

Responsive Content-Centric Delivery in Large Urban Communication Networks: A LinkNYC Use-Case

Hassan Sinky, Bassem Khalfi, Bechir Hamdaoui, and Ammar Rayes[†]
 Oregon State University, Corvallis, OR 97331, sinkyh,khalfib,hamdaoui@oregonstate.edu
[†] Cisco Systems, San Jose, CA 95134, rayes@cisco.com

Abstract—Large urban communication networks such as smart cities are an ecosystem of devices and services cooperating to address multiple issues that greatly benefit end users, cities and the environment. LinkNYC is a first-of-its-kind urban communications network aiming to replace all payphones in the five boroughs of New York City (NYC) with kiosk-like structures providing free public Wi-Fi. We consolidate these networks with standalone edge cloud devices known as cloudlets and introduce geographically distributed content delivery cloudlets (CDCs) to store popular Internet content closer to end users; essential in environments with diverse and dynamic content interests. A content-centric and delivery framework is proposed leveraging NYC’s population densities and CDCs for interest-based in-network caching. Analysis shows that although the adoption of multiple CDCs dramatically improves overall network performance, advanced caching policies are needed when considering increased content heterogeneity. Thus, we propose popularity-driven (pLFU) and cooperation-based (sLFU) caching policies at individual CDCs to account for user and content dynamics over time. The amalgamation of urban population densities, multiple CDC placements and smarter caching techniques helps exploit the ultimate benefits of a content-centric urban communications network and dramatically improves overall network performance and responsiveness. Our proposed solutions are validated using LinkNYC as a use-case.

I. INTRODUCTION

The proliferation of Internet devices has resulted in efforts to integrate various wireless access technologies for improved performance, increased services and inter-connectivity of end users. The recent growth in data demand has prompted researchers to come up with new wireless techniques (e.g., MIMO [1], [2], cooperative communication [3], femtocells [4], etc.) and develop new technologies (e.g., cognitive radio [5], LTE [6], etc.) to be able to meet this high demand. This integration allows for large geographic locations to be serviced providing millions of end users with continuous connectivity and optimal quality of experience (QoE). However, the world has seen unprecedented urban population growth over the years. In fact, the number of urban residents has increased by nearly 60 million a year. By 2050, it is estimated that 70% of the world’s population will be living in cities¹. Urban communication networks and content delivery networks have been introduced to leverage these technologies to better service cities and users alike. Content delivery networks are designed

to improve overall network performance by bringing data closer to the geographical locations of users.

Urban communication networks have evolved over the years to address urban challenges through the use of information, communication technology and the Internet. Building such a network infrastructure capable of adequately servicing urban locations has become increasingly difficult due to the sheer number of Internet devices and users (e.g., rapidly increasing numbers of Internet of Things (IoT) devices). Research shows that global mobile data traffic will increase sevenfold reaching 49 exabytes per month by 2021, most of which will be mobile video content, with a percentage projected to reach up to 78% by 2021 [7]. Thus, not only pushing content closer to the end user but also smarter content placement and timely delivery of content are of the utmost importance.

Traditionally, content delivery nodes or datacenters are geographically distributed throughout the world servicing different regions. In large urban networks the same content may be requested by multiple users resulting in the content traversing the network multiple times, over large distances, to and from a remote content delivery node hosting the content. As detailed in Section II, we leverage the LinkNYC infrastructure, a large urban communications network currently underway providing free public Wi-Fi to New York City (NYC) residents, to help propose and design techniques for efficient content delivery in large urban networks. Thus, the focus of this work is twofold. First, we propose and analyze the placement of multiple content delivery cloudlets (CDC), based on user population densities within a large urban communications network, to bring content closer to consumers. Second, smarter CDC storage techniques are proposed to leverage content-centric in-network caching principles based on content popularity for efficient and timely deliveries. Contributions of this work are as follows:

- Performance analysis of the proposed shift of LinkNYC’s infrastructure from a traditional communications network to a content-centric network.
- Designs and evaluates a CDC placement heuristic which exploits NYC population densities acquired from the 2012 LandScanTM dataset.
- Establishes that large urban communication networks vastly improve with not only a content-centric shift but also the placement of multiple content delivery cloudlets geographically distributed throughout the network.
- Designs and evaluates CDC storage techniques for timely

This work was supported in part by the US National Science Foundation (NSF) under NSF award CNS-1162296.

¹World population data sheet: <http://www.prb.org/>

and responsive content delivery within large networks.

- To our knowledge, this work is the first to analyze and leverage LinkNYC’s urban communications network through content delivery cloudlets improving overall network performance and stability as well as mobile content delivery and service continuity.

The rest of this paper is organized as follows. In Section II content-centric urban communication networks and the LinkNYC infrastructure are introduced, constructed and analyzed. Next, a CDC placement heuristic is proposed in Section III. Section IV introduces potential content storage solutions for responsive content delivery. Proposed solutions are then analyzed and validated using LinkNYC in Section V. Finally, the article is concluded in Section VI.

II. CONTENT-CENTRIC URBAN COMMUNICATION NETWORKS

Coupling urban communication networks with content-centric and delivery principles greatly benefit content producers, consumers and the cities they reside in. Improving their infrastructure using practical approaches to provide more reliable and responsive communications can assist in the technology’s overall success. The underlying concept behind content-centric communication networks is to allow a consumer to focus on the desired named content rather than referencing the physical location or named hosts (IP) where that content is stored [8]–[10]. This shift is a product of empirical research resulting in the fact that the vast majority of Internet usage involves data being disseminated from a source to multiple users. The potential benefits of a content-centric adoption include in-network caching to reduce congestion, improved delivery speeds, simpler network configuration and network security at the data level [11]. This paper combines content-centric and content delivery principles to improve the performance and reliability of urban networks such as LinkNYC.

A. Related Works

In [12]–[16] content caching in content-centric networks (CCN) is considered to exploit in-network caching and improve end user’s QoE. Different caching policies have been proposed which can be classified into three groups: autonomous, centralized, and cooperative [16]. Some research focuses on exploiting popularity to enhance caching in CCN with the aim of reducing the amount of cached items while maintaining a high cache hit rate [14], [17]. Other works [18], [19] focus on virtual network function placement approaches while accounting for mobility [18] and service [19] aspects. In [20], [21] social information is exploited by identifying influential users and proactively caching their content in the network. Combining popularity with cooperation has also been considered [15], [22]–[25]. However, a major issue with these approaches is the communication overhead due to added information exchange and signaling, especially in large networks, as well as requiring many network parameters to be optimal. For this reason autonomous caching techniques are often favored. Our work proposes a framework that leverages (i) content delivery techniques to bring content closer to



Figure 1: Microsoft Azure CDN point of presence locations

users within large urban communication networks and (ii) content-centric principles for cooperative in-network caching to improve overall network performance and user QoE.

B. Content Delivery Networks

The principle behind content delivery networks (CDN) is to bring data as close to the geographic location of the user as possible to improve overall network performance. This helps eliminate the need to traverse the Internet for content which reduces infrastructure and bandwidth costs while improving network robustness and quality of experience (QoE). For instance, Microsoft Azure’s content delivery network consists of 36 points of presence locations² distributed throughout the world as shown in Figure 1. However, in addition to being geographically limited, CDN nodes are generally placed at the Internet edges over multiple backbones servicing different regions remotely. Although the concept of bringing content closer to the consumer through caching copies in various geographic locations improves overall network performance; mobile users in large urban networks naturally endure additional latency due to increased mobility, congestion and hops traversed within the network. This can be improved by placing content even closer to the requesting consumer through content-centric networking and delivery principles.

C. Content-Centric Networks

Different Internet architectures have been introduced for content-centric networking which shifts from the standard IP data packets to Named-Data Networking (NDN). NDN is a future Internet architecture focusing on a content-centric Internet as opposed to today’s host-centric network architecture [11], [26], [27]. These architectures rely on named-content within the Internet to route and direct the flow of data within a network. In these architectures the content is the focus rather than the physical location where the content is stored. NDN consists of content consumers and producers. Consumers generate interest requests for specific content chunks whereas producers respond to interest requests with data packets. Every node is equipped with a Content Store (CS), Pending Interest Table (PIT), and Forwarding Information Base (FIB). When an interest packet is received it is first placed into the PIT and the CS is checked for data correlating to this interest.

²Figure 1 provided by Microsoft Azure: <https://azure.microsoft.com/en-us/documentation/articles/cdn-pop-locations>

If there is a match the interest request is discarded and the corresponding data packet from the CS is returned. Otherwise the interest packet is forwarded based on information in the FIB. Incoming data packets corresponding to pending interests in the PIT are stored in the node's CS. Otherwise, the data packets are dropped. These concepts are used in designing our responsive content delivery framework for large urban networks. Before presenting our solutions, we first introduce and construct the LinkNYC network which is the basis behind our design choices.

D. LinkNYC Framework

In November 2014, LinkNYC announced a project plan to provide a first-of-its-kind communications network offering super fast free gigabit Wi-Fi to everyone in New York City (NYC) through the replacement of thousands of payphones with kiosk-like structures called **Links**. Once completed, LinkNYC will be the largest and fastest free public Wi-Fi network in the world. The **Links** are designed as an update to the standard phone booth and act as Wi-Fi hotspots while also providing basic services such as advertisements, free phone calls, device charging, touchscreen for Internet browsing to access city services, maps and directions. Revenue generated by the **Links**, through kiosk Ads that are displayed on 55 inch displays, is used to maintain the LinkNYC infrastructure. Each Link is equipped with 802.11ac Wi-Fi technology yielding real world download and upload speeds of 300 and 320 Mbps respectively with an average latency of 5 ms and coverage area of up to 45 meters depending on location [28]. This promising self-contained urban communications network provides cities with revenue and analytics while offering consumers free, continuous and reliable connectivity. Our proposed framework is designed for large urban communication networks and tested on a virtual construction of the LinkNYC network using their payphone locations, NYC population densities and hardware expectations based on specifications provided to the public by LinkNYC.

E. Borough Analysis and Topology Construction

Naturally, the number of potential CDCs depends on the number of currently installed payphone locations³ summarized in Table I. Manhattan is the most dense of the five boroughs with 3,409 payphones and an average distance between them of 43 meters. On the other hand, Staten Island is the most sparse with only 51 payphones and an average distance of 606 meters.

LinkNYC cloudlets are categorized based on their borough as shown in Figure 2. Since the connectivity of NYC's payphone backhaul is unknown, we assume that cloudlets are physically connected (e.g., by fiber optic cables) to their nearest neighbors. Given a particular NYC borough (e.g. Brooklyn), we construct an Euclidean minimum spanning tree (EMST) as shown in Figure 5b using Prim's algorithm [29], [30] where edge weights are equal to the geographic distance between cloudlets. EMSTs are useful for telecommunications

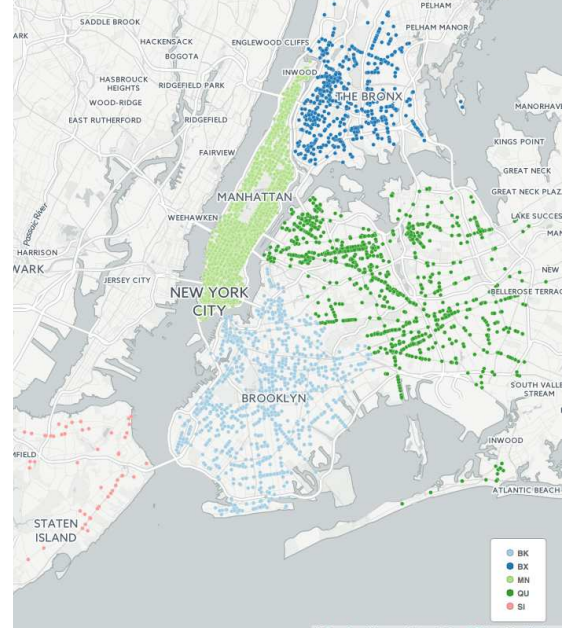


Figure 2: Geographic locations of payphones in NYC

companies to decide for e.g. where to deploy fiber optic cables considering longer cables are more costly. This results in a network topology with $N - 1$ edges where N is the number of cloudlets in a particular borough. This constructed LinkNYC infrastructure will be used later to validate our proposed solutions.

In order to leverage content-centric and content delivery networking with large urban communications networks such as LinkNYC's infrastructure, this work proposes an architectural addition of edge cloud devices referred to as content-delivery cloudlets (CDCs) to better service end users. CDCs are small-scale cloud datacenters that aim to bring data closer to mobile users and are typically located near the edge of the Internet [31]–[35]. Upgrading, for example, LinkNYC's **Links** with these CDCs offers a different dimension to content delivery networking. Unlike traditional content delivery networks, where a limited number of remote servers or nodes are distributed throughout the world, the proposed approach selects and designates a subset of LinkNYC's **links** as CDCs (i.e., producers of content). These selected CDCs are each equipped with an $L2$ storage cache that is much larger than the $L1$ storage cache available on other cloudlets. Content is locally cached as it traverses the network to its destination. For this paper, we assume storage capabilities are limited to CDCs only (i.e. $L2$ cache only) as to avoid excessive deployment costs. In order to decide which CDCs will be selected for providing content delivery services a hierarchical clustering technique is applied to the borough topologies based on NYC's population density distribution detailed in the following section.

III. CONTENT DELIVERY CLOUDLET SELECTION AND PLACEMENT

As consumers become increasingly mobile with dynamic content interests the placement of CDCs in large cities be-

³NYC Open Data: <https://nycopendata.socrata.com>




Figure 3: LinkNYC network construction using minimum spanning trees

comes both crucial and challenging. Specifically, mobile users that encounter multiple cloudlets across a path and undergo frequent handoffs experience service quality continuity issues [36], [37]. With added dynamic content interests, if CDCs' placement is not carefully designed, these issues may be exacerbated. Relying on a single CDC is insufficient to meet the demand of dynamic consumers, and thus, having content readily available in multiple nearby CDCs is indispensable to ensure responsive content delivery and maintain good QoE. This avoids the additional costs of traversing the entire network or querying the original publisher for the content. In what follows, we investigate clustering approaches to efficiently select and decide on the placement of multiple CDCs to enable content-centric networking and delivery in smart cities while considering the LinkNYC network as our use-case for applying and evaluating such approaches.

A. Population Distribution

In order to reflect realistic applications in practice a complete content-centric LinkNYC framework is designed based on NYC's population density distribution using the United States Department of Energy's ORNL (Oak Ridge National Laboratory) LandScanTM 2012 dataset. LandScanTM uses spatial data, satellite imagery analysis and a multi-variable dasy-metric modeling approach to disaggregate census counts to provide accurate ambient population per square kilometer over a 24 hour period. This results in an average population density during the daytime. Nighttime population densities are acquired using the Socioeconomic Data and Applications Center (SEDAC) Gridded Population of the World (GPW) dataset which uses census and satellite data. Average population densities of LinkNYC are illustrated in Figure 4; Note, NYC experiences an influx of 4.5 million people during the day which is mostly concentrated in the Manhattan borough. The datasets were then disaggregated to obtain constant population density estimates immediately surrounding individual LinkNYC cloudlets. Given the population estimates a clustering heuristic is applied to each borough for CDC placement.

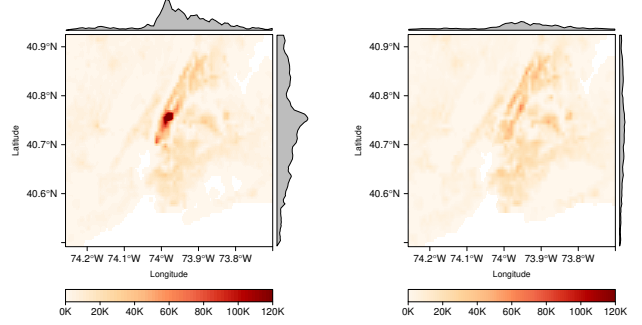


Figure 4: Daytime versus nighttime population densities

B. CDC Hierarchical Clustering Heuristic

Cluster analysis is an NP-hard problem, thus, efficient heuristic algorithms are commonly employed that converge to a local optimum depending on the application. For this purpose, we propose a population density based clustering heuristic for large urban communication networks and CDC placement. Once CDC placement is decided it is assumed to be static. That is, CDCs are upgraded **Links** with enhanced hardware. Thus, frequently changing CDC locations is not practical in a dynamic environment.

Initially, cloudlets are part of the same membership and form a single community. In order to introduce content delivery services to the borough network, a cloudlet is chosen as a CDC based on its average hop count to the remaining cloudlets within its respective community. The probability that requests are initiated from each cloudlet, i , is assumed to be proportional to its respective surrounding population density, γ_i , and is defined as $r_i = \frac{\gamma_i}{\sum_{j=1}^{N_i} \gamma_j}$, where N_i is the number of Links in the entire borough network (e.g., Brooklyn). The population densities of all LinkNYC's Links— γ_i for Link i —are estimated as described in Section III-A. Let's now denote by S the shortest path matrix that contains the length of the shortest path to and from each Link in the borough network. Given the cloudlet request probability vector, $\mathbf{r} = (r_1, r_2, \dots, r_{N_i})$, a weighted average shortest path vector, $\bar{s} = S \cdot \mathbf{r}$, is computed per community and whose entry values represent the weighted average hop count provided by each cloudlet to and from each other. Then, the cloudlet with the minimum sum of weighted average hop counts to remaining cloudlets is selected as the CDC, i.e., $\arg \min \bar{s}$. This ensures content is placed as close to the geographic location of potential consumers within

a community as possible. Once selected, the incident edge between the CDC and the cloudlet with the minimum average hop count is removed forming two disjoint communities. Then, for each community, a single cloudlet is selected as a CDC by recomputing vector \bar{s} . This results in a set of \bar{s} -vectors where the community experiencing a higher average hop count is chosen for additional CDC placement, i.e., $\arg \max (\min \bar{s}^{(1)}, \min \bar{s}^{(2)}, \dots, \min \bar{s}^{(n)})$ where n represents the current number of communities. Edges are progressively removed from the original topology until no edges remain, resulting thus in a set of communities equal to the number of cloudlets.

Borough	# Payphones	Avg. distance	CDCs
Manhattan	3409	43.2 m	50
Queens	1042	136.8 m	25
Brooklyn	1004	150.8 m	25
Bronx	591	125.5 m	20
Staten Island	51	606 m	10
Total	6097	212.5 m	130


Table I: Payphones in the five boroughs of NYC

The intuition behind this heuristic is to assign more CDCs to highly populated areas. Traditionally, large urban networks rely on a regional CDN (i.e., CDCs = 0) for content caching which provides inadequate responsiveness as content would not only still need to be fetched from the original publisher but also traverse the large urban network incurring additional latencies. Incorporating in-network caching through the addition of a single CDC immediately improves latency by storing content closer to users. However, we can see from Figure 5a that a single CDC still does not facilitate timely content delivery as the average latency to access the CDC is high. To determine an adequate number and initial placement of borough CDCs during the daytime, the elbow in the CDC curve in Figure 5a is estimated and used to decide on the number of CDCs to be chosen and are summarized in Table I. Figure 5b illustrates the resulting communities and CDCs for Brooklyn's topology (i.e., CDCs = 25).

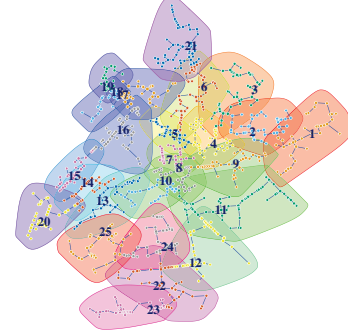
This analysis shows that large urban networks, such as LinkNYC, can benefit greatly with not only the adoption of a content-centric infrastructure (i.e., single CDC) but also with the incorporation of multiple content delivery cloudlets within its urban communications network. Overall a cluster-based placement approach exhibits promising and more stable results yielding lower hop counts and achieving on average a higher rate of cache hits which in turn improves content delivery speeds. Next, we propose content popularity-driven and cooperative caching techniques deployed by CDCs for improved infrastructure performance.

IV. POPULARITY-DRIVEN AND COOPERATIVE CONTENT-CENTRIC CACHING

Internet content of interest to users far surpasses the storage capacity of the CDCs, thereby resulting in frequent content requests from the original content publisher. When content is not stored at intermediate cloudlets, it must be fetched from a remote datacenter incurring additional latency and service



(a) Average hop count as a function of CDCs



(b) Clustering of Brooklyn's LinkNYC network

Figure 5: Borough latency and CDC placement

continuity issues [38]. Naturally, by increasing cache sizes or the number of CDCs, cloudlets will be able to store and push content closer to end users, but this cannot be done without additional hardware and cache deployment costs. Depending on their needs and resource availability, city officials, policy makers, and network administrators manage to find the balance between acceptable QoE and deployment costs. In this work, we opt for techniques that can improve network performance and QoE while avoiding hardware costs. Hereinafter *cloudlet* and *CDC* will be used interchangeably.

A. Popularity-Driven Content Caching

Traditionally, caching consists of fetching content upon request and storing it locally based on some cache replacement technique, such as First-In-First-Out (FIFO), Random Replacement (RR), Most Recently Used (MRU), Least Recently Used (LRU), Least Frequently Used (LFU), etc [39], in anticipation of future content requests. Current cache information (i.e., content, timestamps, queue, etc) is then used to make replacement decisions (i.e., evicting least used content). LRU is among the most traditional reactive cache replacement policies and is still used in practice today [40]. These traditional solutions, however, are not suitable for Internet content delivery in highly populated cities, such as NYC, mainly due to the diversity, heterogeneity, volume, and dynamic nature of Internet content. In what follows from this section, we propose a content-centric LFU caching method suitable for these highly populated cities that incorporates and relies on the popularity of content encountered by CDCs to decide on which content to be cached.

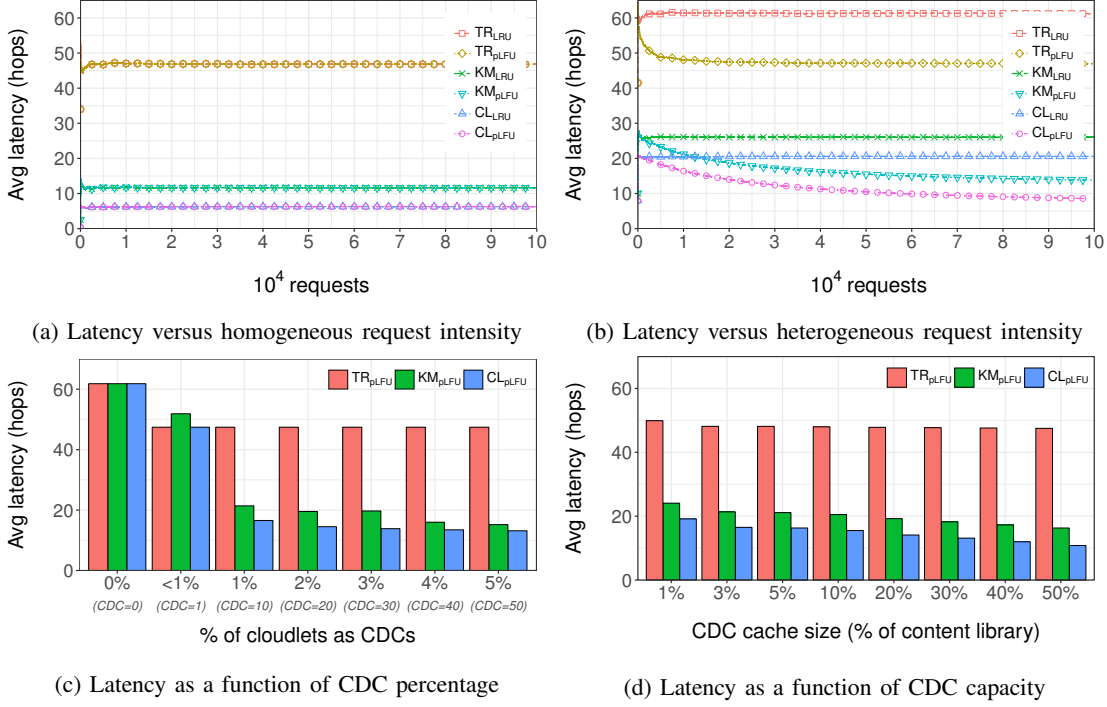


Figure 6: **(a)** and **(b)** Content-centric LinkNYC performance results comparing LRU to pLFU considering homogeneous and heterogeneous content request intensity and using different clustering techniques. **(c)** Average latency as the number of CDCs increases. **(d)** Average latency as the capacity of CDCs increases.

Our proposed popularity-based content caching method, termed pLFU, works as follows. Each cloudlet computes and maintains an estimate of the average number, $\bar{c}_f^{(\kappa)}$, of each content f 's requests encountered by the cloudlet at the κ^{th} period. We propose to compute $\bar{c}_f^{(\kappa)}$ as a weighted moving average $\bar{c}_f^{(\kappa)} = \alpha \bar{c}_f^{(\kappa-1)} + (1-\alpha) c_{f,w}^{(\kappa)}$ where $c_{f,w}^{(\kappa)}$ is the number of requests for content f encountered at the κ^{th} period during window w , and α is a weighting design parameter set between 0 and 1. The parameter κ refers to the current index of the content observation window w . In practical scenarios content popularity mainly depends on current social trends and time of day. In our work, however, content popularity is assumed to follow the Zipf distribution. Zipf parameters are periodically altered to help mimic changing popularities over time. CDC caching policies in large urban communication networks must be able to quickly adapt to these changes. For this reason a small observation window w (i.e., 10) is needed to quickly detect changes in content popularity over time. The *popularity index* of content f is then computed as $p_f^{(\kappa)} = \frac{\bar{c}_f^{(\kappa)}}{\sum_{g \in \mathcal{F}} \bar{c}_g^{(\kappa)}}$ where \mathcal{F} is the set of contents the cloudlet has encountered. Subject to cache sizes, CDCs store encountered content with the highest popularity indexes.

Figure 6 illustrates the benefit of a content-centric shift within Brooklyn's LinkNYC network. For this preliminary analysis we leverage Brooklyn's 25 CDCs, from Table I, where we assume CDCs are capable of storing up to 3% of the network's total popular content. We also assume the latency required to fetch content from the original content publisher is minimal. Figures 6a and 6b depict the average

latency experienced of LRU and pLFU caching policies while considering content homogeneity and heterogeneity respectively. Although content homogeneity, where all users request the **same** content through their respective CDCs, is not a realistic case, it does highlight the performance benefit of our population based clustering and CDC placement heuristic from Section III-B, *CL*, compared to traditional, *TR*, and k-means, *KM*, approaches. Since content requests are the same across CDCs, pLFU does not provide any added benefit over LRU. However, when more realistic heterogeneous content requests are introduced we see the performance boost of a pLFU approach. A traditional single regional CDC employing pLFU (TR_{pLFU}) provides a near 25% reduction in average latency per request as opposed to TR_{LRU} . Introducing in-network caching to LinkNYC provides great improvements to a traditional infrastructure, however, the average latency remains high (47.5 hops) and is not ideal in dynamically changing environments resulting in sub-par user QoE. Incorporating multiple CDCs within LinkNYC addresses this issue by pushing content closer to the edge. Both KM_{LRU} and CL_{LRU} provide immediate latency reductions of 58% and 67% compared to TR_{LRU} . Although CDCs have limited capacities, content popularity can be learned in anticipation of future requests. Over time, and as content requests intensify, KM_{pLFU} and CL_{pLFU} exhibit promising results yielding a 65% and 85% latency reduction compared to TR_{pLFU} . Using the pLFU policy, Figure 6c illustrates average latency as the number of CDCs deployed within Brooklyn's topology is increased and a content request intensity of 20×10^3 . Naturally, latency improves as more CDCs are deployed, however, the

improvement is minimal as the percentage of CDCs reaches around 5%. The effect of increasing CDC cache sizes is shown in Figure 6d. Results show that increasing cache sizes alone may not necessarily improve latency since contents still need to traverse to end users. We promote that there exists a balance between the number of CDCs and a CDC's cache size.

In Sections III and IV-A, a cloudlet-based content placement architecture coupled with a popularity-driven content caching scheme that stores content based on local CDC popularity indexes was proposed. It is shown that integrating content popularity when making content placement and caching decisions reduces both the average downloading time and the network backhaul traffic as opposed to pure reactive policies. Although the popularity-driven caching scheme, pLFU, yields promising results, it does not guarantee content placement that minimizes the global average latency, as it does not account for content availability and popularity in neighboring CDCs. An amalgam of user mobility, content heterogeneity, and in-network storage introduces several complex challenges. More advanced content-centric solutions must be sought in order to reduce network issues and address underlying limitations such as rising content heterogeneity, user mobility, backhaul load, network failure and deployment costs. Next, we propose a potential solution for overcoming these challenges. Specifically, we propose to exploit and rely on cooperation and information sharing among neighboring CDCs for providing faster content access and lesser backhaul traffic.

B. Cloudlet Cooperation for Faster Content Access

Content caching and placement decisions should depend not only on local but also neighboring CDC conditions and observations, such as content popularity, storage capacity, content availability in the neighborhood, user population, and link/network condition (congestion, data rates, etc.). Intuitively, when a new content is requested within some local CDC, the decisions on (i) whether to cache the new content or not, (ii) which CDC to cache the content at, and (iii) which existing cache content to evict should all be community-based in that both local and neighboring CDCs should all cooperate and be involved in making such decisions so that globally optimal placements that satisfy users' QoE while accounting for resource availability constraints can be made. For example, if the new content is available at a nearby CDC, then there might not be a need for caching it again at the local CDC, thus saving local cache resources yet while still allowing users to receive the content with acceptable latency. Now if the new content is not available locally, nor in neighboring CDCs, then the decision to whether to cache or not should depend on its community popularity, not just its local popularity. If this content is popular enough to cache, then the decision to where it should be cached should weigh in its popularity indexes at the different CDCs within the network. That is, even if the content is being requested by a user located within a CDC i , it might be more efficient to cache it at a neighboring CDC if popularity indexes show that the content is more likely to be requested within the neighboring CDC and/or the neighboring CDC has more available cache space.

Designing cooperative content caching and placement approaches that consider the aforementioned performance aspects has not been fully addressed by the research community. And deriving models that capture the various content and network aspects influencing these decisions, such as content popularity, storage availability, content availability, user population, and network condition, is a challenging task that requires careful study. In this work, we propose to introduce a score function $\mathcal{S}_{f,i}^{(\kappa)}$ that each CDC i maintains for each of its encountered content f , updated every period κ , and used to make content placement and caching decisions. Here, $\mathcal{S}_{f,i}^{(\kappa)}$ represents the cost associated with caching content f at CDC i at the update period κ . We propose that this function captures and models the following aspects:

- **Content popularity ($p_{f,i}^{(\kappa)}$):** Popularity index of content f as observed by CDC i during update window κ .
- **Content availability ($a_{f,i}^{(\kappa)}$):** An availability binary index of content f , where 1 indicates that f is cached in CDC i during update window κ , and 0 otherwise.
- **Population density (r_i):** This reflects the population density of CDC i , as described in Section III-A.
- **Inter-CDC delay ($\ell_{i,j}$):** It represents the delay experienced by a user belonging to a CDC i requesting content cached at a neighboring CDC j . It essentially captures the number of hops, as well as the link bandwidth capacity of each hop, connecting CDCs i and j .
- **Intra-CDC delay (ℓ_i):** It represents the average delay experienced by a user requesting content from its community CDC i .

We propose to model $\mathcal{S}_{f,i}^{(\kappa)}$ as a weighted average of a neighborhood score and a local score as such:

$$\mathcal{S}_{f,i}^{(\kappa)} = \beta \times \mathcal{S}_{neigh f,i}^{(\kappa)} + (1 - \beta) \times \mathcal{S}_{local f,i}^{(\kappa)}$$

where $\mathcal{S}_{neigh f,i}^{(\kappa)} = \frac{\sum_{j \in \mathcal{N}_i} r_j p_{f,j}^{(\kappa)} \ell_{i,j}^{(\kappa)}}{|\mathcal{N}_i|}$ and $\mathcal{S}_{local f,i}^{(\kappa)} = \frac{r_i p_{f,i}^{(\kappa)}}{\ell_i^{(\kappa)}}$ where \mathcal{N}_i is the set of CDC i 's neighboring CDCs and β is a design parameter set between 0 and 1. $\mathcal{S}_{f,i}^{(\kappa)}$ captures the local and neighborhood benefit, from the perspective of the deciding CDC i , of caching content f . That is, as content traverses the network each CDC decides to store the traversing content based on its score value. Considering local and neighborhood characteristics allows for our score function to better service dynamic content environments. Deciding on the number of neighboring CDCs to consider is a design choice discussed in Section V.

Note that additional score functions can be modeled to capture CDC resources (processing, storage, memory, energy, etc.) more accurately and are left for future investigation. The neighborhood score function, $\mathcal{S}_{neigh f,i}^{(\kappa)}$, essentially represents a weighted average latency that user requests generated within CDC i will experience if content f is cached among its neighboring CDCs. $\mathcal{S}_{local f,i}^{(\kappa)}$ represents an average latency if content f is stored locally in CDC i . For the extreme case when $\beta = 1$, content resulting in higher potential latencies among our neighbors will be locally stored more often. On the other hand, when $\beta = 0$, our score function stores content with the highest local popularities and low latencies. That is,




Figure 7: Popularity indexes as a function of s

highly popular content with low latencies to fetch are favored over others.

C. Content Popularity Skewness

Naturally, content popularity within communities can vary over time and in order to account for this dynamicity, tuning the design parameter β , in $\mathcal{S}_{f,i}^{(\kappa)}$, becomes especially useful and can provide a more responsive storage mechanism. We assume content popularity follows the Zipf distribution, $f(\tau, s, M) = \frac{1/\tau^s}{\sum_{m=1}^M (1/m^s)}$, where τ is a content's rank in terms of popularity, M is the total number of contents, and s controls the skewness of the distribution. The Zipf distribution is widely used in the literature to describe popularity [41]–[43].

It is clear from Figure 7 that as parameter s increases so does the skewness in content popularity especially when $s \geq 1$. We propose a dynamic parameter β , in the score function $\mathcal{S}_{f,i}^{(\kappa)}$, that is inversely proportional to parameter s . That is, for the extreme case when contents have equal popularities (i.e., $s = 0$) simply storing the most popular content locally, as done with pLFU, will not suffice. However, querying content popularities and availabilities among neighboring CDCs for storage decisions intuitively improves responsiveness by effectively consolidating storage capabilities. Thus, a higher β value is preferred to favor our neighborhood in our cache decisions. Conversely, when content skewness is high (i.e., $s = 2$), it is more efficient to locally store the most popular content within a community. In other words, as content popularities are more skewed our score function favors storing the most popular content locally. For this reason a dynamic β parameter is necessary to balance between local and neighborhood based storage decisions depending on the s parameter in the Zipf distribution. This allows $\mathcal{S}_{f,i}^{(\kappa)}$ to be adjusted accordingly and adapt to network changes over time.

D. Cooperative Content Placement

The parameter β allows us to balance between the need for having responsive content delivery locally versus among our CDC neighbors. This is especially useful for dynamic content popularity within large urban communication networks. As these above network and content conditions change over time, each CDC must periodically maintain and compute score function values for encountered contents. This could be

done by having CDCs query neighboring CDCs for content popularity indexes, population densities, CDC latency and resource availability, and use this information for updating these values, which are then used as follows for content placement decisions:

- Content popularity indexes, $p_{f,i}^{(\kappa)}$, are maintained for each encountered content f during period κ as described in Section IV-A
- For each content f encountered at CDC i , the $\mathcal{S}_{neigh f,i}^{(\kappa)}$ and $\mathcal{S}_{local f,i}^{(\kappa)}$ portions of the content score value, $\mathcal{S}_{f,i}^{(\kappa)}$, are computed and maintained during period k
- For each CDC j in our set of neighbors $\in \mathcal{N}$, inter-CDC latency, $\ell_{i,j}^{(\kappa)}$, is computed based on content f 's availability among our neighbors. Note, the additional latency required to fetch content from the original publisher is also captured:

$$\ell_{i,j}^{(\kappa)} = \begin{cases} \ell_{i,j} + \ell_{j,origin}^{(\kappa)} & a_{f,i}^{(\kappa)} = 0 \\ \ell_{i,j} & a_{f,i}^{(\kappa)} = 1 \end{cases}$$

- Intra-CDC latency, $\ell_i^{(\kappa)}$, is computed during period κ as the average latency (i.e, hops) from each cloudlet to its community CDC i . Similarly, in the event that $a_{f,i}^{(\kappa)} = 0$, $\ell_{j,origin}^{(\kappa)}$ is accounted for to capture the additional latency required to fetch content from the original publisher
- Design parameter β is set based on an estimation of the s during period κ and maintained per CDC
- For each content f encountered at CDC i , it is stored locally if $f \in \mathcal{S}_{max}$ where \mathcal{S}_{max} is the set of contents with the highest local content scores, subject to CDC storage capacity. Otherwise no caching takes place. If CDC i 's storage is full and $f \in \mathcal{S}_{max}$ then the content with the minimum score value is evicted from the local cache

Considering both local and neighborhood conditions and content availability in modeling $\mathcal{S}_{f,i}^{(\kappa)}$ allows us to account for changing content popularities within different communities over time. Our intuition indicates that storage decisions must rely heavily on a CDC's content popularity distribution. That is, as content popularity becomes more uniform, considering content availability among our neighbors improves performance. However, as content popularity is more skewed, storage decisions should be based on local CDC content popularities. We discuss our design choices and analyze our results next.

V. FRAMEWORK ANALYSIS: A LINKNYC USE-CASE

A detailed examination using LinkNYC's urban communications network is performed to investigate the fundamental benefits of shifting from a traditional IP network to a content-centric network where popular content is cached as it disseminates throughout the network. ndnSIM is an NDN simulator based on NS-3 and was used to analyze the proposed urban communications network infrastructure. Traditionally, content delivery nodes are distributed globally to bring content as close to the geographic location of the user. However, promising results show that having multiple content-delivery cloudlets

Parameter	Value
Borough	Brooklyn
Cloudlets	1004
CDCs	25
Content Categories	Sports, Education, Politics, Movies
$ M $	600
C_i	20
Number of requests	10^6
ℓ_{origin}	[250, 500] hops
$ \mathcal{N}_i $	24
s range	[0, 2]
\hat{s} observations	10^3
w	10^2
α	.2
β	$1 - \frac{1}{1 + e^{-20(s-.5)}}$

Table II: Simulation parameters


within LinkNYC’s infrastructure coupled with smarter caching techniques dramatically improves overall network performance and stability. In an increasingly mobile and dynamic environment, where network conditions and service continuity issues are amplified due to user mobility and varying content interests; network stability, efficiency and content delivery responsiveness are essential to a user’s QoE.

A. Setup

In this analysis we focus on NYC’s Brooklyn borough. Among the topology’s 1004 cloudlets, 25 CDCs are geographically selected based on population densities using our technique from Section IV. Content requests are then generated following a Zipf distribution with varying s parameters per community to realistically simulate community-specific interests and changing content popularity. We assume LinkNYC’s popular content library, M , contains 600 contents categorized by type (i.e., Sports, Entertainment, Politics and Education) where each CDC’s storage capacity, C_i , is set to 20 contents. Content sizes are assumed to be constant where popularity indexes follow the Zipf distribution. However, fluctuating population densities and varying link capacities are currently in progress and have been left for future work. Requesting content from its original publisher also yields varying latencies per community (i.e., between 250 and 500 hops). In addition, popularity skewness and interests are periodically shuffled among communities to capture time variation within the LinkNYC infrastructure. Finally, design parameters α and β are chosen to control (i) the framework’s sensitivity to popularity changes and (ii) the degree to which we balance between neighborhood and local content storage decisions respectively. Our simulation parameters are summarized in Table II.

B. CDC Neighborhood Size

Neighborhood size can play an important role in overall network performance. Increasing the size of CDC i ’s neighborhood, \mathcal{N}_i , improves performance by effectively consolidating



(a) Latency as a function of sLFU neighborhood size


(b) Latency as a function of different s parameters

Figure 8: (a) sLFU performance as CDC neighborhood size is increased. (b) Comparing traditional cache replacement policies with pLFU and sLFU considering static content popularity skewness.

CDC capabilities. Figure 8a emphasizes this performance increase with respect to latency as the neighborhood size is grown to 24. This allows for efficient cooperation among CDCs when making content storage decisions. Since the number of CDCs are relatively low, the negative impact of increasing neighborhood size is often negligible. However, in networks comprising of hundreds of thousands or even millions of devices, such as Massive Internet of Things (MIoT) networks, increasing the neighborhood size can be disadvantageous in terms of complexity due to the added overhead injected into the infrastructure through the exchange of content stores and popularity indexes [44]. LinkNYC network administrators should decide on appropriate neighborhood sizes depending on the application. For instance, during high profile sporting events or tense political climates, many contents may be popular throughout the entire NYC borough and thus using a neighborhood size equal to the total number of CDCs may be preferred. Conversely, geographically restricted interests such as social events may cause only one or two communities to share common interests and thus smaller or even dynamic neighborhood sizes may be beneficial to the overall performance of the network.

C. Impact of Content Popularity

Figure 8b illustrates the effect content popularity has within LinkNYC. We compare our proposed techniques, pLFU and sLFU, with traditional LRU and LFU cache replacement policies. Here we assign the same s value to all Brooklyn communities to emphasize the importance of tuning the design

parameter β in $\mathcal{S}_{f,i}^{(\kappa)}$ based on the content distribution. With s values close to zero, β is adjusted accordingly allowing for sLFU to cooperate more with neighboring CDCs to make storage decisions, drastically reducing latency by up to 33% compared to LFU and pLFU, and up to 43% compared to LRU. However, as we increase s , and in turn content popularity skewness, sLFU appropriately relies more on local CDC popularity indexes for storage decisions. That is, when content popularity is very skewed (i.e. $s \geq 2$), it is sufficient to only store the most popular content locally; hence, traditional methods also perform well.

D. Content Dynamicity and Estimation

Figure 9 introduces content dynamicity within our content-centric LinkNYC network. Content popularity skewness and interests are periodically shuffled among communities to capture time and content variation. This depicts shifting user interests per community over time (i.e. different sporting events, changing political climates, etc.). During period κ , where w contents have been encountered, popularity indexes can be leveraged to estimate the s parameter of the content distributions. Each CDC i can then incorporate its estimated s value for tuning design parameter β of the score function $\mathcal{S}_{f,i}^{(\kappa)}$. We use maximum likelihood estimation (MLE) to find s based on our current content observations (i.e. 10^3 observations) allowing sLFU to adapt to changing popularities over time.

Figures 9a and 9b depict LinkNYC performance as content requests intensify and while popularity skewness is relatively low. The LRU cache replacement policy is the worst performer and results in high average latency, over 200 hops, that is not suitable for a dynamic and mobile environment. LFU and pLFU perform similarly, however, pLFU is able to converge quicker to an average latency of nearly 180 hops. sLFU's and sLFU's cooperation with neighboring CDCs allows them to perform noticeably better achieving over a 30% improvement compared to traditional methods even in less than ideal conditions where content popularity is relatively similar across communities. When maximum content popularity skewness is increased (i.e. $s = 2$), as shown in Figure 9c, traditional methods perform better but are still no match for sLFU and sLFU which still provide a 30% reduction in average latency. Figure 9d illustrates the cache hit percentages when different s parameters are used. sLFU consistently achieves a high cache hit percentage than traditional methods constantly achieving a near 80% hit rate even with low s values. sLFU's cooperation-based technique is a novel approach for content placement in large urban communication networks. Implementing sLFU at CDCs provides an efficient and cooperative LinkNYC infrastructure for responsive content delivery in dynamically changing environments; essential for mobile users where timely delivery is of the essence.

E. User Mobility

Mobile users within LinkNYC are bound to experience intermittent connectivity as multiple cloudlets are encountered risking frequent disconnects and disruptions in a user's service especially near the edges of a cloudlet's coverage area [36].

This causes potential packet losses, increased response times and re-transmissions. User mobility and dynamic content interests introduce added complexity to an already challenging issue. In order to accommodate mobile users as well as avoid interfering with popular community content we propose the consolidation of CDC storage with a smaller, temporary $L1$ storage strictly used to store prefetched mobile user content on candidate CDCs located on a mobile user's path. Once a mobile user consumes the prefetched content the $L1$ cache is immediately vacated in preparation for other mobile users. In our previous work, we designed a cloudlet-aware mobility based solution to address service continuity issues as multiple cloudlets are encountered across a path. In [38] a proactive algorithm is designed and applied by a mobile user as it moves within a community of cloudlets to prefetch content chunks. Prefetching content on candidate CDCs or cloudlets across a path has been shown to address mobility induced service continuity issues.

VI. CONCLUSION AND FUTURE DIRECTIONS

In this article, we investigated the fundamental benefits of shifting a large urban communications network to a content-centric network where popular Internet content is cached as it disseminates throughout the network. Different caching techniques were designed and evaluated. Initially, a technique to learn and store the most popular content locally at individual CDCs is proposed. However, results show that performance can be improved through cooperative interaction among neighboring CDCs in order to consolidate storage capabilities. Promising results show that an amalgamation of in-network caching, the placement of multiple CDCs and advanced storage techniques employed within an urban network's infrastructure dramatically improves overall network performance and responsiveness. This type of network unlocks a variety of applications and open challenges. An economical aspect can be introduced where CDC storage space is rented to various businesses to serve popular Internet content or customer services while providing an additional source of revenue for the network. Moreover, LinkNYC can be exploited through mobile cloud computing by augmenting CDC computing capabilities. The focus is still latency minimization while allowing users to offload heavily computational tasks. Our framework lays the foundation for a first-of-its-kind content-centric urban communication network with the focus of minimizing latency and leveraging interest-based content caching through popularity learning and cooperative content placement.

REFERENCES

- [1] B. Hamdaoui and P. Ramanathan, "Cross-layer optimized conditions for QoS support in multi-hop wireless networks with MIMO links," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 4, 2007.
- [2] B. Hamdaoui and K. G. Shin, "Characterization and analysis of multi-hop wireless mimo network throughput," in *Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2007, pp. 120–129.
- [3] W. Su, A. K. Sadek, and K. R. Liu, "Cooperative communication protocols in wireless networks: performance analysis and optimum power allocation," *Wireless Personal Communications*, vol. 44, no. 2, pp. 181–217, 2008.

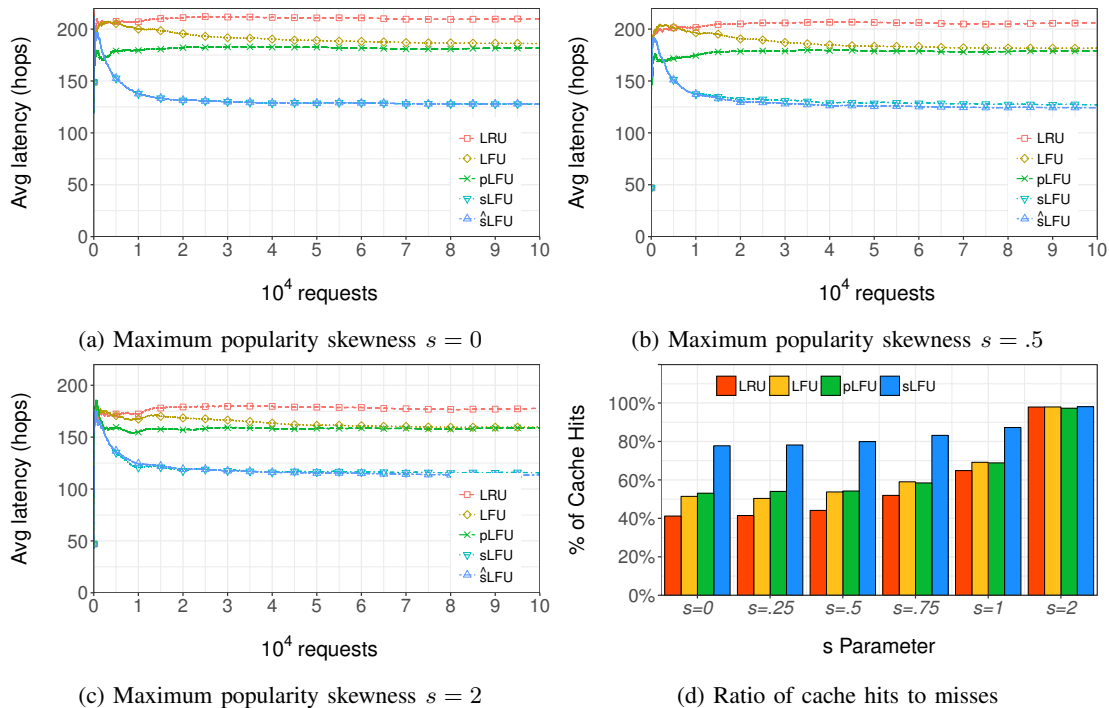


Figure 9: (a), (b), and (c) Illustrate content-centric LinkNYC performance results comparing LRU, LFU, pLFU and sLFU considering a dynamic network with periodically fluctuating content interests and popularities where s is changed per community with a range between 0 and 2. (d) The ratio of cache hits to cache misses with respect to s .

- [4] N. Chakchouk and B. Hamdaoui, "Uplink performance characterization and analysis of two-tier femtocell networks," *Vehicular Technology, IEEE Transactions on*, vol. 61, no. 9, pp. 4057–4068, 2012.
- [5] P. Venkatraman, B. Hamdaoui, and M. Guizani, "Opportunistic bandwidth sharing through reinforcement learning," *Vehicular Technology, IEEE Transactions on*, vol. 59, no. 6, pp. 3148–3153, 2010.
- [6] D. Dechene and A. Shami, "Energy-aware resource allocation strategies for lte uplink with synchronous harq constraints," *IEEE Transactions on Mobile Computing*, no. 99, pp. 1–1, 2014.
- [7] C. V. N. Index, "Global mobile data traffic forecast update, 2016–2021," *Cisco white paper*, 2017.
- [8] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *IEEE Communications Magazine*, vol. 50, no. 7, 2012.
- [9] S. K. Fayazbakhsh, Y. Lin, A. Tootoonchian, A. Ghodsi, T. Koponen, B. Mags, K. Ng, V. Sekar, and S. Shenker, "Less pain, most of the gain: Incrementally deployable icn," in *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, ACM, 2013, pp. 147–158.
- [10] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A survey of information-centric networking research," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 2, pp. 1024–1049, 2014.
- [11] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, k. claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang, "Named data networking," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 66–73, Jul. 2014.
- [12] D. Rossi and G. Rossini, "Caching performance of content centric networks under multi-path routing (and more)," *Relatório técnico, Telecom ParisTech*, 2011.
- [13] S. Retal, M. Bagaa, T. Taleb, and H. Flinck, "Content delivery network slicing: Qoe and cost awareness," in *Proc. IEEE ICC*, 2017.
- [14] C. Bernardini, T. Silverston, and O. Fester, "Mpc: Popularity-based caching strategy for content centric networks," in *Proc. of IEEE International Conference on Communications (ICC)*, 2013, pp. 3619–3623.
- [15] M. Zhang, H. Luo, and H. Zhang, "A survey of caching mechanisms in information-centric networking," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 3, pp. 1473–1499, 2015.
- [16] A. Ioannou and S. Weber, "A survey of caching policies and forwarding mechanisms in information-centric networking," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 2847–2886, 2016.
- [17] K. Thar, T. Z. Oo, C. Pham, S. Ullah, D. H. Lee, and C. S. Hong, "Efficient forwarding and popularity based caching for content centric network," in *Information Networking (ICOIN), 2015 International Conference on*. IEEE, 2015, pp. 330–335.
- [18] T. Taleb, M. Bagaa, and A. Ksentini, "User mobility-aware virtual network function placement for virtual 5g network infrastructure," in *Communications (ICC), 2015 IEEE International Conference on*. IEEE, 2015, pp. 3879–3884.
- [19] M. Bagaa, T. Taleb, and A. Ksentini, "Service-aware network function placement for efficient traffic handling in carrier cloud," in *Wireless Communications and Networking Conference (WCNC), 2014 IEEE*. IEEE, 2014, pp. 2402–2407.
- [20] C. Bernardini, T. Silverston, and O. Fester, "Socially-aware caching strategy for content centric networking," in *Proc. of IEEE Networking Conference*. IEEE, 2014, pp. 1–9.
- [21] C. Yanqing, W. Muqing, Z. Min, and W. Kaili, "Socially-aware noderank-based caching strategy for content-centric networking," in *Proc. of IEEE International Symposium on Wireless Communication Systems (ISWCS)*, 2016, pp. 297–303.
- [22] Z. Ming, M. Xu, and D. Wang, "Age-based cooperative caching in information-centric networking," in *Proc. of IEEE International Conference on Computer Communication and Networks (ICCCN)*, 2014, pp. 1–8.
- [23] S. Wang, J. Bi, and J. Wu, "Collaborative caching based on hash-routing for information-centric networking," in *Proc. of ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, 2013, pp. 535–536.
- [24] J. M. Wang, J. Zhang, and B. Bensaou, "Intra-as cooperative caching for content-centric networks," in *Proc of ACM SIGCOMM workshop on Information-centric networking*. ACM, 2013, pp. 61–66.
- [25] J. Li, H. Wu, B. Liu, J. Lu, Y. Wang, X. Wang, Y. Zhang, and L. Dong, "Popularity-driven coordinated caching in named data networking," in *Proceedings of the eighth ACM/IEEE symposium on Architectures for networking and communications systems*. ACM, 2012, pp. 15–26.
- [26] A. Afanasyev, I. Moiseenko, and L. Zhang, "ndnSIM: NDN simulator for NS-3," NDN, Technical Report NDN-0005, October 2012. [Online]. Available: <http://named-data.net/techreports.html>
- [27] S. Mastorakis, A. Afanasyev, I. Moiseenko, and L. Zhang, "ndnSIM 2.0: A new version of the NDN simulator for NS-3," NDN, Technical Report NDN-0028, January 2015.
- [28] D. Hardawar. (2016, Jan.) Linknyc's free gigabit wifi is here, and it

is glorious. [Online]. Available: <https://www.engadget.com/2016/01/19/linknyc-gigabit-wifi-hands-on/>

- [29] J. C. Gower and G. Ross, "Minimum spanning trees and single linkage cluster analysis," *Applied statistics*, pp. 54–64, 1969.
- [30] S. Manen, M. Guillaumin, and L. Van Gool, "Prime object proposals with randomized prim's algorithm," in *Proceedings of the IEEE International Conference on Computer Vision*, 2013, pp. 2536–2543.
- [31] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," *IEEE Pervasive Computing*, vol. 8, no. 4, pp. 14–23, Oct 2009.
- [32] M. Chen, Y. Qian, J. Chen, K. Hwang, S. Mao, and L. Hu, "Privacy protection and intrusion avoidance for cloudlet-based medical data sharing," *IEEE Transactions on Cloud Computing*, 2016.
- [33] Z. Jiang and S. Mao, "Energy delay tradeoff in cloud offloading for multi-core mobile devices," *IEEE Access*, vol. 3, pp. 2306–2316, 2015.
- [34] M. Chen, Y. Zhang, Y. Li, S. Mao, and V. C. Leung, "Emc: Emotion-aware mobile cloud computing in 5g," *IEEE Network*, vol. 29, no. 2, pp. 32–38, 2015.
- [35] Y. Xu and S. Mao, "A survey of mobile cloud computing for rich media applications," *IEEE Wireless Communications*, vol. 20, no. 3, pp. 46–53, 2013.
- [36] H. Sinky, B. Hamdaoui, and M. Guizani, "Handoff-aware cross-layer assisted multi-path TCP for proactive congestion control in mobile heterogeneous wireless networks," in *GLOBECOM 2015*, pp. 1–7.
- [37] H. Sinky and B. Hamdaoui, "Cross-layer assisted tcp for seamless hand-off in heterogeneous mobile wireless systems," in *Globecom Workshops (GC Wkshps)*, 2013 IEEE. IEEE, 2013, pp. 4982–4987.
- [38] —, "Cloudlet-aware mobile content delivery in wireless urban communication networks," in *Global Communications Conference (GLOBECOM)*, 2016 IEEE. IEEE, 2016, pp. 1–7.
- [39] H. Al-Zoubi, A. Milenkovic, and M. Milenkovic, "Performance evaluation of cache replacement policies for the SPEC CPU2000 benchmark suite," in *Proceedings of the 42nd annual Southeast regional conference*. ACM, 2004, pp. 267–272.
- [40] D. Lee, J. Choi, J.-H. Kim, S. H. Noh, S. L. Min, Y. Cho, and C. S. Kim, "LRFU: a spectrum of policies that subsumes the least recently used and least frequently used policies," *IEEE transactions on Computers*, vol. 50, no. 12, pp. 1352–1361, 2001.
- [41] N. Sarrar, S. Uhlig, A. Feldmann, R. Sherwood, and X. Huang, "Leveraging zipf's law for traffic offloading," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 1, pp. 16–22, Jan. 2012. [Online]. Available: <http://doi.acm.org/10.1145/2096149.2096152>
- [42] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," *SIGCOMM Comput. Commun. Rev.*, vol. 29, no. 4, pp. 251–262, Aug. 1999. [Online]. Available: <http://doi.acm.org/10.1145/316194.316229>
- [43] L. A. Adamic and B. A. Huberman, "Zipf's law and the Internet," *Glottometrics*, vol. 3, pp. 143–150, 2002. [Online]. Available: <http://www.hpl.hp.com/research/idl/papers/ranking/adamicglottometrics.pdf>
- [44] A. Osseiran, J. F. Monserrat, and P. Marsch, *5G Mobile and Wireless Communications Technology*. Cambridge University Press, 2016.



Bassem Khalfi (S'14) is currently a Ph.D. candidate in ECE at Oregon State University. He received his Graduate Engineer Diploma (2012) from the Higher School of Communication of Tunis (SUP'COM) in Ariana, Tunisia, and his M.S. degree (2014) from the National School of Engineers, Tunis, Tunisia. His research focuses on various topics in the area of wireless communication and networks, including dynamic spectrum access and sensing, RF energy harvesting, content-centric networking, and IoT.



Bechir Hamdaoui (S'02-M'05-SM'12) is presently an Associate Professor in the School of EECS at Oregon State University. He received the Diploma of Graduate Engineer (1997) from the National School of Engineers at Tunis, Tunisia. He also received M.S. degrees in both ECE (2002) and CS (2004), and the Ph.D. degree in ECE (2005) all from the University of Wisconsin-Madison. His research interest spans various areas in the fields of computer networking, wireless communications, and mobile computing, with a current focus on distributed optimization,

parallel computing, cognitive networks, cloud computing, and Internet of Things. He has won several awards, including the ICC 2017 Best Paper Award, the 2016 EECS Outstanding Research Award, and the 2009 NSF CAREER Award. He is presently an Associate Editor for IEEE Transactions on Wireless Communications (2013-present). He also served as an Associate Editor for IEEE Transactions on Vehicular Technology (2009-2014), Wireless Communications and Mobile Computing Journal (2009-2016), and for Journal of Computer Systems, Networks, and Communications (2007-2009). He also served as the chair for the 2017 IEEE INFOCOM Demo/Posters program, the 2016 IEEE GLOBECOM Mobile and Wireless Networks symposium, the 2014 IEEE ICC Communications Theory symposium, the 2011 ACM MOBICOM's SRC program, and several other IEEE symposia and workshops, including IWCMC 2009-2017, CTS 2012, and PERCOM 2009. He also served on technical program committees of many IEEE/ACM conferences, including INFOCOM, ICC, and GLOBECOM. He has been selected as a Distinguished Lecturer for the IEEE Communication Society for 2016 and 2017. He is a Senior Member of IEEE, IEEE Computer Society, IEEE Communications Society, and IEEE Vehicular Technology Society.



Ammar Rayes is a Distinguished Engineer at Cisco's Advanced Services Technology Office. His research interests include Network Analytics, IoT, machine learning and security. He has authored three books (latest: <http://www.springer.com/us/book/9783319448589#aboutBook>), over 100 publications in refereed journals and conferences on advances in software & networking related technologies and over 25 patents. He is the Founding President and board member of the International Society of Service Innovation Professionals www.issip.org,

Adjunct Professor at SJSU, Editor-in-Chief of "Advances of Internet of Things" Journal and Editorial Board Member of the European Alliance for Innovation - Industrial Networks and Intelligent Systems. He has served as Associate Editor of ACM Transactions on Internet Technology and Wireless Communications and Mobile Computing Journals, Guest Editor of multiple journals and several IEEE Communication Magazine issues, co-chaired the Frontiers in Service Conference and appeared as Keynote speaker at several IEEE and industry Conferences. At Cisco, Ammar is the founding chair of Cisco Services Research and Cisco Services Patent Council. He received Cisco Chairman's Choice Award for IoT Excellent Innovation & Execution. He received his BS and MS Degrees in EE from the University of Illinois at Urbana and his Ph.D. degree in EE from Washington University in St. Louis, Missouri, where he received the Outstanding Graduate Student Award in Telecommunications.



Hassan Sinky is presently an Assistant Professor in the College of Computer and Information Systems at Umm Al-Qura University in Makkah, Saudi Arabia. He received his M.S. and Ph.D. degrees from Oregon State University. His research interests include large urban wireless communication networks, content-delivery and content-centric networks, quality of service and quality of experience methods, cross-layer assisted multi-path TCP and seamless handoffs in wireless mobile scenarios.