Computer Networks 55 (2011) 936–947

Contents lists available at ScienceDirect

Computer Networks

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

MultiCache: An overlay architecture for information-centric networking

Konstantinos Katsaros *, George Xylomenos, George C. Polyzos

Mobile Multimedia Laboratory, Department of Informatics, Athens University of Economics and Business, 76 Patision Str., Athens 104 34, Greece

ARTICLE INFO

Article history:
Received 17 May 2010
Received in revised form 1 November 2010
Accepted 13 December 2010
Available online 19 December 2010

Keywords:
Information centric
Content centric
Publish-subscribe
Peer to peer
Multicast
Caching

ABSTRACT

It has become apparent for quite some time that the Internet has evolved from a network connecting pairs of end-hosts to a substrate for information dissemination. While this shift towards information centric networking has been clearly demonstrated by the proliferation of file sharing and content delivery applications, it has not been reflected in a corresponding shift in network architecture. To address this issue, we designed MultiCache, an information-centric architecture aiming at the efficient use of network resources. MultiCache is based on two primitives: multicast and caching. It exploits overlay multicast as a means for content delivery and takes advantage of multicast forwarding information to locate, in an anycast fashion, nearby caches that have been themselves fed via multicast. We evaluate MultiCache against a widespread file sharing application (BitTorrent) with respect to both network resource consumption and end-user experience.

1. Introduction

The Internet was originally designed as a communication substrate enabling the delivery of data between pairs of end-hosts. Unfortunately, this end-point centric model seems to no longer cater to current communications needs: while users focus on the desired information, the underlying network focuses on the end-to-end communication between end-hosts. Inevitably, a translation between the information domain and the networking domain must take place, typically consisting of the establishment of a delivery path between the data provider and the data consumer. This translation is usually inefficient, as it is based on end-point centric data delivery overlays that neglect network topology, data location and data popularity, ultimately over-consuming network resources. For example, in *Peer to Peer* (P2P) file sharing, it has been shown that a major part of the incurred traffic crosses *Internet Service Provider* (ISP) boundaries, even though the desired information could have been retrieved locally [1].

We believe that at the heart of this problem lies the lack of information awareness inside the network, that is, the fact that only the end-points are aware of *what* is being delivered. Due to this deficiency, attempts to more efficiently use network resources, such as via caching and multicast, are unlikely to succeed, as decisions are made at network end-points based on coarse grained information. We believe that only an information-centric network model can enable the efficient use of network resources, thus better reflecting user needs. In this context, by knowing *what* is being delivered, in addition to *where* it originates from or is destined to, the network becomes inherently capable of forming targeted and efficient delivery structures. In addition, in this model users can directly express their interest in pieces of information, rather than engaging in the aforementioned domain translation.

Towards this direction, we present *MultiCache*, an overlay network architecture that brings information into focus. MultiCache takes advantage of information-awareness to improve network utilization via resource sharing. To achieve this, network operators deploy and control proxy overlay routers that enable the joint provision of multicast and caching, targeting both synchronous and asynchronous requests. End-hosts interact with this infrastructure by simply providing flat, location independent identifiers for

* Corresponding author. Tel.: +30 210 8203 693; fax: +30 210 8203 686.
E-mail addresses: ntinos@aueb.gr (K. Katsaros), xgeorge@aueb.gr (G. Xylomenos), polyzos@aueb.gr (G.C. Polyzos).

1389-1286/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2010.12.012

the desired content, without engaging in the process of locating an end-host providing the data. Inside the network, the Scribe overlay multicast scheme [2] is employed to transport the content from its origin in a publish/subscribe fashion, thus serving synchronous requests (e.g., flash crowds) and feeding in-network caches. By taking advantage of the locality awareness of the Pastry routing substrate [3], anycast queries, based on the already established overlay multicast forwarding state, are later used to locate nearby caches that can serve asynchronous requests via unicast.

In this paper, we highlight the benefits of information centric networking, as realized in MultiCache. In particular, we demonstrate how the proposed architecture takes advantage of information-awareness in order to improve both the utilization of network resources and the end-user experience. To this end, we compare a MultiCache-based content distribution application with the popular BitTorrent application. Our simulation results demonstrate a significant reduction of traffic load, in conjunction with considerably lower download times. To the best of our knowledge, this is the first head to head comparison of an information centric architecture against the current Internet model in the context of content distribution, with respect to network resource utilization and end user experience.

The remainder of this paper is organized as follows: In Section 2 we provide a thorough description of MultiCache, including both protocol functionality and deployment issues. We then present our MultiCache-based content distribution application, comparing it with BitTorrent in Section 3. In Section 4, we analytically investigate the scaling properties of MultiCache with respect to the overlay multicast trees created and the forwarding state required by MultiCache. We contrast MultiCache with other approaches in Section 5, concluding and presenting our next steps in Section 6.

2. Proposed architecture

2.1. Overview

MultiCache aims to establish an information-centric model of communication that better reflects current Internet usage patterns, in order to facilitate the deployment of resource sharing mechanisms. The introduction of information-awareness into the network enables the network (in addition to the end points) to identify pieces of information. In consequence, request similarities can be detected, thus allowing request aggregation and, eventually, network resource sharing. This enables the deployment of resource sharing mechanisms inside the network, rather than at the end-points. At the same time, information-awareness simplifies the usage model, as it allows end-hosts to directly denote the desired piece of information to the network, rather than engaging in the process of locating an end-point providing it.

In this context, the proposed architecture is realized as an overlay network, based on the deployment of additional infrastructure inside access networks. Namely, ISPs deploy and control *Overlay Access Routers* (OARs) at their *Points-of-Presence* (PoPs). These OARs establish an overlay routing substrate for the forwarding of data based on information identifiers, using the Pastry routing scheme [3]. Pastry provides a location-aware key-based routing fabric (see Section 2.2.1) which allows the network location of any piece of information to be discovered, given a globally unique information identifier (ID) in the *Pastry* namespace. Based on this substrate, the proposed architecture adopts and extends the Scribe scheme [2] to enable the joint operation of overlay multicast and caching, subject to the temporal characteristics of the requests (see Sections 2.2.2, 2.3 and 2.4). Scribe establishes multicast trees to serve multiple synchronous requests, while feeding in-network caches also located at the OARs. The already established multicast forwarding state is reused to allow later requests to reach nearby caches, leading to a hybrid protocol where data is delivered either via multicast or from a caching location via unicast.

At the edges of the network, end-hosts access the overlay network through a *proxy* OAR, designated during network attachment. The usage model of the proposed architecture closely follows the publish/subscribe paradigm [4]. Data consumers (i.e., subscribers) send a subscription request message towards their proxy OAR declaring the ID of a desired piece of information. The proxy OAR is then responsible to fetch the requested item using multicast and/or caching. On the other hand, data providers (i.e., publishers) advertise their content to the network by submitting an advertisement message to their proxy OAR, which is then responsible to locate the corresponding *Scribe* tree, using *Pastry* routing (see Section 2.2). Despite not participating in any of the overlay protocols (i.e., Pastry, Scribe and MultiCache protocols), end-hosts interact with the network in a simplified manner that does not include a translation between the desired data and its location.

The proposed deployment of overlay functionality inside access networks is motivated by several factors. First, it is expected to significantly improve performance, as it results in multicast trees that avoid forwarding content over the, typically lower bandwidth, access uplinks [5]. Moreover, the overlay character of the architecture facilitates deployment, as it does not require the replacement of existing infrastructure, thus also allowing the unobstructed operation of established services and applications. By deploying MultiCache inside access networks, content is cached close to the clients, facilitating the discovery of caches in the clients' networking vicinity and therefore enabling the localization of traffic (see Sections 2.4.3 and 3). At the same time, the deployment on top of ISP owned, dedicated servers is expected to provide lower churn rates, mitigating the maintenance overhead of the routing substrate, in contrast with typical end host deployment scenarios [6]. Finally, as discussed in [7], placing caches close to the end points of the network avoids incentive incompatibilities regarding inter-domain relationships. A simple deployment example is given in Fig. 1, with OARs collocated with the corresponding access routers.

2.2. Background

We provide below an introduction to the *Distributed Hash Table* (DHT) based Pastry overlay routing substrate