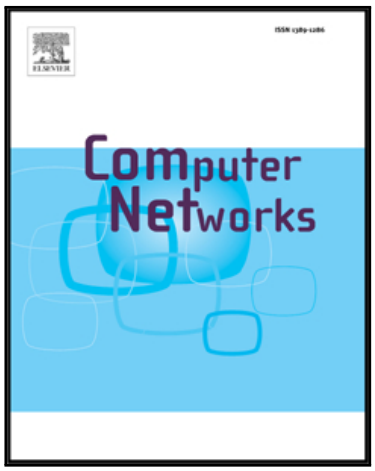

Greedy Caching: An Optimized Content Placement Strategy for Information-centric Networks

Bitan Banerjee, Adita Kulkarni, Anand Seetharam

PII: S1389-1286(18)30190-7
DOI: 10.1016/j.comnet.2018.05.001
Reference: COMPNW 6480

To appear in: *Computer Networks*

Received date: 7 November 2017
Revised date: 8 March 2018
Accepted date: 1 May 2018

Please cite this article as: Bitan Banerjee, Adita Kulkarni, Anand Seetharam, Greedy Caching: An Optimized Content Placement Strategy for Information-centric Networks, *Computer Networks* (2018), doi: 10.1016/j.comnet.2018.05.001

Greedy Caching: An Optimized Content Placement Strategy for Information-centric Networks

[1]Bitan Banerjee*, [2]Adita Kulkarni*, [2]Anand Seetharam
[1]Department of Electrical and Computer Engineering, University of Alberta, Canada
[2]Computer Science Department, SUNY Binghamton, USA
bitan@ualberta.ca, akulka17@binghamton.edu, aseethar@binghamton.edu

Abstract

Most content placement strategies in information-centric networks (ICN) primarily focus on pushing popular content to the network edge, fail to effectively utilize the caches in the network core and provide limited performance improvement. In this paper, we propose *Greedy Caching*, a content placement strategy that determines the set of content to be cached at each network node so as to maximize the network hit rate. *Greedy Caching* caches the most popular content at the network edge, recalculates the relative popularity of each piece of content based on the request miss stream from downstream caches and then determines the content to be cached in the network core. We perform exhaustive simulation in the Icarus simulator [1] using realistic Internet topologies (e.g., GARR, GEANT, WIDE, scale-free networks) as well as real-world request stream traces, and demonstrate that *Greedy Caching* provides significant improvement in content download delay (referred to as latency) over state-of-the-art dynamic caching and routing strategies for ICN for a wide range of simulation parameters. Our simulation results suggest an improvement of 5-28% in latency and 15-50% improvement in hit rate over state-of-the-art policies for synthetic traces.

1. Introduction

Content caching at storage-enabled network nodes is one of the most attractive features of information-centric networking (ICN), a new networking paradigm that aims to evolve today's Internet from a host-centric model to a content-centric one. In-network content caching enables requests for content to be served from intermediate caches, in addition to the content custodians (origin servers). Serving a request from an intermediate cache has several benefits such as reduced content download delay, increased throughput and decreased network congestion.

Video delivery companies (e.g., YouTube, Netflix) already use simple forms of popularity based in-network caching in today's Internet to improve user performance. These video delivery applications primarily determine the popularity of multimedia content based on parameters such as release date, viewership of past series of a show and push popular content to the network edge [2]. In recent years, proposed caching policies have also identified that caching popular content within the network is essential to improve performance [3, 4, 5]. However, existing policies focus mainly on the network edge and fail to effectively leverage caches in the network core. Deploying network-wide caches in ICN is likely to be expensive. Therefore, it is important to design efficient caching and routing policies that maximize cache utilization, both at the network edge and in the network core and minimize unnecessary content duplication. While it is tempting to think that determining what content to cache at a node only requires local information, content cached at downstream nodes drastically impacts the request stream seen by upstream caches and may ultimately reduce network-wide cache utilization. We use the words cache and node interchangeably in this paper.

To this end, in this paper, we propose *Greedy Caching*, a content placement strategy that determines the optimized set of content to be cached at each network node based on relative content popularity so as to maximize overall network hit rate. *Greedy Caching* estimates the relative content popularity at each node based on the request stream from directly connected users as well as the request miss stream from downstream nodes and then greedily caches the topmost relatively popular content at that node. The difficulty of the problem stems from the fact that different pairs of network nodes can forward requests to one another, resulting in interdependencies, and cycles in the underlying graph, thereby making it difficult to estimate the relative content popularity. We demonstrate via extensive experiments that *Greedy Caching* not only maximizes network level metrics such as hit rate, but also results in reduced content download delay (referred to as latency).

The main contributions of this paper are given below.

- We assume that the network has an underlying routing policy for forwarding requests for content toward the custodian. We propose *Greedy Caching*, a content placement policy that greedily caches the most popular content at each cache based on their relative content popularity. To estimate relative content popularity, *Greedy Caching* first leverages routes provided by the routing algorithm to cre-

*The first two authors have made equal contribution