Computer Networks 91 (2015) 577–592

Contents lists available at ScienceDirect

Computer Networks

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

A distributed in-network caching scheme for P2P-like content chunk delivery

Xu Zhang*, Ning Wang, Vassilios G. Vassilakis, Michael P. Howarth

Faculty of Engineering and Physical Sciences, University of Surrey, Guildford, Surrey GU27XH, UK

ARTICLE INFO

Article history:
Received 10 October 2014
Revised 30 March 2015
Accepted 27 August 2015
Available online 12 September 2015

Key Words:
Information-centric networking
In-network caching
Peer-to-peer (P2P) systems

ABSTRACT

In-network content caching has recently emerged in the context of information-centric networking (ICN), which allows content objects to be cached at the intermediate router side. In this paper, we specifically focus on in-network caching of peer-to-peer (P2P)-like content objects for improving both service and operation efficiencies. We propose a fully distributed in-network caching protocol for P2P-like content chunks, aiming to reduce P2P based content traffic load and also to achieve improved content distribution performances. Toward this end, the proposed holistic decision-making logic takes into account context information of both the underlying network and the P2P characteristics, such as chunk availability, popularity and peer distances. In addition, we also analyse the benefit of coordination between neighbouring content routers when making caching decisions in order to avoid duplicated P2P chunk caching nearby. An analytical modelling framework is developed to quantitatively evaluate the efficiency of the proposed in-network caching scheme. Extensive experiments are also conducted to validate the analytical results.

1. Introduction

Information- or more specifically content-centric networking (ICN/CCN) [1,4] has been introduced in recent years to effectively deliver named content objects in the network rather than relying on the traditional host-to-host communication model. A key feature of ICN/CCN is that content objects can be cached within the network for local access by future interested clients. On the other hand, given the limited capacity of network devices comparing to the large content catalogue, the management of content distribution among caches at the router side, or the *in-network caching* becomes crucial. Current in-network caching schemes mainly focus on generic client-server based paradigms, while how in-network caching operates for chunk-based peer-to-peer (P2P)-like content distribution has not yet been addressed in the literature. In particular, content is divided into small-sized chunks in P2P applications (e.g., BitTorrent [34]), so as to facilitate the content downloading experiences by accessing multi-source peers at end clients, in comparison to a single *always-on* server in the CDN scenario. In this context, selfish behaviours of peers may lead to unexpected churns, resulting in unstable availabilities of demanded content objects that may deteriorate users' perceived service quality. As such, more "unselfish" caching at the network side can be intuitively regarded as a useful approach to mitigate such adverse effects. It is worth mentioning that in-network caching for chunk-based P2P content poses distinct technical challenges, which has not yet been thoroughly addressed until now. Important factors for P2P chunk caching, such as the content chunk *availability* (due to peer join/leaving churns), need to be specifically considered. In addition, another technical issue is *where* to cache between ad hoc source peers and requesting peers in order to maximise the caching efficiency. Other conventional context information, such as *router cache*

* Corresponding author. Tel: +4407448391396.
E-mail addresses: xu.zhang@surrey.ac.uk (X. Zhang), n.wang@surrey.ac.uk (N. Wang), v.vassilakis@surrey.ac.uk (V.G. Vassilakis), m.howarth@surrey.ac.uk (M.P. Howarth).

http://dx.doi.org/10.1016/j.comnet.2015.08.035
1389-1286/

space and P2P content chunk *popularity*, also needs to be taken into account when making comprehensive P2P content caching/replacement decisions.

In this paper, we aim to holistically address the aforementioned research challenges related to in-network caching of chunk-based P2P content, and this is achieved through the design of a fully distributed in-network P2P caching protocol. The decision-making logic behind such caching operations takes into account a series of important factors concerned from both the underlying network as well as the P2P content itself, with the objective of assuring caching efficiency and content distribution quality. We start from a simple scenario with *uncoordinated* decision-making by individual content routers (CRs) for caching passing-through P2P content chunks. Such a scenario takes into account a variety of metrics for enabling *locally* optimised caching decision by each CR, including chunk popularity/availability and cache capacity, etc. Out of these parameters, chunk availability is specific for the P2P content distribution scenario. In order to evaluate the in-network caching policy, a dedicated mathematical model based on a single CR is developed to analyse such content management in individual CRs. In addition, we further enhanced the scheme with locally *coordinated* caching techniques in order to avoid unnecessarily duplicated caching between nearby CRs. This requires coordinated CRs to periodically exchange information on their locally cached P2P chunks. Bloom filters (BFs) are adopted for aggregating information on cached chunks in signalling between coordinated CRs in order to achieve scalability. A theoretical analysis is carried out in order to derive an upper bound gain compared to the uncoordinated scenario. We show that, contrary to (always-on) CDN scenarios, a popularity-based caching policy can lead to decreased caching performance for P2P content chunks, while the proposed caching scheme can achieve more desirable performance. Last but not least, our work also offers insights into provisioning strategies on CRs' caching resources for supporting caching of P2P content chunks in dynamic environments. We validate the analytical results via extensive simulations and prove the effectiveness of the proposed model.

The rest of this paper is organized as follows. Related works are discussed in Section 2. In Section 3, we propose our efficient in-network caching policy, and generic stochastic models are also developed in this section on caching operations, with both independent and coordinated caching concerns. Simulation results are given in Section 4 to prove the effectiveness of the proposed caching policy and also to validate the analytical results. And we conclude the paper in Section 5.

2. Related work

In the recent IETF proposals, authors in [2,3] and [31] propose two options to cache content chunks within the network, either *implicitly* or *explicitly*. With implicit caching, CRs passively cache *passing-through* content chunks, while the explicit policy enables routers to explicitly request for content chunks to be cached locally. Song et al. [2] explored in-network caching related to P2P systems. A random caching policy was adopted at each router, while the issue of what to cache is not addressed.

In recent research works in the context of ICN, authors in [4,5] promote a universal in-network caching strategy. However, high caching redundancy can be introduced as observed in [6], and thus the authors suggested a probabilistic caching algorithm for distribution of content caching along a content delivery path, based on caching capacities of CRs for a given path per unit time. A simpler policy is proposed in [15], in which caching decisions are taken uniformly random with a fixed probability. Similarly, a random autonomous caching scheme is adopted in [8], such that a simple load-sharing between routers can be provided. Psaras [32] develops a Markov model to evaluate the in-network caching performance, and they show the performance gain under universal caching compared to non-caching scenario. User-assisted in-network caching is examined in [39] (or uCache) to enable users to share downloaded contents, apart from potential caching at network nodes (e.g., routers). It focuses on caching capabilities across users, considering both centralised and distributed caching. The uCache is based on a universal caching among users or ICN caches, while in our work, we focus on caching policy exploration and show that intelligent caching with cooperation among CRs can achieve considerable network and content distribution efficiency, compared to content-blind universal caching.

In the context of coordinated in-network caching, there have been a number of relevant works in the literature [9–14, 16,28]. The authors in [12–14] mainly explored the forwarding and routing scheme, while Xie et al. [16] focused on traffic engineering's point of view to promote collaboration in a two-level hierarchical manner, with the adoption of a universal caching policy. Cho et al. [10] and Ming et al. [11] both promoted implicit coordination to cache popular chunks at routers along a delivery path. And Guo et al. [9] proposed a collaboration scheme that goes beyond the en-route routers, the caching decision of which is based on solely content popularity measured in a wide scope of the whole network. Authors in [28] mainly explored the optimality of coordination among CRs of CCN, by developing an optimisation model to deal with the trade-offs between network performance and the coordination cost.

The authors of [7,18] and [19] are the only identified research works related to P2P in-network caching. Authors in [7] suggested to deploy in-network caches inside access devices at the network edge, in order to localize overlay traffic inside a network to avoid costly inter-domain traffic, and to achieve the improved performance of end users. While Yamamoto [18] established a test platform to show the basic benefits of in-network caching in P2P systems in terms of traffic redundancy elimination, the content distribution at routers still followed a random way. Ye et al. [19] mainly explored the cache placement policy along the delivery paths without considering caching decisions.

Different from the above research works, we focus on the development of efficient in-network caching operations of P2P contents in this paper, from both independent and coordinated perspectives. Specifically, due to dynamic P2P features, content availability has become a critical issue for P2P systems [20], and this is specifically taken into account in our work, along with the considerations of the cache capability at individual routers and also the content popularity. Although authors in [21] explored an optimal content placement, they