On the Economic Value of Vehicular Caching

Yichen Ruan
Carnegie Mellon University
yichenr@andrew.cmu.edu

Carlee Joe-Wong
Carnegie Mellon University
cjoewong@andrew.cmu.edu

ABSTRACT

The economic value of a new mobile caching method utilizing vehicles is studied. An optimization model is built using stochastic geometry tools. Two possible choices of utility functions are discussed together with some preliminary results.

KEYWORDS

Vehicles, Edge Caching, Heterogeneous Network

1 INTRODUCTION

Due to the widespread use of smart devices, recent years have witnessed the rapid growth in the volume of mobile data traffic. In response, efforts have been made to cache popular contents on edge devices such as small cells or user devices. These contents can then be retrieved directly from edge devices, without passing through the network core. Recently, some works [4] have proposed the use of vehicles as cache carriers and relay nodes. Compared to small cells, vehicular points are expected to have higher density and lower operating cost in an urban environment. Vehicles' mobility also allows them to more flexibly meet the dynamic content demands across time and space. While the locations of small cells are fixed once they are installed, vehicles can dynamically migrate to areas with heavier data traffic. This helps to improve the rate of utilization of the cached contents, thus increasing the value of these caches.

Most existing works on mobile caches focus on the optimization of the contents allocated to the caches. In this work, we instead quantify the economic value of having vehicular caches in the first place, as the cache demands vary over time and space. Our results can help guide Internet service providers (ISPs) that want to relieve the traffic in their network, or content delivery network (CDN) operators that provide caching services to content providers.

The major challenges of the research include 1) Modeling the mobility of vehicles; 2) Quantifying the physical and economic interaction of vehicular caching with other caching methods, e.g., interference with stationary small cell caches; and 3) Solving for the optimal cache provisioning, which is a non-convex problem.

2 SYSTEM MODEL

Suppose there is a competitive caching market comprised of two tiers of caching products: small cell caching and vehicular caching. The buyers (operators) subscribe to a combination of the two products to maximize their utilities. The products are sold in terms of per unit intensity, which is the expected number of devices in a unit area. I.e., the buyers decide the intensity of small cells L_s and that of vehicles L_v to be deployed in their networks. We assume an existing macrocell tier with fixed, exogenous intensity L_m, leading to a 3-tier heterogeneous network. We focus only on the caching market and do not consider dual-purpose products, e.g., the backhaul access provided by femtocells that also act as caches. Our future work will generalize the current findings to other types of caches, e.g., including device-to-device caching.

To address the temporal and spatial dynamics of the caching system, we divide the decision space into T time slots and D regions. For each $(t, d) \in [T] \times [D]$, assume the locations of all communication equipments (i.e. caching points, macro cells and vehicles) follow some homogeneous Poisson point process (PPP) with constant intensity value $L_k(t, d), k \in \{m, s, v, u\}$ where u denotes users and $\{m, s, v\}$ the caching tiers. Since macro cells and small cells are unable to move, unlike vehicles, their intensities are fixed over time. I.e. $\forall t_1, t_2, L_k(t_1, d) = L_k(t_2, d), k \in \{m, s\}$. The price P_k per intensity that the operator must pay to use the caches, may depend on (t, d) as well, and is fixed in advance by the cache devices. We will consider finding the optimal prices in our future work.

To facilitate the analysis, contents are assumed to be cached as chunks with the same size. Each chunk is associated with an exogenous preference (probability of being requested) $f_i, i = 1, 2, \ldots$ Without loss of generality, we assume $f_1 \geq f_2 \geq \ldots$ Two chunks may come from the same content source yet have different preferences, e.g., some segments of a video may be more popular than others [4]. The preferences can also be a function of (t, d), which allows the temporal and spatial differences of user demands. In that case, the index i can itself depend on time and region. But we generally omit the (t, d) notation for the purposes of clarity.

All devices within the same tier share the same configuration, i.e. they have the same price $P_k(t, d)$, transmit power p_k and storage capacity N_k. Communication inside a region is assumed to only interfere with devices within that region. The content allocation policy is modeled in terms of the caching probability. In tier k, the content i is cached with probability $H_{k,i}(t, d)$. We can therefore model different allocation policies by choosing different values of H. It follows that all chunks are "cached" with probability 1 in the macro cell tier, i.e. $H_{m,i}(t, d) \equiv 1$.

We assume caching tiers are transparent to the users. I.e. user requests are always directed to the cell with the highest signal-to-interference-plus-noise ratio (SINR). This guarantees that the users always benefit from caching. The assumption can be relaxed without changing the framework of this study by assigning to each tier a connecting preference [2].

3 PROFIT MAXIMIZATION

The operators solve the maximization problem as follows to find the optimal intensities L_v^*, L_s^* which they should subscribe.

$$\begin{aligned} \underset{L_v, L_s \in \mathbb{R}^{[T] \times [D]}}{\text{maximize}} \quad & \gamma U(L_v, L_s) - \operatorname{tr}(P_v^T L_v) - \operatorname{tr}(P_s^T L_s) \\ \text{subject to} \quad & L_v(t, d) \geq 0, L_s(t_1, d) = L_s(t_2, d) \geq 0. \end{aligned}$$

Here $U(L_v, L_s)$ is some utility function, γ is a scaling coefficient, and the remaining terms represent the cost of using the cache. In the remaining part of this section, we will discuss two possible

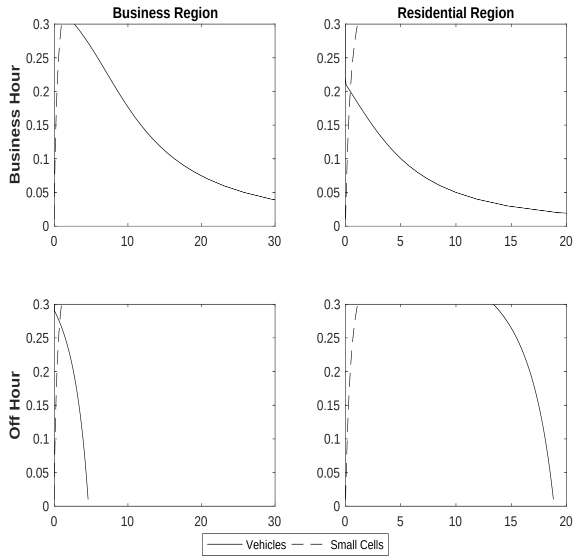

Figure 1: As the price $P_v(1, \cdot)$ varies (y-axis), the optimal vehicle intensity at each time (x-axis) also varies, allowing the CDN to achieve a higher profit than with only small cells.

utility models under the following naive content allocation policy:

$$H_{k,i}(t, d) = \begin{cases} 1, & i \le N_k \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

Following the naive allocation policy, each tier sequentially caches the most popular chunk ($i = 1$), the second most popular chunk ($i = 2$), etc., until reaches its capacity N_k. This policy is generally not optimal, but it can be easily implemented without the need for additional information required by other policies.

3.1 Utility from Cache Hits

CDN operators are typically paid for the amount of data offloaded to their infrastructure. They thus earn more profits with a higher rate of cache hits. Thus, we define their utility as the probability that a typical request is offloaded to the cache (Eq. 2). The γ coefficient then denotes the marginal value of cache hits, which is proportional to the per-byte monetary payoff to the CDN service.

$$U = \sum_{t,d} \frac{L_u(t,d)}{\|L_u\|_1} \sum_i f_i(t,d) \sum_{k \in \{s,v\}} \mathbb{P}(S = k | t, d, i) \tag{2}$$

Here $S = k|t, d, i$ denotes the event that tier k is selected for a typical request for content i originating from (t, d). It can be shown that under the PPP distribution and the transparent network assumption, the probability $\mathbb{P}(S = k|t, d, i)$ has the form

$$\mathbb{P}(S = k|t, d, i) = \frac{p_k^{2/\alpha} H_{k,i}(t,d) L_k(t,d)}{\sum_{j \in \{m,s,v\}} p_j^{2/\alpha} H_{j,i}(t,d) L_j(t,d)} \tag{3}$$

Here α is the path loss coefficient. Applying the naive allocation policy, and assuming that small cells have larger capacity ($N_s > N_v$), we can derive closed-form solutions for the optimization problem.

As a demonstration, we consider an area with a business region and a residential region. A typical day is divided into two time slots: the business hours and the off hours. For simplification, suppose prices are only functions of time, i.e. both regions share the same price at each time slot. Figure 1 shows the change of the optimal demands for both products with respect to the price of vehicles in the business hours. The price in business hours also affects the demands in off hours. We see that the optimal vehicular intensity greatly varies in the business and off hours, reflecting the value of having mobile caches that can relocate to different areas.

3.2 Utility from Downlink Rate

Unlike CDNs, ISPs do not directly profit from cache hits. Instead, their revenue comes from mobile users. In a typical competitive market, an ISP attracts more subscribers by improving the performance (e.g. coverage, downlink rate, uplink rate etc.) of its network, and caching is one way to improve the downlink rate. Thus, as is suggested by [2], we define the utility function as the logarithm of the downlink rate. Its marginal value (i.e. the γ term) can be evaluated using tools such as the discrete choice model [1].

Let W_k be the bandwidth associated with tier k, β be the portion of bandwidth allocated to the typical user, τ be the SINR threshold, C be the probability the typical user is in coverage. The logarithm of the downlink rate $\log R$, can be written as

$$U = \log R = \log\log(1+\tau) + \log(\beta) + \log(W) + \log(C) \tag{4}$$

With this utility function, the ISP's resulting optimization problem is generally non-convex. But it can be efficiently solved by fractional programming algorithms [3], and we find that the optimal intensity varies similarly to Figure 1 for CDNs.

4 CONCLUSION AND FUTURE WORK

We propose to use vehicles as caching points to improve the performance of the current wireless network. The mobility of vehicles is modeled by discretizing the space into several regions and time slots and using stochastic geometry tools. To study the economic value of the vehicular caching, we formulate a profit maximization problem with small cell caching as a competing product.

For the next step, we plan to refine the current methods, and look into some key problems that can be solved by our models. Some possible insights include: 1) How is the demand affected by prices in other time slots and regions; 2) How much do the CDN operators/ISPs/users benefit from vehicular caching; and 3) How much lower can vehicle intensity be (thus potentially reducing caching costs) while yielding the same utility as a small cell cache deployment.

REFERENCES

[1] Pierre Coucheney, Patrick Maillé, and Bruno Tuffin. 2013. Impact of competition between isps on the net neutrality debate. *IEEE Transactions on Network and Service Management* 10, 4 (2013), 425–433.

[2] Yicheng Lin, Wei Bao, Wei Yu, and Ben Liang. 2015. Optimizing user association and spectrum allocation in HetNets: A utility perspective. *IEEE Journal on Selected Areas in Communications* 33, 6 (2015), 1025–1039.

[3] Kaiming Shen and Wei Yu. 2018. Fractional Programming for Communication SystemsâĂŤPart I: Power Control and Beamforming. *IEEE Transactions on Signal Processing* 66, 10 (2018), 2616–2630.

[4] Luigi Vigneri. 2017. *Vehicles as a mobile cloud: modelling, optimization and performance analysis.* Ph.D. Dissertation. Université Côte d'Azur.