Computer Networks 110 (2016) 104–117

Contents lists available at ScienceDirect

Computer Networks

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

Optimal multipath congestion control and request forwarding in information-centric networks: Protocol design and experimentation

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

[a] Cisco Systems, 11 rue Camilles Desmoulins, 92130, Issy Les Moulineaux, France
[b] Bell Labs, Nokia, 7 Route de Villejust, 91620 Nozay, France

ARTICLE INFO

Article history:
Received 26 May 2015
Revised 24 January 2016
Accepted 18 September 2016
Available online 21 September 2016

Keywords:
Information-Centric networking
Multipath congestion control
Request forwarding

ABSTRACT

In this paper we consider the problem of joint congestion control and request forwarding in Information-Centric Networks, namely the named-data networking architecture (NDN). The network architecture we consider is based on information retrieval natively pull-based, driven by user requests, point-to-multipoint and intrinsically coupled with in-network caching. We formalize the problem as global optimization with non-linear objectives and linear constraints with the twofold objective of maximizing user throughput and minimizing overall network cost. We solve it via decomposition and derive a family of optimal congestion control strategies at the receiver and of distributed algorithms for dynamic request forwarding at network nodes. An experimental evaluation of our proposal is carried out in different network scenarios using realistic workloads, to assess the performance of our design and to highlight the benefits of an ICN approach. The experimentation is carried out using the NDN software router implementation on a large grid infrastructure deployed to enable experimental research.

1. Introduction

The TCP/IP core of the Internet architecture has remained stable over the last decades, while accommodating unexpected innovation at lower and upper layers. Recent advances in wireless and optical networking technologies have empowered heterogeneous mobile connectivity and boosted access network capacity. The diffusion of web services, cloud and social networks has in turn driven Internet usage towards information delivery and imposed a new communication model based on Information exchange, caching and real time processing.

Current information delivery solutions are deployed as overlays on top of the existing IP network infrastructure leveraging CDNs, transparent caching etc. However, the inefficiency of the overlay approach in terms of performance guarantees/ resource utilization has been pointed out in the research community, together with question of the sustainability of such evolution model [19,22,27,30]. Some evidence has been provided that "architectural anchors" like IP addressing, host-centric point-to-point communication, inherently curb the Information-driven network evolution.

A mismatch exists between the location addressing of IP and the location-agnostic content addressing by name realized by Information services (e.g. via HTTP). It calls for a paradigm shift of the Internet communication model as proposed by Information-Centric Networking (ICN) architectures [19,22,26,30,41]. The common objective is to embed content awareness into the network layer to enable efficient, mobile and connectionless information-oriented communication. The twist associated to ICN communication model lies on a small set of principles. ICN advocates a *name-based* communication, controlled by the receiver (*pull-based*), realized via name-based routing and forwarding of user requests in a *point-to-multipoint* fashion and supported by the presence of a pervasive *network-embedded caching* infrastructure.

The novel ICN's transport paradigm holds considerable promises for enhanced flexibility in terms of mobility management, improved end-user performance and network resources utilization. This is due to the coupling between (i) a 'connectionless' communication with unknown sender(s), (ii) a unique endpoint at the receiver and (iii) dynamic request routing, decided by network nodes aiming at localizing temporary copies of the content, in a hop-by-hop fashion.

The definition of multipath transport control mechanisms adapted to ICN is at an early stage in the literature [8,23] as well as that of ICN dynamic forwarding mechanisms [13,39]. A comprehensive analysis of the joint problem of optimal congestion control and request forwarding still lacks.

* Corresponding author.
E-mail addresses: gcarofig@cisco.com (G. Carofiglio), massimo.gallo@nokia.com (M. Gallo), lumuscar@cisco.com, luca.muscariello@gmail.com (L. Muscariello).

http://dx.doi.org/10.1016/j.comnet.2016.09.012

In this paper, we tackle the problem of a joint multipath congestion control and request forwarding for ICN. The key contributions of the paper are the following:

- We formulate a global optimization problem with the twofold objective of maximizing user throughput and minimizing overall network cost. We decompose it into two subproblems for congestion control and request forwarding and solve them separately.
- We derive a family of optimal multipath congestion control strategies to be performed at the receiver and prove the optimality of the proposal.
- We derive a set of optimal dynamic request forwarding strategies and design a distributed algorithm to be implemented by ICN nodes in conjunction with the proposed receiver-driven congestion control.
- We design a mechanism managing fine grained forwarding to high popular content while keeping aggregate based forwarding to less popular content. The mechanisms allows to achieve optimal performance while keeping the system scalable.
- Finally, we implement the proposed joint congestion control and request forwarding in NDN[41] and set up a testbed for large scale experimentation on the Grid'5000 facility ([1]) The performance are assessed through experiments under realistic traffic demand, synthesized by a large set of traffic measurements, and in different network scenarios to appreciate the convergence to the optimal equilibrium, the role of in-path caching and the benefits of ICN multipath. The protocols designed in this paper have been developed and released in open source for the NDN code base and made available at http://systemx.enst.fr/lurch.

In this paper we do not consider several research problems that still requires solutions namely:

- the set of routed faces for a given name prefix are precomputed by a routing algorithm. In this paper routing is part of the assumptions.
- The set of data names that the client issues into the network requires a protocol to manage that. For instance by mean of a content manifest or by using a discovery protocol. This is part of the assumptions made in this paper.

The remainder of the paper is organized as follows. Problem statement, design goals and choices are summarized in Section 2. The formulation of the global optimization problem is presented in Section 3, while in Section 4.1 we solve it by decomposing it into two separate sub-problems related to multipath congestion control and to request forwarding. From the theoretical solution, the optimal rate and congestion controller and the forwarding strategy are derived. Section 4.2 presents the details of the transport protocol and of the design of the distributed forwarding algorithm. The performance assessment by means of experiments is shown in Section 6 and Section 7. Finally, Section 8 surveys related work, while conclusions are drawn in Section 9.

2. Problem statement

ICN transport model fundamentally differs from the current TCP/IP model. To comment on the main differences, let us summarize the basic ICN communication principles.

2.1. System description

Our work is primarily based on CCN/NDN proposal [22,41], though the defined mechanisms have broader applicability in the context of ICN (http://tools.ietf.org/group/irtf/trac/wiki/icnrg) and more generally of content delivery networks employing a similar transport model (e.g. some HTTP-based CDNs). In ICN, *Information objects* are split into Data packets, uniquely identified by a name, and permanently stored in one (or more) repository(ies). Users express packet requests (*Interests*) to trigger Data packets delivery on the reverse path followed by the routed requests.

Interests are *routed by name* towards one or multiple repositories, following one or multiple paths. Rate and congestion control is performed at the end user. Intermediate nodes keep track of outstanding Interests in data structures called PIT (Pending Interest Table), to deliver the requested data back to the receiver on the reverse path. Each PIT entry has an associated timer, so that all requests for the same Data, during such time interval are not forwarded upstream as long as the first query is outstanding. In addition, nodes temporarily store Data packets in a local cache, called *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* via Longest Prefix Matching returns the interface where to forward the Interest (selected among the possible ones). FIB entries are associated to name prefixes (see [22]). As a consequence, Data may come from the repository, or from any intermediate cache along the path with a temporary copy of the Data packet.

2.2. Potential and challenges of ICN transport model

As a result of the addressing-by-name principle, ICN transport model overcomes the static binding between an object and a location identifier: the receiver issues name-based packet requests over possibly multiple network interfaces with *no a priori knowledge of the content source* (hitting cache or repository). The content-awareness provided by names to network nodes enables a different use of buffers, not only to absorb input/output rate unbalance but for temporary caching of in-transit Data packets. Even without additional storage capabilities in routers, the information access by name of ICN allows two new uses of in-network wire speed storage: (i) *Reuse*: subsequent requests for the same Data can be served locally with no need to fetch data from the original server/repository; (ii) *Repair*: packet losses can be recovered in the network, with no need for the sender to identify and retransmit the lost packet.

Under such assumptions, the notion of connection between two endpoints, as in the TCP/IP model, is no longer required, worse, connections would prevent (i) the early reply of the request at the first unknown hitting cache on the path or (ii) the dynamic path switching due to end-user mobility. If a first step toward a multiparty content-oriented communication has been made by Swift (http://libswift.org/), ICN directly introduces a *connectionless, yet stateful* transport model where just the receiver keeps a flow state associated to an ongoing content retrieval. As a consequence, ICN simplifies mobility/connectivity disruption management, not requiring any connection state migration in case of end-user mobility or path failure.

ICN transport model brings new challenges that TCP, even in its multipath versions, can not address and, instead, motivates the definition of novel transport control and forwarding mechanisms. Indeed, TCP is connection-based and multipath solutions (e.g. [37]) perform load-balancing over static paths, precomputed by a routing protocol. Instead, in ICN, not only the receiver, but also in-path nodes to the repository may perform dynamic packet-by-packet re-