# Bulk of Interest: Performance Measurement of Content-Centric Routing

Matthias Wählisch
Freie Universität Berlin
waehlisch@ieee.org

Thomas C. Schmidt
HAW Hamburg, Informatik
t.schmidt@ieee.org

Markus Vahlenkamp
HAW Hamburg, Informatik
markus@vahlenkamp.net

## ABSTRACT

The paradigm of information-centric networking subsumes recent approaches to integrate content replication services into a future Internet layer. Current concepts foster either a dynamic mapping that directs content requests to a nearby copy, or an immediate routing on content identifiers. In this paper, we evaluate in practical experiments the performance of content routing, which we analyze with a focus on conceptual aspects. Our findings indicate that the performance of the content distribution system is threatened by a heavy management of states that arise from the strong coupling of the control to the data plane in the underlying routing infrastructure.

## Categories and Subject Descriptors

C.2.1 [**Computer-Comm. Networks**]: Network Architecture and Design

## General Terms

Measurement, Security

## Keywords

Performance, Experimental Evaluation, Routing

## 1. INTRODUCTION

One major dedication of today's Internet is the global distribution of content in huge amounts. Content distribution networks (CDNs) facilitate an efficient, wide-area replication of static data for selected content providers, whereas the end-to-end design of TCP/IP does not foresee implicit replication and in-network storage. There is no openly available standard solution for the asynchronous, global replication of popular content in the current Internet. Currently, *Information-Centric Networks* (ICN) [1] propose to fill this gap.

Essentially two approaches to routing exist in current ICN proposals, an evolutionary path that routes on IP, and 'clean slate' concepts that route on content identities. NetInf extends the current Internet by a resolution service that maps content names to topological IDs like IP addresses. TRIAD, DONA and NDN perform content retrieval by routing on names. Route responses and the data itself are then forwarded along reverse paths (RPF), either by using IP as a lower layer, or without IP but by dedicated RPF states.

Operating on the content itself forces network infrastructure into a content awareness. A mapping service is not only required to resolve names to source locations, but must advising a nearby replica, the existence of which it learned from the data distribution system. Content routers need to rely on names in its interface tables and – for RPF-based forwarding schemes – a reverse state for every data unit. This control information is highly dynamic and requires regular updates from the data plane. The ICN paradigm thereby opens up the control plane to continuous modifications from the data plane. This is in contrast to the current Internet, where DNS and routing states remain unaltered by data-driven events such as transmitted file names or data locations.

In this poster, we briefly present our study [2] of joint control and data plane behaviour under varying data conditions for the example of NDN [3] (§ 2). We are in particular interested in the response to load of the ICN infrastructure. Experiments are performed in test networks running PARC's CCNx software and reveal a flaw of performance (by concept, not implementation) with increasing demands to the routing system (§ 3).

## 2. BASIC MEASUREMENT SETUP

For our measurement study, we use the CCNx implementation version 0.5.1 [4], i.e., the client library to announce content interests, the content repository to store data, and the `ccnd` to forward subscription and data. The following analysis focuses on the effects on the router side. Even though the measurements relate mainly to the `ccnd`, we do *not* evaluate the implementation but use it as one real-world instance of the information-centric network deployment to illustrate protocol mechanisms. To gain a fine-grained view, we concentrate on the local system, as well as inter-router dependencies.

The basic network topology is represented by a chain of two CCNx routers directly interlinked at 100 Mbit/s, one end connects the content consumer, the other the content repository. We leave default values for all CCNx parameters. CCNx routers communicate via TCP.

## 3. RESULTS

To analyze the performance of content consumption in our experiment, the content receiver initiates parallel download of multiple 10 Mbit files over a constant time. We consider


*SIGCOMM'12,* August 13–17, 2012, Helsinki, Finland.
ACM 978-1-4503-1419-0/12/08.

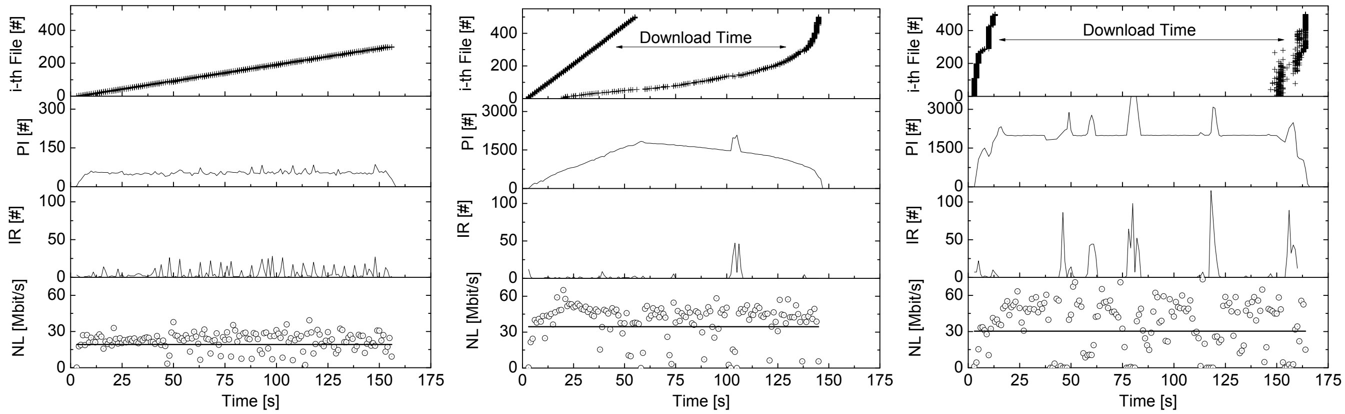


(a) 2 files per second (b) 10 files per second (c) 100 files per second

**Figure 1: Parallel download of 10 Mbit files: Start and stop time of the download per file at the receiver & resource consumption at its designated router [Pending Interests (PI), Interest Retransmits (IR), and Network Load (NL) including the mean goodput]**

three extremes, the request of 2 files, 10 files, and 100 files per second. Fig. 1 shows the start and completion time of the download per file (top graph), the Pending Interest Table (PIT) size, the effective number of Interest retransmissions, and the traffic load including the mean goodput at the first hop. For visibility reasons, we rescaled one y-axis (PI) in Fig. 1(a).

With an increasing number of parallel downloads, not only the download time increases significantly, but also the interval of the request and receive phase grows in the scenarios of overload. While the download time is almost constant for two files per second (cf., Fig. 1(a)), the stop time diverges non-linearly from the beginning of the download in the cases of excessive parallelism (cf., Fig. 1(b),(c)). 150 *s* are needed to download *each* single file in the worst case (Fig. 1(c)), while the link capacity would permit to retrieve *all* files in about 10 *s*.

The reason for this performance flaw is visualized in the subjacent graphs. A higher download frequency leads to an increasing number of simultaneous PIT entries, which require coordination with the data plane. Each file request will be split into the request of multiple chunks, in which the generation of corresponding interest messages will be pipelined. As soon as the content traverses, Interest states dissolve and thus release memory. These operations cause a continuous maintenance of states triggered by data traversal, and a simultaneous burst in CPU load (CPU exhaustion not shown). Finally it results in growing Interest retransmits after droppings or timeouts (shown in second lowest graphs), which again leads to retransmissions of data chunks. As an overall net effect, the network utilization fluctuates significantly, but does not adapt to actual user demands: Even though data requests could fill the links easily, the average load remains about constant at 30 %.

# 4. CONCLUSIONS & DISCUSSIONS

We have analyzed content-centric routing performance under varying loads and closely examined the implications of data-driven state management. The exhaustion of memory and processing resources following excessive state allocations was identified as one major reason for service degradation. It is worth noting that these effects are not shortcomings of the CCNx implementation, but driven by concept. Resource-intensive state management opens the floor widely to DoS attacks. Threats arrive from *resource exhaustion*, *state decorrelation*, and *path and name infiltration* [2].

An obvious approach to mitigate the resource exhaustion problem is to limit the rates of state injection into the network. Applying restrictions per user, though, will require addressing and tracking of end nodes, and lead to traffic shaping and bandwidth restrictions. As content states will accumulate in the network, and inter-provider deployment almost surely will lead to a heterogeneous, unbalanced network transitions, rate limiting may milden, but cannot effectively prevent the resource exhaustion problems discussed in this paper. Conceptual seem necessary to make content centric networking scalable, robust, and resistant to infrastructure attacks.


# Acknowledgment

This work has been supported by the German Federal Ministry of Education and Research (BMBF) within the project HAMcast (see http://hamcast.realmv6.org).


# 5. REFERENCES


[1] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlmann, “A Survey of Information-Centric Networking (Draft),” Dagstuhl Seminar Proceedings, Tech. Rep. 10492, 2011.

[2] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, “Backscatter from the Data Plane — Threats to Stability and Security in Information-Centric Networking,” Open Archive: arXiv.org, Technical Report arXiv:1205.4778v1, 2012. [Online]. Available: http://arxiv.org/abs/1205.4778v1

[3] V. Jacobson, D. K. Smetters, J. D. Thornton, and M. F. Plass, “Networking Named Content,” in *Proc. of the 5th Int. Conf. on emerging Networking EXperiments and Technologies (ACM CoNEXT’09)*. New York, NY, USA: ACM, Dec. 2009, pp. 1–12.

[4] PARC, “The CCNx Homepage,” www.ccnx.org, 2012.