

Coordinating In-Network Caching in Content-Centric Networks: Model and Analysis

Yanhua Li*

Haiyong Xie^{†‡}

Yonggang Wen[§]

Zhi-Li Zhang*

*University of Minnesota, Minneapolis, MN 55414, USA

[†]University of Science and Technology of China, Suzhou 215123, China

[‡]Central Research Institute, Huawei Technologies, Shenzhen 518129, China

[§]Nanyang Technological University, Singapore

Abstract—In-network content storage has become an inherent capability of routers in the content-centric networking architecture. This raises new challenges in utilizing and provisioning the in-network caching capability, namely, how to optimally provision individual routers’ storage to cache contents, so as to balance the trade-offs between the network performance and the provisioning cost. To address this problem, we first propose a holistic model to characterize the network performance of routing contents to clients and the network cost incurred by globally coordinating the in-network storage capability. We then derive the optimal strategy for provisioning the storage capability that optimizes the overall network performance and cost, and analyze the performance gains via numerical evaluations on real network topologies. Our results reveal interesting phenomena; for instance, different ranges of the Zipf exponent can lead to opposite optimal strategies, and the trade-offs between the network performance and the provisioning cost have great impacts on the stability of the optimal strategy. We also demonstrate that the optimal strategy can achieve significant gain on both the load reduction at origin servers and the improvement on the routing performance.

Keywords—in-network caching, content-centric networks, coordinated caching

I. INTRODUCTION

Internet has become a ubiquitous, large-scale content distribution system. To date, not only traditional Web contents, but also an increasingly large number of video contents have been delivered through the Internet (see, *e.g.*, [1], [2], [3]); moreover, video content delivery over the Internet is expected to grow even more tremendously in the next few years [4], [5]. These have posed significant challenges to the Internet, *e.g.*, how to store and disseminate the large-scale contents to support more robust, efficient, and expedited services for the users.

To address these challenges, Content Delivery Networks (CDNs) with built-in large-scale, distributed content caching mechanisms have been adopted in the Internet. CDNs are typically deployed and operated independently by third-party CDN carriers (*e.g.*, Akamai [6]), where CDNs are

interdomain overlays spanning across multiple underlying networks; each of such underlying networks may be operated by different Internet Service Providers (ISPs). Some other CDNs are deployed by individual ISPs within their own networks for intradomain content dissemination (*e.g.*, AT&T [7] and Level3 [8]). In both cases, content caching is one of the key mechanisms that make CDNs successful. However, content caching is only deployed as an overlay service rather than an inherent network capability, due to lack of the storage capability at individual routers.

Recently, in-network storage (and caching) as an inherent and integral network capability has been proposed in the emerging content-centric networking (CCN) architecture [9], [10]. CCN is a content-oriented future Internet architecture where content stores are available in routers; thus, content caching and dissemination become an inherent feature of routers. In such networks, users focus only on contents, rather than the physical locations from which contents can be retrieved. Moreover, the network routing and in-network storage are most likely provisioned by the same network carrier in a content-centric manner.

Content caching in CCN can be either coordinated or non-coordinated, similar to that in the Internet (see, *e.g.*, [11], [12], [13], [14], [15], [16]). On one hand, coordinated caching mechanisms require that CCN routers store contents in a coordinated manner, which allows more contents to be efficiently cached in the “cloud” of content stores closer to the users, thus improving the overall content delivery performance. Hence, the cost of coordinating content caching and the performance of routing traffic become main concerns in CCN. On the other hand, non-coordinated caching mechanisms store only the locally most popular contents at each CCN router, without coordination with other routers. Therefore, such mechanisms not only incur less coordination cost but also are more likely to store less distinct contents due to lack of coordination. Furthermore, studies have shown that the popularity of both Web and video contents follows the Zipf distribution [17], [18], [19], and that user-generated contents distributed through social networks are expected to become one of the most significant contributors to Internet traffic [5]; hence, a dominant portion of contents are *not*

Corresponding authors: Haiyong Xie (haiyong.xie@ustc.edu) and Zhi-Li Zhang (zhzhang@cs.umn.edu).

This study was done partly when Yanhua Li was a research intern at Center for Innovation, Huawei Central Research Institute.

popular. As a result, non-coordinated caching mechanisms are likely to suffer from the long-tail distribution, due to that contents are more likely fetched from distant origin servers that serve these contents.

Therefore, there exist clear trade-offs between the network performance and the coordination cost when designing in-network caching mechanisms for CCN. More specifically, coordinated caching mechanisms may trade the coordination cost for the network performance (*e.g.*, lower average latency), while non-coordinated caching mechanisms may incur a significantly lower cost on provisioning in-network caching and may degrade the network performance due to lack of fine-grained control on where contents are cached, retrieved and routed to users.

In this paper, we focus on in-network caching mechanisms and their trade-offs in content-centric networks, where routers possess both the routing and the in-network storage capabilities. We make the first attempt to address the new challenge in CCN, namely, how to optimally provision CCN routers' storage capability and investigate the trade-offs between the network performance and the coordination cost.

More specifically, we develop a simple holistic model, which allows us to systematically analyze the optimal strategy for provisioning the in-network storage capability. We provide rigorous proofs for the existence and uniqueness of the optimal strategy, which guides us to investigate the trade-offs between the network performance and the coordination cost. For the special case where the coordination cost is not a concern, we derive the closed-form solution of the optimal strategy and quantify the performance gains obtained when applying the optimal strategy.

We summarize our contributions as follows:

- We develop a simple holistic model to capture the network performance of routing contents to clients and the network cost incurred by globally coordinating the provision of the in-network storage capability.
- We derive the optimal strategy of provisioning the in-network storage capability to optimizes the overall network performance and cost, with mild conditions under which the optimal strategy is guaranteed to be unique.
- Through numerical analysis, we observe interesting phenomena. In particular, we observe that the stability of the optimal strategy is sensitive to key factors such as the the parameter of the content popularity distribution and the trade-off weights for the network performance and the coordination cost.

The rest of the paper is organized as follows. In Section II, we motivate our studies using a simple example. In Section III, we develop a holistic model that characterizes the overall network performance and cost to facilitate the analysis. In Section IV, we derive and analyze the optimal strategy for provisioning the in-network storage capability. In section V, we present the numerical evaluations and

results to analyze the performance gains of the optimal strategy. In Section VI, we present the related work. We conclude the paper with future work in Section VII.

II. MOTIVATION

We first motivate our study through an illustrative example shown in Figure 1, which shows an intradomain network consisting of three routers R_0 , R_1 and R_2 , and one origin server O serving two content objects a and b . The network belongs to a single administrative domain (represented by the cloud). All routers have the routing capability to forward contents to peer routers or clients. Moreover, R_1 and R_2 have storage capacity to store one single content object only, whereas R_0 does not have any available capacity for storing a or b .

We assume that there are two sets of clients (not shown in the figure) sending two request flows to their first-hop routers R_1 and R_2 , respectively. The two request flows are identical, represented by a repeating sequence $\{a, a, b\}$. We assume that the performance (*e.g.*, latency) of fetching contents from a peer router is much better than from the origin server O , since in real network, the origin may reside quite far away from the network. Then, apparently, storing contents at the routers R_1 and R_2 would reduce the overall delay and improve the network performance that clients experience. Due to the limited storage capacity, the problem is how to select contents to store at each router so as to improve the network performance and reduce the coordination cost.




Figure 1. A motivating example.

We consider the following two in-network caching strategies and their trade-offs: (1) *Non-coordinated caching*: R_1 and R_2 work independently, where they both adopt the canonical caching policy based on frequency or historical usage. Assume that the content popularity distribution is consistent, and that routers R_1 and R_2 have already cumulated the information that a is requested more often than b . In this case, both R_1 and R_2 store a . (2) *Coordinated caching*: R_1 and R_2 work jointly and always prefer each other over the origin server whenever possible. In this case, R_1 and R_2 may store a and b respectively. Without loss of generality, we assume that R_1 stores a , and R_2 stores b . Then, on cache misses, a requested content will always be retrieved from either R_1 or R_2 , rather than the origin server O .

We compare the coordinated and non-coordinated caching strategies when the network is in the steady state (*i.e.*, the in-network storage at R_1 and R_2 has been steadily populated), by using three metrics: *the load on origin*, *the routing hop count*, and *the storage coordination cost*. Note that the first two metrics can be used to measure the network performance, while the third metric can be used to measure the cost of provisioning the in-network storage capability. We summarize the comparison results in Table I.

Table I
COMPARING THE COORDINATED AND NON-COORDINATED STRATEGIES.

	Non-coordinated caching	Coordinated caching
Load on origin	33%	0%
Routing hop count	≈ 0.67	0.5
Coordination cost	0	1

First of all, the load on origin is measured by the percentage of all requests served directly by the origin server O . With the non-coordinated strategy, the requests for content a will be directly served by R_1 or R_2 (recall that both R_1 and R_2 store a in this case), while the requests for b will have to be served by the origin server. This means a total $1/3$ of all requests from two flows incur the traffic load on the origin server. However, with the coordinated strategy, since both a and b are stored locally (*i.e.*, at R_1 and R_2 respectively), all requests from the clients can be served by either R_1 or R_2 . Hence, the load on origin is 0 when using the coordinated strategy, much less than that when using the non-coordinated strategy.

Secondly, the routing hop count is measured by the average number of network hops traversed when fetching contents (we focus only on the links between R_0 , R_1 , R_2 and O). Using the non-coordinated strategy, clients requesting for a can directly fetch a from R_1 or R_2 without going through any peer router, while requests for b have to go to the origin which is two hops away via router R_0 (*i.e.*, the total hop count is 2). Therefore, the average routing hop count for non-coordinated strategy is $\frac{1}{3} \cdot 2 \approx 0.67$ per request. In contrast, using the coordinated strategy, only requests for b sent to R_1 and requests for a sent to R_2 trigger content fetching from their one-hop peer router, namely, R_2 and R_1 , respectively. Hence, the average routing hop count is $\frac{2}{6} \cdot 1 + \frac{1}{6} \cdot 1 = 0.5$ per request.

Moreover, the coordination cost is measured by the number of messages that have to be exchanged among routers in order to reach consensus on the caching decision. Apparently, in the non-coordinated strategy, routers decide which contents to store purely based on their local information; therefore, non-coordinated caching does not incur any coordination cost. However, to implement coordinated caching, non-trivial communication costs are necessary to coordinate the caching decisions of both R_1 and R_2 . In this example, to ensure that R_1 and R_2 store different contents, at least one message has to be exchanged between them.

In this example, the coordinated caching strategy leads to

a lower load on origin and a lower routing hop count, while the non-coordinated strategy incurs a lower coordination cost. This suggests that there exist trade-offs between the coordination cost and the network performance. Hence, it is important to investigate how to provision the in-network caching capability and understand the trade-offs.

III. NETWORK AND PERFORMANCE-COST MODEL

In this section, we develop a holistic approach to quantifying the overall network performance of routing traffic and the cost of coordinating the in-network storage capability.

A. A Simple, Holistic Network Model

We consider a simple, holistic network model for content-centric networks. We focus on the network of a single administrative domain (*e.g.*, an autonomous system), where a set of routers with both the routing and storage capability serve content requests originated from end users, as shown in Figure 2. The origin server O stores all content objects, referred to as the “*origin*”; therefore, requests for any content object can always be satisfied by O . Note that O is an abstraction of multiple origin servers (in practice, there are multiple origin servers hosting different contents).




Figure 2. A simple, holistic network model.

We assume that there are n routers in the network and the number of contents N is sufficiently large. To simplify the analysis, we also assume that contents are equally large and all routers have the same storage capacity c ; therefore, we are able to normalize the content size to one unit with respect to routers’ storage capacity. Note that in the recently proposed content-centric networking architecture [9], [10], contents are segmented into smaller pieces, each of which is treated as an individually named content object, to allow flexible distribution and flow control. Content segmentation has also been adopted in many existing overlay content distribution systems, *e.g.*, BitTorrent [20] and eMule [21]. These observations suggest that a homogeneous content model is reasonable in content-centric networks.

Many studies have shown that the content popularity follows the Zipf distribution (see, *e.g.*, [17], [18], [19]). The Zipf’s law predicts that out of a population of N elements,

the frequency of elements of rank i , denoted by $f(i; s, N)$, is

$$f(i; s, N) = \frac{1/i^s}{\sum_{j=1}^N (1/j^s)} = \frac{1/i^s}{H_{N,s}}, \quad i = 1, 2, \dots \quad (1)$$

where s is the Zipf exponent and $H_{N,s} = \sum_{j=1}^N j^{-s}$ is the N -th generalized harmonic number of order s . Note that s is a key parameter of the Zipf distribution and is close to 1 but in general not equal to 1. We consider $s \in (0, 1) \cup (1, 2)$ in our analysis. In other words, $f(i; s, N)$ is the likelihood of the i -th ranked content object being requested.

The storage capability of CCN routers can be provisioned in either a non-coordinated or a coordinated manner. In the non-coordinated provision case, each router stores only the most popular contents locally so that clients can fetch them from directly connected routers. Since no coordination is necessary, routers make caching decisions independently and do not incur any cost of coordination. On the contrary, in the coordinated provision case, routers store popular contents in a coordinated and collaborative manner. Hence, more contents will be cached at peer routers in the network, and clients may experience less delay when fetching contents.

However, the coordination among routers comes at certain costs. Suppose that there exists a conceptually centralized coordinator (*i.e.*, C) in Figure 2. Then in order to manage the coordinated caching across the network, the coordinator has to collect the information of content stores from and disseminate necessary messages to all routers in the network. Note that the coordinator is conceptually centralized; in practice, it can be implemented in a fully distributed manner.

In the following subsection, we will develop a performance-cost model to characterize the network performance and costs. Our model is general and unifies the coordinated and non-coordinated caching mechanisms.

B. Performance-Cost Model

We primarily consider the network routing performance and the coordination cost in CCN.

1) *Routing performance*: The network routing performance refers to the performance of routing contents from one point to another in a content-centric network. The network carriers define their own routing performance metrics, for instance, the average total number of hops all traffic traverses in the network, or the average latency experienced by end users. In this paper, we use the average latency as the main routing performance metric. When there is no ambiguity, we also refer to the average latency as the routing performance. Note that our model is applicable to other metrics such as the average hop count.

As shown in Figure 2, we denote by d_0 the average latency of serving requests from clients' closest routers (which store the requested contents locally). d_1 represents the average latency of serving a request from a peer router in the given network, namely, a directly connected router does not have the requested content but can fetch it from a peer router

in the same administrative domain. Moreover, d_2 denotes the average latency of fetching contents from the origin. Note that d_1 includes two types of latency: the average latency between a client and its corresponding router (*i.e.*, d_0), and the average latency of transferring contents from peer routers. Therefore, $d_1 > d_0$. Similarly, $d_2 > d_1$. Note also that d_0, d_1 and d_2 collectively reflect the average latency incurred by routing contents in the network. We further define $t_1 = \frac{d_1}{d_0}$ as the *first-tier latency ratio*, $t_2 = \frac{d_2}{d_1}$ as the *second-tier latency ratio*, and $\gamma = \frac{d_2 - d_1}{d_1 - d_0}$ as the *ratio of tiered latency* (or *tiered latency ratio* for short).

We consider a unified general model to formulate both coordinated and non-coordinated caching mechanisms by introducing a parameter $x \in [0, c]$, which denotes the amount of storage capacity allocated for coordinated caching mechanisms at each router. Each router stores in its local storage (*i.e.*, the $c - x$ portion) the top ranked contents in a non-coordinated manner, and all routers collaboratively store $n \cdot x$ contents that are ranked from $c - x + 1$ to $c - x + nx$. We use $f(k; s, N)$ to characterize the probability of the k -th ranked content. Moreover, we compute the overall probability of requesting for the top k contents by

$$F(k; s, N) = \sum_{i=1}^k f(i; s, N) = \frac{H_{k,s}}{H_{N,s}}, \quad k = 1, 2, \dots$$

where $H_{k,s}$ and $H_{N,s}$ are the k -th and N -th harmonic numbers of order s . Therefore, the average latency of serving a content request is

$$\begin{aligned} T(x; d_0, d_1, d_2) &= F(c - x; s, N) \cdot d_0 \\ &\quad + [F(c - x + xn; s, N) - F(c - x; s, N)] \cdot d_1 \\ &\quad + [1 - F(c - x + xn; s, N)] \cdot d_2. \end{aligned} \quad (2)$$

The rationale is that each router uses the $c - x$ portion of its storage to store the most popular contents, and use the remaining x portion to store (distinct) contents in a coordinated manner. As a result, the total number of unique contents stored in all routers is $(c - x) + x \cdot n$ (recall that the content object size is normalized to 1).

2) *Coordination cost*: The coordination cost refers to the cost incurred by coordinated provisioning of the storage capability among all participating routers. We consider three types of costs incurred to coordinate the in-network content caching decisions, including the *computational cost* of calculating the optimal storage provisioning policy for all routers and all contents, the *communication cost* of collecting statistics from and distributing optimal policies to all routers, and the *enforcement cost* of implementing the optimal policy at each individual router.

Among these three types of costs, the communication cost is a function of x . More specifically, the states of the coordinated storage at each router should be communicated to other routers in order for all routers to collectively compute the optimal policy. Such communication cost can

contribute non-negligible amount of traffic. Many studies suggest that ISPs tend to define their own piece-wise linear functions to capture such cost (see, *e.g.*, [22]); therefore, we adopt a linear function to capture the communication cost.

Note that the computational cost is dependent on the number of contents (*i.e.*, N), coordinated contents per router (*i.e.*, x), and many other factors such as the network topology and content popularity distribution. Recall that the number of contents is typically extremely large and is most likely to dominate other factors. Note also that the enforcement cost is independent of x , for instance, the complexity of hash-based algorithms for matching requests with stored contents does not depend on the number of stored contents. Therefore, we consider both the computational cost and the enforcement cost as constants, and characterize the overall coordination cost in CCN by

$$W(x; w, \hat{w}) = w \cdot n \cdot x + \hat{w}, \quad (3)$$

where \hat{w} is the invariant computational and enforcement cost, w is the expected communication cost per content per router (referred to as the *unit coordination cost* for short), and $w \cdot n \cdot x$ is the overall communication cost.

IV. PROBLEM FORMULATION AND ANALYSIS

In this section, we formulate the problem of how CCN routers' storage capability should be provisioned as an optimization problem, and systematically study the optimal solution, *i.e.*, the optimal provisioning strategy for the in-network storage capability. More specifically, we provide a rigorous proof for the existence and uniqueness of the optimal strategy.

A. Problem Formulation

In practice, the network routing performance and the coordination cost may not be well aligned. Inspired by many studies where there exist multiple types of network performance and costs (see, *e.g.*, [23], [24]), we introduce a *trade-off weight parameter* $\alpha \in [0, 1]$ and formulate the overall performance/cost as a convex combination of the routing performance¹ and the coordination cost:

$$T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2) = \alpha \cdot T(x; d_0, d_1, d_2) + (1 - \alpha) \cdot W(x; w, \hat{w}). \quad (4)$$

The goal of coordinating in-network caching is to find the optimal x^* that minimizes T_w , namely,

$$x^*(\alpha) = \arg \min_x T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2). \quad (5)$$

We define $\ell(\alpha) = \frac{x(\alpha)}{c}$ as the *coordination level* and refer to $\ell^*(\alpha) = \frac{x^*(\alpha)}{c}$ as the *optimal strategy*, namely, the optimal percentage of coordinated storage.

¹Recall that we use the average latency to measure the routing performance.

In order to ease the analysis and derive meaningful results, we apply the assumption that N is sufficiently large and approximate $F(x; s, N)$ using a continuous function

$$F(x; s, N) \approx \frac{\int_1^x t^{-s} dt}{\int_1^N t^{-s} dt} = \frac{x^{1-s} - 1}{N^{1-s} - 1}, \quad s \in (0, 1) \cup (1, 2). \quad (6)$$

B. Existence of Optimal Strategy

By checking the existence of the first-order derivative and the positivity of the second-order derivative of $T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2)$, both with respect to x , we can formally prove the following lemma, which suggests the existence of the optimal strategy (see the Appendix for a complete proof):

Lemma 1: $T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2)$ is a convex function of x . The optimal solution to (5) exists, if the following conditions for system parameters hold:

- $0 \leq x \leq c$ and $c > 0$,
- The number of contents is sufficiently large ($N \gg 1$),
- the number of routers $n > 1$,
- $0 < s < 2$ and $s \neq 1$, and
- $d_0 < d_1 \leq d_2$.

We remark that the conditions for guaranteeing the existence of the optimal strategy are reasonable and are most likely to hold in practice. The number of contents is typically large, *i.e.*, $N \gg 1$, and s is typically a positive number between 0 and 2 (see, *e.g.*, [17], [18], [19]). The number of routers n could range from a dozen to a couple of hundred in an administrative domain.

Additionally, as far as the latency is concerned, the condition $d_0 < d_1 \leq d_2$ is most likely to hold in realistic networks. First, d_0 can be approximated by the latency between the end users and their first-hop routers. Its typical values are about 100 milliseconds in cellular networks (see, *e.g.*, [25]), 10–20 milliseconds in cable access networks (see, *e.g.*, [26]), and 30 milliseconds in ADSL access networks (see, *e.g.*, [27]). Second, $d_1 - d_0$ can be approximated by the latency between routers in the same administrative domain, and its values typically range from a few to 20 milliseconds on average, depending on the geographical coverage of the network (*e.g.*, [27]). Last, d_2 typically ranges from more than one hundred to a couple of hundred milliseconds with heavy-tailed distribution (see, *e.g.*, [28]).

C. Uniqueness of Optimal Strategy

The following lemma characterizes the optimal strategy ℓ^* , which can be proven by letting the first-order derivative of $T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2)$ equal to zero:

Lemma 2: The optimal strategy ℓ^* satisfies the following equation:

$$a\ell^{-s} = (1 - \ell)^{-s} + b, \quad (7)$$

where $a \approx \gamma \cdot n^{1-s}$ and $b \approx \frac{1-\alpha}{\alpha} \cdot \frac{N^{1-s}-1}{1-s} \cdot \frac{(n-1)w}{d_1-d_0} c^s$, $\alpha \in [0, 1]$, $\gamma > 0$, $s \in (0, 1) \cup (1, 2)$, $n > 0$, $N > 0$, $c > 0$, $w > 0$, and $d_1 - d_0 > 0$.

We next apply Lemma 2 to prove the uniqueness of ℓ^* in Theorem 1 as follows.

Theorem 1: There exists a unique solution to (7).

Proof: Given a particular trade-off weight parameter α , we define $y(\ell) = a\ell^{-s}$ and $z(\ell) = (1-\ell)^{-s} + b$, respectively. Firstly, we show that within $\ell \in (0, 1)$, both $y(\ell)$ and $z(\ell)$ are continuous and monotonically decreasing and increasing, respectively.

Continuity. Since for any $\ell \in (0, 1)$, both derivatives of $y(\ell)$ and $z(\ell)$ exist; namely, $\frac{dy(\ell)}{d\ell} = -a\ell^{-s-1}$ and $\frac{dz(\ell)}{d\ell} = -s(1-\ell)^{-s-1}$, $y(\ell)$ and $z(\ell)$ are continuous with respect to $\ell \in (0, 1)$.

Monotonicity. Given any $0 < \ell_1 < \ell_2 < 1$ and $s \in (0, 1) \cup (1, 2)$, we have

$$y(\ell_1) - y(\ell_2) = a(\ell_1^{-s} - \ell_2^{-s}) = \frac{a}{\ell_1^s \ell_2^s} (\ell_2^s - \ell_1^s) > 0,$$

thus $y(\ell)$ monotonically decreases, when increasing $\ell \in (0, 1)$. Similarly, we have $1 > 1 - \ell_1 > 1 - \ell_2 > 0$ and the following inequality holds

$$z(\ell_1) - z(\ell_2) = \frac{\ell_1^s - \ell_2^s}{(1-\ell_1)^s (1-\ell_2)^s} < 0,$$

which in turn proves that $z(\ell)$ monotonically increases, when increasing $\ell \in (0, 1)$.

Moreover, we observe that $\lim_{\ell \rightarrow 0} y(\ell) = \infty$, $\lim_{\ell \rightarrow 1} y(\ell) = a$, $\lim_{\ell \rightarrow 0} z(\ell) = 1 + b$, and $\lim_{\ell \rightarrow 1} z(\ell) = \infty$. Hence, $y(\ell)$ and $z(\ell)$ must have a unique intersection point in the range $(0, 1)$. ■

D. Optimal Strategy for Routing Performance Optimization

We next focus on the optimal strategy when the routing performance is the dominant concern (*i.e.*, $\alpha = 1$). We will derive the closed-form optimal strategy and analyze the impacts of various system parameters.

Theorem 2: When $\alpha = 1$, the unique optimal strategy for (5) is

$$\ell^* = \frac{x^*}{c} \approx \frac{1}{\gamma^{\frac{1}{s}} n^{1-\frac{1}{s}} + 1}. \quad (8)$$

Proof: Let $\frac{\partial T(x; d_0, d_1, d_2)}{\partial x} = 0$, we have

$$\frac{\alpha(1-s)(d_2-d_1)}{N^{1-s}-1} (\gamma(c-x)^{-s} - (n-1)(c+(n-1)x)^{-s}) = 0.$$

Since n is typically sufficiently large, both $(n-1)^{-1} \approx 0$ and $n-1 \approx n$ hold. Then the above equation can be further simplified as

$$(1-\ell)^s = \gamma \cdot (n-1)^{-1} \cdot (1+(n-1)\cdot\ell)^s \approx \gamma \cdot n^{s-1} \cdot \ell^s, \quad (9)$$

where $\gamma > 0$, $s \in (0, 2)$ and $s \neq 1$, $n > 0$. Solving (9) yields the optimal strategy ℓ^* for $\alpha = 1$. ■

It is important to note that ℓ^* is a function of the tiered latency ratio γ (*i.e.*, ratio between $d_2 - d_1$ and $d_1 - d_0$), rather than the absolute values of the individual latencies (such as d_0 , d_1 and d_2). We refer to this property as the *latency scale free* property (or *scale free* for short). This property is particularly desirable and helpful in designing, deploying and provisioning storage capability optimally in a network.

Note that in real networks, the average latencies (*e.g.*, d_0 , d_1 , and d_2) are all bounded in a few to a hundred milliseconds; thus, γ is bounded between 1 and 100 in general, while the number of routers n can scale up dramatically as the network size increases. Hence, we consider how increasing n impacts the optimal strategy ℓ^* while taking γ as a bounded constant. More specifically, when $s \in (0, 1)$, the optimal strategy ℓ^* quickly approaches 1 as n increases; in other words, all routers should dedicate all their storage capacity to coordinated caching when the number of routers is large. However, when $s \in (1, 2)$, the optimal strategy ℓ^* converges to 0 as n increases, meaning that all routers' storage capacity should be dedicated to non-coordinated caching instead.

This observation reveals that $s = 1$ is a singular point; $s \in (0, 1)$ and $s \in (1, 2)$ lead to opposite optimal strategies. We will evaluate and discuss in more details how various factors affect the optimal strategy ℓ^* in the next section.

E. Performance Gain

We now quantify the performance gain as a result of the optimal strategy ℓ^* . We consider two types of performance gain, the origin load reduction G_O from the origin server's perspective, and the routing performance improvement G_R from the network carrier's perspective.

1) *Origin load reduction G_O :* G_O is the total load reduction on the origin server, namely, the improvement on the total traffic load incurred on the origin server under the optimal strategy compared to the non-coordinated caching strategy. Based on the assumption of unit-size contents, the traffic load that the origin server sees can be directly expressed as the ratio between the number of contents served by the origin server when using the optimal strategy and when using the non-coordinated strategy.

More specifically, the traffic demand at the origin using the optimal caching strategy is $1 - F(c + (n-1)x^*; s, N)$, while the demand at the origin using the non-coordinated caching strategy is $1 - F(c; s, N)$. Therefore, the ratio of expected load at the origin with the optimal strategy over the non-coordinated strategy is

$$\begin{aligned} G_O &= 1 - \frac{1 - F(c + (n-1)x^*; s, N)}{1 - F(c; s, N)} \\ &= \frac{(c + (n-1)x^*)^{1-s} - c^{1-s}}{N^{1-s} - c^{1-s}} \end{aligned}$$

2) *Routing Performance Improvement G_R :* G_R is the total improvement on the routing performance, namely, the

Table II
TOPOLOGIES USED IN EVALUATIONS

Topology	$ V $	$ E $	Region	Type
Abilene	11	28	North America	Educational
CERNET	36	112	East Asia	Educational
GEANT	23	74	Europe	Educational
US-A	20	80	North America	Commercial

improvement on the overall routing performance under the optimal strategy versus the non-coordinated strategy.

Note that when routers are non-coordinated (*i.e.*, $x = 0$), the routing performance in (2) is

$$T(0; d_0, d_1, d_2) = \frac{(N^{1-s} - c^{1-s}) \cdot d_2 + (c^{1-s} - 1) \cdot d_0}{N^{1-s} - 1}.$$

Therefore, the overall routing performance improvement is

$$G_R = 1 - \frac{T(x^*, d_0, d_1, d_2)}{T(0, d_0, d_1, d_2)}.$$

V. EVALUATIONS

In this section, we quantify the optimal strategy and the performance gain through numerical evaluations on four real network topologies. More specifically, we first introduce the four real network datasets and our evaluation (parameter) settings. Then, we evaluate how various factors affect the optimal strategy ℓ^* and the performance gain obtained when applying ℓ^* .

A. Datasets and Evaluation Setup

We use four real network topologies in our evaluations, namely, Internet2 (the Abilene Network) [29], CERNET [30], GEANT [31], and an anonymized tier-1 network carrier US-A in North America.

In particular, Abilene is a high-performance backbone network established by the Internet2 community in the late 1990s. The old Abilene network was retired and became the Internet2 network in 2007. It has 11 regional network aggregation points and the backbone connections among them are primarily OC192 or OC48. CERNET is the first nation-wide education and research network in China. It is funded by the government and managed by the Ministry of Education in China. It is constructed and operated by Tsinghua University and other leading universities in China, with 36 aggregation points and OC192 links. GEANT is a pan-European data network dedicated to the research and education community. Together with Europe's national research networks, GEANT connects 40 million users in over 8,000 institutions across 40 countries. GEANT has 23 aggregation points with links ranging from OC3 to OC192.

Each network topology, denoted by $G = (V, E)$, has the location information for each router $i \in V$ (the total number of routers $n = |V|$). We also obtain the pair-wise latency d_{ij} for every pair of routers $i, j \in V$ in each topology.

Let d_{ij} denote the average latency between two routers i and j . We estimate the unit coordination cost w by taking the maximum expected latency among routers, namely,

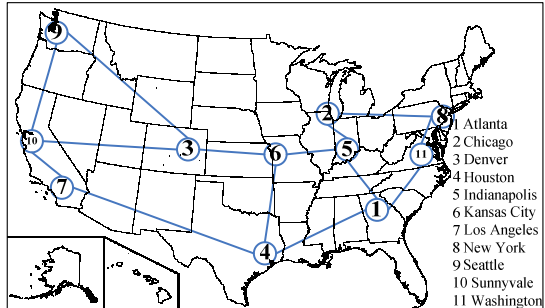


Figure 3. The Abilene network topology

$w = \max_{i,j \in V} d_{ij}$, since the communications among routers (or between the conceptual centralized coordinator and all routers) can be implemented in parallel, and the maximum latency plays a key role in determining the speed of converging to the optimal strategy.

Additionally, let h_{ij} denote the hop count of the shortest path between i and j . The average routing performance, measured by the average hop counts of the shortest paths among router pairs, is $(d_1 - d_0) = \frac{1}{|V|^2} \sum_{i,j \in V} h_{ij}$. Note that the routing performance can also be measured by the other metrics, *e.g.*, the average pair-wise latency $(d_1 - d_0) = \frac{1}{|V|^2} \sum_{i,j \in V} d_{ij}$. In our evaluations, we applied both metrics and observed similar results; thus we only present the results for the routing performance measured by the hop count.

We summarize the statistics of the four networks in Table II. We show the topological structure of the Abilene network in Figure 3, and omit the other three networks for brevity. Table III lists the topological parameters obtained from four real networks.

Table III
TOPOLOGICAL PARAMETERS

Topology	n	w (ms)	$d_1 - d_0$ (ms)	$d_1 - d_0$ (hops)
Abilene	11	22.3	14.3	2.4182
CERNET	36	33.3	16.2	2.8238
GEANT	23	27.8	16.0	2.6008
US-A	20	26.7	15.7	2.2842

We list in Table IV the general empirical ranges of network parameters, as well as detailed parameter settings in our evaluations. Note that we choose n , w , and $d_1 - d_0$ from the real network topologies, as listed in Table III. We obtain similar results for all four network topologies, so we only present the results for the topology of US-A for brevity. Moreover, in order to investigate how the topological parameters affect the optimal strategy, we also vary the number of routers (n) and the communication cost (w) in our evaluations.

We comment that the exact values that each parameter takes can vary over time across different networks; however, the overall trends are less likely to change.

B. Optimal strategy ℓ^*

We first evaluate how various parameters affect the optimal strategy ℓ^* .

Table IV
SYSTEM PARAMETERS USED IN ANALYSIS.

Parameters	α	γ	s	n	N	c	$w(\text{ms})$	$d_1 - d_0(\text{hops})$
Ranges	[0, 1]	1 ~ 10	$(0, 1) \cup (1, 2)$	10 ~ 500	10^9 $\sim 10^{12}$	10^6 $\sim 10^9$	10 ~ 100	1 ~ 10
Figure 4, 8, 12	(0, 1)	{2, 4, 6, 8, 10}	0.8	20	10^6	10^3	26.7	2.2842
Figure 5, 9, 13	[0.2, 1]	5	$[0.1, 1) \cup (1, 1.9]$	20	10^6	10^3	26.7	2.2842
Figure 7, 11	[0.2, 1]	5	0.8	20	10^6	10^3	10 ~ 100	2.2842
Figure 6, 10	[0.2, 1]	5	0.8	10 ~ 500	10^6	10^3	26.7	2.2842




Figure 4. The trade-off parameter




Figure 5. The Zipf exponent




Figure 6. The network size




Figure 7. The unit coordination cost

1) *Trade-off parameter α* : We first investigate the impacts of the trade-off weight parameter α to the optimal strategy ℓ^* in Figure 4.

We observe that when α increases, namely, the routing performance is weighted more than the coordination cost, the optimal strategy ℓ^* increases monotonically from 0 to 1. This happens because as the routing performance becomes more dominant in the objective function (4), an increasingly larger portion of the storage should be dedicated to the coordinated caching in order to optimize the overall network performance and cost.

We also observe that for the same α , a higher γ leads to a higher level of coordination. Moreover, given a certain γ , when α is relatively small, ℓ^* increases slowly over α . However, when α is sufficiently large, ℓ^* grows rapidly and becomes more sensitive to changes of α .

These interesting phenomena suggest that α should be adjusted carefully when it is in the sensitive range, which is governed by other parameters, *e.g.*, γ . For example, as shown in Figure 4, when $\gamma=2$, the sensitive range is around $\alpha \in [0.2, 0.4]$, and the range shifts to $[0.6, 0.8]$ when $\gamma=10$.

2) *Zipf exponent s* : We observe in Figure 5 that as s increases, ℓ^* exhibits various trends over s . Note that $s=1$ is a singular point, and is taken away from the range of s , because $s=1$ leads to a constant routing performance $T(x; d_0, d_1, d_2) = d_2$, which is invariant to the coordination level ℓ .

We make the following observations. Firstly, for $\alpha=1$, *i.e.*, only the routing performance is considered, the optimal strategy ℓ^* decreases from 1 to 0.35, as s changes from 0 to 2. This observation confirms our theoretical results presented in Theorem 2, namely, for $s \in (0, 1)$ (resp. $s \in (1, 2)$), ℓ^* converge to 1 (resp. 0), with an increasing n .

Secondly, when $\alpha < 1$, the optimal strategy ℓ^* converges to 0, namely, non-coordinated caching mechanisms are more preferred, when s approaches 0. This happens because

caching is becoming less effective (due to less contents are popular enough to stay in routers' storage) and the coordination cost is gradually dominating the routing performance when using coordinated caching mechanisms. Moreover, for $0 \leq \alpha < 1$, there exists a maximum ℓ^* around $0.5 \sim 0.9$; while in reality, s turns out to be approximately around $0.5 \sim 0.9$ (see, *e.g.*, [17], [18], [19]). This illustrates that in practice, the optimal strategy ℓ^* usually indicates a higher coordination level.

Lastly, the optimal strategy ℓ^* decreases when α is decreasing; namely, the higher the weight on the coordination cost is, the lower the optimal coordination level is. This means that when the coordination cost is the major concern, non-coordinated caching mechanisms are more preferred.

3) *Network size n* : Figure 6 shows how ℓ^* changes with a varying size of an intradomain network (*i.e.*, the number of routers n).

We observe that the optimal strategy ℓ^* decreases as n increases, because the more routers a network has, the higher the coordination cost is. Moreover, for a given network size, ℓ^* increases drastically as we put a higher weight on the routing performance (*i.e.*, α increases), suggesting that a higher coordination level can help to reduce more traffic and thus to further improve the routing performance.

4) *Unit coordination cost w* : We observe in Figure 7 that when the routing performance dominates in (4), *i.e.*, $\alpha=1$, ℓ^* is a constant close to 1, whereas for small α , *e.g.*, $\alpha < 0.4$, ℓ^* decreases drastically as the unit coordination cost w increases. This suggests that a low coordination level can help improve the overall network performance and cost when w is large. Moreover, a larger α leads to a larger ℓ^* for the same w , which confirms the results presented in Figure 4. This trend is also similar to the observation we made in Figure 6.




Figure 8. α vs. G_O




Figure 9. s vs. G_O




Figure 10. n vs. G_O




Figure 11. w vs. G_O

C. Performance Gain

We next evaluate the performance gain of the optimal strategy from both the origin's and the carrier's perspectives.

1) *Origin load reduction G_O* : We observe in Figure 8 that as the trade-off parameter α increases, the gain on origin load reduction increases, due to the fact that a higher ℓ^* allows routers to store more contents. Note that a higher γ leads to a higher overall origin load reduction. We also observe in Figure 9 that for a relatively smaller α , the overall origin load reduction is higher and reaches the maximum at around $s = 1.3$. Note that $s = 1$ is a singular point.

Figure 10 illustrates how the total number of routers affects the load reduction at the origin server. When α is relatively small, the origin load reduction stays roughly constant over n , and a higher α leads to a higher origin load reduction. However, when α is approaching 1, the effect of the network size emerges; namely, the origin load reduction increases with an increasing n . This observation indicates that when the coordination cost is not dominated by the routing performance (*i.e.*, α is small), the network size n has nearly no effect on the origin load reduction.

Moreover, Figure 11 indicates that when α is small (*e.g.*, $0 \leq \alpha < 0.4$), the origin load reduction decreases rapidly as the unit coordination cost increases. The reason is that when the unit coordination cost increases, the optimal coordination level ℓ^* decreases drastically, meaning that routers can store a much smaller number of distinct contents, and eventually the origin server has to serve more requests due to cache misses at routers. This phenomenon implies that for a large w , the gain on origin load reduction is low. In addition, when α is relatively large, or in other words the routing performance is weighted more, the origin load reduction becomes almost invariant with respect to a varying unit coordination cost.

2) *Routing performance improvement G_R* : We observe in Figure 12 that as we increase the weight of the routing performance (*i.e.*, α increases), the overall routing performance improvement G_R increases, and a higher γ will further raise the overall level of improvement. In particular, the routing performance improvement can be as significant as 60–90% when the trade-off parameter and the tiered latency ratio are reasonably large (*e.g.*, $\alpha \geq 0.5$ and $\gamma \geq 8$).

Additionally, Figure 13 shows that when s is further away from 1, *i.e.*, closer to 0 or 2, the routing performance




Figure 12. α vs. G_R




Figure 13. s vs. G_R

improvement is smaller; whereas for s close to 1 ($s = 1$ is a singular point), the routing performance improvement is large (reaching the maximum at around $s = 1$), which suggests that for those scenarios with the Zipf exponent s closer to 1, the optimal strategy is more efficient since more significant improvement on the routing performance can be achieved. When varying the parameters w and n , we observe results similar to Figure 10 and 11, therefore we omit them here for brevity.

VI. RELATED WORK

Content caching has been a key component of Internet-based services for many years (see, *e.g.*, [6], [32], [33]), and there have been many studies in literature on content caching techniques (see, *e.g.*, [11], [12], [15], [16], [34], [35], [36], [37]). In particular, coordinated (or collaborative) content caching has been studied extensively. Researchers have investigated the effectiveness of collaborative caching (see, *e.g.*, [11], [38]) and proposed numerous collaborative caching schemes for both general networks and networks with specific structures, including general Internet-based content distribution (see, *e.g.*, [12], [16], [39]), delivering special types of contents (*e.g.*, [13]), content caching in networks with special topological structures (*e.g.*, [15]), and content caching in ad hoc networks (*e.g.*, [34]), mobile broadcast environment (*e.g.*, [35]), 3G networks (*e.g.*, [37]), and peer-to-peer networks (*e.g.*, [36]).

Our work differs from these studies in two ways. First, our network model for content-centric networks is novel, with storage at router level as an inherent in-network capability. As a result, we formulate the problem by focusing on the overall network performance and cost from the network carriers' perspectives. Most of previous studies assumed different network models (*e.g.*, overlay network models) and did not characterize such overall network performance and cost. Secondly, we consider both the routing performance

and the coordination cost, and investigate the trade-offs between them.

Additionally, there exists a line of recent work on emerging Content-Centric Networking [9] and Named Data Networking (NDN) [10], where content storage becomes an inherent capability of network routers. CCN and NDN are closely related, with the latter focusing more on fundamental research. CCN/NDN has become one of the representative alternatives for the future Internet architecture. Both CCN and NDN have attracted much attention. There has been an increasingly large body of literature on CCN and NDN, to name a few, naming and name resolution (*e.g.*, [40], [41]), data transfer (*e.g.*, [42]), flow and traffic control (*e.g.*, [43]), routing and router design (*e.g.*, [44], [45]), mobile and ad hoc networks (*e.g.*, [46], [47]), privacy (*e.g.*, [48]), and caching (*e.g.*, [49], [50], [51], [52], [53]). In particular, in [50], Xie *et al.* proposed a traffic-engineering-guided content placement and caching algorithm for CCN; and in [51], Sourlas *et al.* proposed content placement and caching algorithms to minimize overall traffic cost of content delivery, specifically designed for CCN. However, none of the existing work addresses the optimal strategy of coordinated content caching and investigates the trade-offs between the routing performance and the coordination cost in the context of CCN/NDN. To the best of our knowledge, our work is the first attempt to formally investigate and providing insights in addressing these issues.

VII. CONCLUSION

In content-centric networks, routers possess both the routing and the in-network storage capability, which raises new challenges in network provisioning, namely, how to optimally provision individual routers' storage capability for content caching, so as to optimize the overall network performance and provisioning cost.

In this paper, we developed a holistic model to quantify the overall network performance of routing contents to clients and the overall provisioning cost incurred by coordinating the in-network storage capability. Based on this model, we derived the optimal strategy for optimizing the overall network performance and cost, and evaluated the optimal strategy using real network topologies. We observed interesting phenomena; for example, different ranges of the Zipf exponent, a key parameter of the content popularity distribution, can lead to opposite optimal strategies, and the trade-off parameter α has great impacts on the stability of the optimal strategy. Our evaluation results also demonstrated significant gain on both the load reduction at origin and the improvement on the routing performance.

There are several directions for our future work. First of all, we plan to generalize the holistic model to capture the network dynamics and design online self-adaptive algorithms to adjust the coordination level. We are also interested in a heterogeneous model where the network has

heterogeneous storage capability and performance metrics (*e.g.*, both the routers' storage capacity and the metrics for the network performance may vary). Last but not least, we will perform extensive Internet-based experiments to understand the implementation challenges.

ACKNOWLEDGMENT

Yanhua Li and Zhi-Li Zhang are supported in part by the US NSF grant CNS-1017092, CNS-1017647, the DTRA grant HDTRA1-09-1-0050, and a DoD ARO MURI Award W911NF-12-1-0385. Haiyong Xie is supported in part by the National Natural Science Foundation of China (NSFC) under Grant No. 61073192, by the Grand Fundamental Research Program of China (973 Program) under Grant No. 2011CB302905, by the New Century Excellent Talents Program under grant No. NCET-09-0921, and by the Fundamental Research Funds for Central Universities under Grant No. WK0110000014. Yonggang Wen is supported in part by the Start-up Grant from NTU and the tier-1 grant (RG31/11) from MOE, Singapore.

REFERENCES

- [1] Netflix, <http://www.netflix.com>.
- [2] YouTube, <http://www.youtube.com>.
- [3] J. Zhou, Y. Li, V. K. Adhikari, and Z.-L. Zhang, "Counting youtube videos via random prefix sampling," in *ACM IMC*, 2011.
- [4] Cisco, "Cisco visual networking index: Forecast and methodology, 2009–2014," White Paper, 2010.
- [5] —, "Hyperconnectivity and the approaching zettabyte era," White Paper, 2010.
- [6] Akamai Technologies, <http://www.akamai.com>.
- [7] AT&T, "AT&T content acceleration service," Product Brief, 2011.
- [8] Level3 Communications, <http://www.level3.com>.
- [9] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *ACM CoNEXT*, 2009.
- [10] L. Zhang, D. Estrin, J. Burke, V. Jacobson, J. D. Thornton, D. K. Smetters, B. Zhang, G. Tsudik, K. Claffy, D. Krioukov, D. Massey, C. Papadopoulos, T. Abdelzaher, L. Wang, P. Crowley, and E. Yeh, "Named data networking (NDN) project," Palo Alto Research Center (PARC), Tech. Rep. NDN-0001, 2010.
- [11] A. Wolman, G. M. Voelker, N. Sharma, N. Cardwell, A. Karlin, and H. M. Levy, "On the scale and performance of cooperative web proxy caching," *SIGOPS Oper. Syst. Rev.*, vol. 33, no. 5, pp. 16–31, 1999.
- [12] L. Fan, P. Cao, J. Almeida, and A. Z. Broder, "Summary cache: a scalable wide-area web cache sharing protocol," *IEEE/ACM Transactions on Networking*, vol. 8, no. 3, pp. 281–293, 2000.
- [13] J. Ni and D. Tsang, "Large-scale cooperative caching and application-level multicast in multimedia content delivery networks," *IEEE Comm. Mag.*, vol. 43, no. 5, pp. 98–105, 2005.
- [14] F. Sailhan and V. Issarny, "Cooperative caching in ad hoc networks," in *Int'l Conf. on Mobile Data Management*, 2003.
- [15] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *IEEE INFOCOM*, 2010.
- [16] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl, "Globally distributed content delivery," *IEEE Internet Computing*, vol. 6, no. 5, pp. 50–58, 2002.
- [17] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: Evidence and implications," in *IEEE INFOCOM*, 1999.
- [18] X. Cheng, C. Dale, and J. Liu, "Statistics and social network of youtube videos," in *IEEE IWQoS*, 2008.

[19] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "Youtube traffic characterization: a view from the edge," in *ACM IMC*, 2007.

[20] BitTorrent, <http://www.bittorrent.com>.

[21] eMule, <http://www.emule-project.net>.

[22] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *IEEE INFOCOM*, 2000.

[23] W. Jiang, R. Zhang-Shen, J. Rexford, and M. Chiang, "Cooperative content distribution and traffic engineering in an ISP network," in *ACM SIGMETRIC*, 2009.

[24] Y. Li, Z.-L. Zhang, and D. Boley, "The routing continuum from shortest-path to all-path: A unifying theory," in *IEEE ICDCS*, 2011.

[25] C. Serrano, B. Garriga, J. Velasco, J. Urbano, S. Tenorio, and M. Sierra, "Latency in broad-band mobile networks," in *IEEE VTC Spring*, 2009.

[26] N. Barakat and T. Darcie, "Delay characterization of cable access networks," *IEEE Communications Letters*, vol. 11, no. 4, pp. 357–359, 2007.

[27] K. C. Kang, K.-D. Nam, D. J. Jeon, and S. Y. Kim, "Delay characteristics of high-speed Internet access network," in *APNOMS*, 2003.

[28] B. Huffaker, M. Fomenkov, D. J. Plummer, D. Moore, and k claffy, "Distance metrics in the Internet," in *IEEE Int'l Telecommunications Symposium*, 2002.

[29] Abilene, "Internet2 IP IGP Metrics," <http://noc.net.internet2.edu/i2network/maps--documentation/maps.html>.

[30] CERNET, "CERNET Topology," <http://www.edu.cn/20060111/3170194.shtml>.

[31] GEANT, "GEANT Topology," http://archive.geant.net/upload/pdf/Topology_Oct_2004.pdf.

[32] Limelight Networks, <http://www.limelight.com>.

[33] Amazon CloudFront Express, <http://www.amazon.com/cloudfront>.

[34] L. Yin and G. Cao, "Supporting cooperative caching in ad hoc networks," *IEEE Transactions on Mobile Computing*, vol. 5, no. 1, pp. 77–89, 2006.

[35] C.-Y. Chow, H. V. Leong, and A. T. S. Chan, "Distributed group-based cooperative caching in a mobile broadcast environment," in *ACM MDM*, 2005.

[36] M. Hefeeda and B. Noorizadeh, "On the benefits of cooperative proxy caching for peer-to-peer traffic," *IEEE Transactions on Parallel and Distributed Systems*, vol. 21, no. 7, pp. 998–1010, 2010.

[37] J. Erman, A. Gerber, M. T. Hajiaghayi, D. Pei, S. Sen, and O. Spatscheck, "To cache or not to cache: The 3G case," *IEEE Internet Computing*, vol. 15, no. 2, pp. 27–34, 2011.

[38] S. Gadde, J. Chase, and M. Rabinovich, "Web caching and content distribution: a view from the interior," *Computer Communications*, vol. 24, no. 2, pp. 222 – 231, 2001.

[39] D. Wessels and K. Claffy, "Internet cache protocol (ICP), version 2," IETF, RFC 2186, 1997.

[40] A. Ghodsi, T. Koponen, J. Rajahalme, P. Sarolahti, and S. Shenker, "Naming in content-oriented architectures," in *ACM ICN Workshop*, 2011.

[41] M. D'ambrosio, C. Dannewitz, H. Karl, and V. Vercellone, "MDHT: A hierarchical name resolution service for information-centric networks," in *ACM ICN Workshop*, 2011.

[42] G. Carofiglio, M. Gallo, L. Muscariello, and D. Perino, "Modeling data transfer in content-centric networking," in *ITC*, 2011.

[43] S. Oueslati, J. Roberts, and N. Sbihi, "Flow-aware traffic control for a content-centric network," in *IEEE INFOCOM*, 2012.

[44] S. DiBenedetto, C. Papadopoulos, and D. Massey, "Routing policies in named data networking," in *ACM ICN Workshop*, 2011.

[45] D. Perino and M. Varvello, "A reality check for content centric networking," in *ACM ICN Workshop*, 2011.

[46] M. Meisel, V. Pappas, and L. Zhang, "Ad hoc networking via named data," in *ACM MobiArch*, 2010.

[47] M. Varvello, I. Rimac, U. Lee, L. Greenwald, and V. Hilt, "On the design of content-centric MANETs," in *Int'l Conf. on Wireless On-Demand Network Systems and Services*, 2011.

[48] S. Arianfar, T. Koponen, S. Shenker, and B. Raghavan, "On preserving privacy in information-centric networks," in *ACM ICN Workshop*, 2011.

[49] I. Psaras, R. Clegg, R. Landa, W. Chai, and G. Pavlou, "Modelling and evaluation of ccn-caching trees," in *IFIP Networking*, 2011.

[50] H. Xie, G. Shi, and P. Wang, "TECC: Towards collaborative in-network caching guided by traffic engineering," in *IEEE INFOCOM*, 2012.

[51] V. Sourlas, L. Gkatzikis, and L. Tassioulas, "On-line storage management with distributed decision making for content-centric networks," in *EURO-NGI*, 2011.

[52] S. Guo, H. Xie, and G. Shi, "Collaborative forwarding and caching in content centric networks," in *IFIP Networking*, 2012.

[53] G. Tyson, S. Kaune, S. Miles, Y. El-khatib, A. Mauthe, and A. Taweel, "A trace-driven analysis of caching in content-centric networks," in *ICCCN*, 2012.

VIII. APPENDIX

A. Proof of Lemma 1

We accomplish the proof by checking the existence of the first-order derivative and the positivity of the second-order derivative of $T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2)$. Note that by using the approximation in (6), $F(x; s, N)$ is differentiable, and its first- and second-order derivatives are $\frac{\partial F(x, s, N)}{\partial x} = \frac{1-s}{N^{1-s}-1} \cdot x^{-s}$ and $\frac{\partial^2 F(x, s, N)}{\partial x^2} = -\frac{s(1-s)}{N^{1-s}-1} \cdot x^{-s-1}$, respectively.

We rewrite (4) as below:

$$T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2) = \alpha(d_0 - d_1) \frac{(c-x)^{1-s} - 1}{N^{1-s} - 1} + \alpha(d_1 - d_2) \frac{(c + (n-1)x)^{1-s} - 1}{N^{1-s} - 1} + \alpha d_2 + (1 - \alpha)(wnx + \hat{w}).$$

Its first-order derivative $\frac{\partial T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2)}{\partial x}$ is

$$\frac{(1-s)\alpha}{N^{1-s}-1} [(d_1 - d_0)(c-x)^{-s} - (d_2 - d_1)(n-1)(c + (n-1)x)^{-s}] + (1 - \alpha)wn, \quad (10)$$

and its second-order derivative $\frac{\partial^2 T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2)}{\partial x^2}$ is

$$\frac{s(1-s)\alpha}{N^{1-s}-1} [(d_1 - d_0)(c-x)^{-s-1} - (d_2 - d_1)(n-1)^2(c + (n-1)x)^{-s-1}].$$

The positivity of (10) can be proven as follows. Since $c + (n-1)x > 0$ holds, we have $(c + (n-1)x)^{-s-1} > 0$, with $c \geq x \geq 1$ and $n \geq 1$. Similarly, $(c-x)^{-s-1} \geq 0$. Moreover, we assume that the number of contents in the network is sufficiently large, *i.e.*, $N \gg 1$. When $s \in (0, 1)$, both $1-s > 0$ and $N^{1-s} - 1 > 0$ hold. In contrast, when $s \in (1, 2)$, both $1-s < 0$ and $N^{1-s} - 1 < 0$ hold. Hence, $\frac{s(1-s)}{N^{1-s}-1} > 0$ always holds true for any $s \in (0, 1) \cup (1, 2)$ and $N > 1$. Therefore, with $d_2 - d_1 \geq 0$, $(n-1)^2 > 0$, and $d_1 - d_0 > 0$, we prove $\frac{\partial^2 T_w(x; \alpha, w, \hat{w}, d_0, d_1, d_2)}{\partial x^2} > 0$.