ARTICLE IN PRESS

Computer Networks xxx (2013) xxx–xxx

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

On the performance of bandwidth and storage sharing in information-centric networks

Giovanna Carofiglio [a], Massimo Gallo [a], Luca Muscariello [b,*]

[a] Bell Labs, Alcatel-Lucent, Route de Villejust 91620, Nozay, France
[b] Orange Labs, France Telecom, 38-40, rue du Gen. Leclerc 92794, Issy-Les-Moulineaux, France

ARTICLE INFO

Article history:
Received 30 November 2012
Received in revised form 6 June 2013
Accepted 21 August 2013
Available online xxxx

Keywords:
Information-Centric Networking
Performance analysis
Transport

ABSTRACT

Information Centric Network (ICN) proposals, have recently emerged to define new network architectures where content, and not its location, becomes the core of the communication model. Such new paradigms push data storage and delivery at network layer and are designed to cope with current Internet usage, mainly centered around content dissemination and retrieval.

In this paper, we develop an analytical model of ICN storage and bandwidth sharing under fairly general assumptions on total demand, topology, content popularity and limited network resources. Our study applies to a class of content oriented networks identified by receiver-driven packet-based communication with in-network storage. We derive a closed-form expression for expected stationary delivery time as a function of hit/miss probabilities at network caches, content popularity and cache sizes. Our analytical results, supported by packet level simulations, can be used to analyze fundamental trade-offs of ICN architectures. They also provide an essential building block for the design and evaluation of ICN protocols.

1. Introduction

Internet usage has significantly evolved in the last years, and today is mostly centered around content dissemination and retrieval. We assist to an exponential growth of digital information diffused over the Internet, eased by cheaper storage and bandwidth supports, and driven by the increasing popularity of highly demanding services, such as cloud computing or video delivery. However, Internet architecture is still based on the original end-to-end model and appears to be unsuited to deal with the aforementioned trends. A large range of over-the-top (OTT) solutions like Content Delivery Networks (CDNs), Peer-to-peer (P2P), have been designed and widely deployed to overcome the mismatch at application layer. Such solutions suffer from a number of limitations in terms of efficient resources utilization, scalability and management cost, mainly because they are engineered as overlay service-dedicated infrastructures at application layer.

Significant research projects have been funded in the last years focusing on the definition of novel architecture designs for the future Internet (e.g. US NSF GENI [1], or EU FIA [2]). Information Centric Network (ICN) proposals, as PARC's CCN [3], PSIRP [4], or DONA [5], aim at redesigning the Internet architecture with named data as the central element of the communication paradigm, instead of its physical location. As a shared principle of ICN proposals, content is uniquely identified, addressed and retrieved by its name independently from its location, and storage capabilities are embedded into the network.

To meet the challenges of ICN proposals, transport protocols need to natively support content dissemination and retrieval based on names instead of end-point addresses, and to exploit content-awareness for efficient usage of bandwidth and storage resources. A preliminary

* Corresponding author. Tel.: +33 1 45 29 60 37; fax: +33 1 45 29 60 69.
E-mail addresses: Giovanna.Carofiglio@alcatel-lucent.com (G. Carofiglio), Massimo.Gallo@alcatel-lucent.com (M. Gallo), luca.muscariello@orange.com (L. Muscariello).

1389-1286/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2013.08.018

Nomenclature

N	number of network nodes (levels in the tree topology)
M	number of content items ($m = M/K$ in each class k)
$x(i)$	content store (Cache) size at node i
λ, $\lambda(i)$	total content request rate at first node, at node $i > 1$
$\mu(i)$	miss rate at node $i > 1$
λ_k	content request rate for class k w/o filtering
λ_k^f	content request rate for class k with filtering
σ	average content item size in number of packets
$q(i)$	content popularity distribution of requests at node i
$p_k(i)$	miss probability for class k at node i
$p_k^f(i)$	miss probability with filtering for class k at node i
R_i	round trip delay between client and node i
γ_i	fair share on route i (Gbps)
VRTT_k	virtual round trip delay of class k w/o filtering
VRTT_k^f	virtual round trip delay of class k with filtering
$\mathrm{RVRTT}_k(i)$	residual VRTT_k at node i w/o filtering
$\mathrm{RVRTT}_k^f(i)$	residual VRTT_k at node i with filtering
Δ	filtering time window
T_k	content delivery time of class k
W	number of parallel packet requests expressed by clients

understanding of transport and caching issues in ICN appears to be necessary to quantify potential benefits and to guide optimized protocol design.

In this paper, we focus on the performance evaluation of ICN networks and develop an analytical model of bandwidth and storage sharing under limited resources. Each user is supposed to implement a receiver-driven flow control protocol yielding fair and efficient bandwidth utilization along the path to the content repository, while caches are managed by least-recently used (LRU) per-packet replacement policy. Under these assumptions, a closed-form characterization of the average content delivery time is provided, which captures the tradeoff between user performance and limited network resources. In this paper we extend the analysis presented in [6,7]. In particular, the manuscript includes a set of complete propositions and related proofs, not included in the previous publications.

The rest of the paper is organized as follows. Section 3 provides system description. Section 2 is devoted to related work, while Section 4 introduces notation and main modeling assumptions. In Section 5 we report the main analytical results about single cache dynamics, while in Section 6 we extend the model to network case. In Section 7 we present some simulation results in order to validate the storage sharing model. Section 8 presents the storage sharing model while Section 8 presents some results considering both bandwidth and storage sharing. Section 9 concludes the paper.

2. Related work

Previous work on information-centric networks has mainly focused on global architecture design (i.e. [3–5,8]) while less effort has been devoted to analyze caching and transport mechanisms in such architectures. Recently, Arianfar et al. [9] analyze the feasibility of caching in routers at line-speed, while Lee et al. [10] consider the benefits of ICN in-network storage in terms of energy efficiency with respect to traditional distribution architectures. Also, Carofiglio et al. show the role played by storage management in ICN by means of experimental evaluation [11]. However, none of the aforementioned works provide an analytical characterization of transport performance, and its interaction with packet-level caching dynamics. In the context of Web caching there have been previous attempts to modeling content-level cache dynamics, most of them related to a single cache scenario under LRU replacement policy. The majority of analytical models of LRU caches start from the relation between the LRU miss probability, and the tail of the search cost distribution for the Move-To-Front (MTF) searching algorithm (i.e. King et al. in [12]). In [13] Flajolet et al. derive an integral expression for the Laplace transform of the search cost distribution function, that needs to be numerically integrated with complexity proportional to the cache size and the number of content items. Alternative combinatorial approaches are developed in [14] by Starobinski and Tse and in [15] by Coffman and Jelenković. In [16], Jelenkovic et al. give an asymptotic characterization, for a large number of content items, of the MTF search cost distribution and hence of the LRU miss probabilities both in the light-tailed and in the heavy-tailed case. Recently, such approach has also been extended, providing asymptotic closed formulas for the FIFO and RND replacement policy in [17] by Simonian et al.

A recent work in [18] by Jelenkovic and Kang provides an analytical characterization of the miss probability and thus miss rate under Poisson assumptions of content requests' arrivals.

It is worth to remark that almost all of these prior studies are devoted to the analysis of LRU-based rules for a single cache and with unit-sized objects. In [19] Che et al. apply a single LRU cache approximated model to study a cache coordination technique, whereas recently, the approximated single cache model of [20] by Dan and Towsley has been applied in [21] by the same author to a network of caches, under the Independence Request Model (IRM) assumption. The approximated model introduced in [19] has also been used in [22] by Laoutaris et al. and more recently in [23] by Roberts et al.

3. System description and problem statement

In this paper we primarily focus on the content-centric networking (CCN) proposal by PARC [3], though the model-

Please cite this article in press as: G. Carofiglio et al., On the performance of bandwidth and storage sharing in information-centric networks, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.08.018