions and Requirements," PURSUIT Project, Deliverable D2.3, 2011.

[3] G. Carofiglio, M. Gallo, and L. Muscariello, "Joint hop-by-hop and receiver-driven interest control protocol for content-centric networks," in *Proceedings of the second edition of the ICN workshop on Information-centric networking*, ser. ICN '12, 2012, pp. 37–42.

[4] P. Jokela, A. Zahemszky, C. Esteve Rothenberg, S. Arianfar, and P. Nikander, "LIPSIN: Line Speed Publish/Subscribe Inter-Networking," in *Proc. ACM SIGCOMM*, 2009.

[5] S. Arianfar, J. Ott, L. Eggert, P. Nikander, and W. Wong, "A transport protocol for content-centric networks," 2010.

[6] S. Arianfar, P. Nikander, and J. Ott, "On content-centric router design and implications," ser. ReARCH, 2010.

[7] R. Briscoe, "Re-feedback: Freedom with Accountability for Causing Congestion in a Connectionless Internetwork," Ph.D. dissertation, UCL, 2009. [Online]. Available: URL: http://www.cs.ucl.ac.uk/staff/B.Briscoe/pubs.html#refb-dis

[8] V. Jacobson, R. L. Braynard, T. Diebert, P. Mahadevan, M. Mosko, N. H. Briggs, S. Barber, M. F. Plass, I. Solis, E. Uzun, B.-J. Lee, M.-W. Jang, D. Byun, D. K. Smetters, and J. D. Thornton, "Custodian-based information sharing," *Communications Magazine, IEEE*, vol. 50, no. 7, pp. 38–43, Jul. 2012.

[9] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," in *Proc. of ACM SIGCOMM*, 2009.