ON A SELFISH CACHING GAME

P. MICHIARDI, C.F. CHIASSERINI, C. CASETTI, C.L. LA, M. FIORE

Introduction : We address the problem of content caching in a hybrid wireless network: mobile nodes can connect to a cellular network (e.g., a mobile broadband network such as 3G) and are able to form a temporary multi-hop network (e.g., a 802.11-based ad hoc network). We assume content to be hosted at an origin server in the Internet, which can only be accessed through the cellular network. Nodes using the cellular network are able to download a fresh version of the content, which will be cached and served to other nodes issuing requests over the multi-hop network. Content popularity drives access behavior: in this work we assume a "flash-crowd" scenario, in which users discover a new content and wish to access it concurrently. As a consequence, *access congestion* determines to a large extent the download performance, for both the cellular and the ad hoc network. For simplicity, here we consider nodes to be interested in a single information object.

The problem of caching has received a lot of attention in the past due to its importance in enhancing performance, availability and reliability of content access in wireless systems. However, this problem has been addressed often under the assumption that nodes would cooperate by following a strategy that aims at optimizing the system performance, regardless of the costs incurred by each individual node. Our goal, instead, is similar to the one in [1], in that we build a model where nodes are selfish, i.e., they choose whether to cache or not the content so as to minimize their own cost. Our work however differs from [1] in how content demand is modeled.

Let $\mathcal{I}$ be a set of nodes uniformly deployed on area $\mathcal{A} = \pi R^2$. We consider a single base station covering the network area, and we denote by $r < R$ the radio range that nodes use for ad hoc communications. Also, let the information object be of size equal to L bytes; the object requires f updates per second from the origin server (in order to obtain a fresh copy) and each update implies the download of U bytes.

We now formulate the caching problem as a (simultaneous move) caching game. $\mathcal{I}$ is the set of players, with $|\mathcal{I}| = I$, and $\mathcal{S}_i = \{1, 0\}$ is the set of all possible strategies for player $i \in \mathcal{I}$. Additionally, let $s_i \in \mathcal{S}_i$ be the strategy of player i, where $s_i = 1 \vee 0$. The array $s = \{s_1, s_2, ..., s_i, ..., s_I\}$ is a strategy profile of the game. Furthermore, let $\mathcal{C} \subseteq \mathcal{I}$ be the set of players whose strategy is to access the object from the origin server and cache it, that is $s_i = 1, \forall i \in \mathcal{C}$, and $\mathcal{N} \subseteq \mathcal{I} = \mathcal{I} \setminus \mathcal{C}$ the set of players whose strategy is to access the cached object through ad hoc communications, that is $s_i = 0, \forall i \in \mathcal{N}$. Finally, let $|\mathcal{C}| = x$ and $|\mathcal{N}| = I - x$. In this game we need to impose $\mathcal{C} \neq \emptyset$: at least one player has to cache the object for otherwise content would not be available to any player. Given a strategy profile s, the cost incurred by player i is defined as:

$$C_i(s) = \beta_i \mathbb{I}_{s_i=1} + \gamma_i \mathbb{I}_{s_i=0}$$

where β_i and γ_i are the air time costs if i obtains the content through the 3G network and via the ad hoc network, respectively, and $\mathbb{I}_{s_i}$ is the indicator function. In this work we focus on access costs, neglecting the energy cost that a node caching the object experiences when serving other nodes.

We now define precisely the two terms β_i and γ_i. To this end, let us introduce the following quantities: R_{3G} and R_h are the bit rates offered, respectively, by the cellular and the ad hoc network; T_c is the time for which node $i \in \mathcal{C}$ caches the object; h is the *average* number of hops required to access the closest cached object, assuming a uniform distribution of nodes on $\mathcal{A}$ and a *uniform distribution of caches*. Formally, $h = \sqrt{\frac{\mathcal{A}}{x\pi r^2}} = \frac{R}{r}\frac{1}{\sqrt{x}}$.

With these definitions at hand, we can now focus on the two cost terms, β_i and γ_i:

$$\beta_i = |\mathcal{C}|\frac{L + T_c f U}{R_{3G}} = |\mathcal{C}|k_1 \qquad \gamma_i = |\mathcal{N}|\frac{hL}{R_h} = \frac{|\mathcal{N}|}{\sqrt{|\mathcal{C}|}}k_2 \ , \ k_2 = \frac{R}{r}\frac{L}{R_h}$$

β_i distinguishes "installation" and "maintenance" costs and models the congestion incurred by nodes trying to access the information object at the same time: the bit rate R_{3G} is inversely proportional to the number of concurrent users accessing a *single* 3G base station [2]. Similarly, γ_i represents the air time consumed to access the current version of the cached object and models the congestion cost created by simultaneous access of nodes operating in ad hoc mode[1].

Social Optimum : The social cost of a given strategy profile is defined as the total cost incurred by all players, namely:

$$C(s) = \sum_{i\in\mathcal{C}}\beta_i + \sum_{i\in\mathcal{N}}\gamma_i \qquad \rightarrow \qquad C(x) = x^2 k_1 + \frac{(I-x)^2}{\sqrt{x}}k_2$$

The minimum social cost is $\hat{C}(s) = \min_s C(s)$. We solve the optimization problem and plot in Fig. 1 the number x of caches that minimizes the social cost, for $I = 100$, $R_{3G} = 2$ Mbps, $L = 1000$ bytes, $U = 100$ bytes, $T_c = 100$ s, and $f = 0.01$ updates/s.

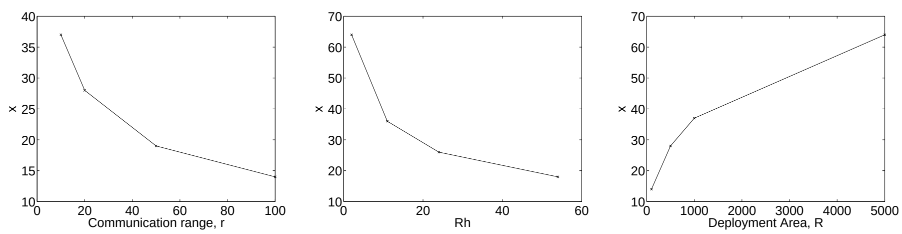

(a) R_h=24 Mbps, R=500 m (b) r=20 m, R=500 m (c) r=20 m, R_h=24 Mbps

FIGURE 1. Number of caches minimizing the social cost, as a function of system parameters: x decreases with r and R_h (improved communication capabilities of the ad hoc network) and increases with R (decreasing nodal density).

Next, we focus on a simple two-player game, derive its equilibrium points and compare their efficiency to the social optimum.

A Two Player Game and its Extensions : The two-player version of the caching game involves two players whose strategy set is $S_i = \{1, 0\}$ as defined above. For clarity, we label $s_i = 1 \rightarrow C$ and $s_i = 0 \rightarrow N$. Fig. 2 illustrates the normal form game[2].

	$p_2(C)$	$p_2(N)$
$p_1(C)$	$\left(4k_1, 4k_1\right)$	$\left(k_1, k_2\right)$
$p_1(N)$	$\left(k_2, k_1\right)$	(∞, ∞)

(a) Cost matrix, c_{ij}

	$p_2(C)$	$p_2(N)$
$p_1(C)$	$\left(\frac{1}{4k_1}, \frac{1}{4k_1}\right)$	$\left(\frac{1}{k_1}, \frac{1}{k_2}\right)$
$p_1(N)$	$\left(\frac{1}{k_2}, \frac{1}{k_1}\right)$	$(0, 0)$

(b) Payoff matrix, π_{ij}

FIGURE 2. Two-player caching game in normal form. $p_i(x)$ indicates player i choosing strategy x.

[1] Our congestion model is more conservative than the capacity scaling law defined in [3].

[2] When no player caches the object, access costs are infinite.

Clearly, strategy N is *strictly dominated* by strategy C if and only if $4k_1 < k_2$: in this case, we would have only one Nash Equilibrium (NE), which is (C, C). Instead[3], when $4k_1 > k_2$, we face an *anti-coordination game*, in which players randomize their strategies. Indeed, there are two conflicting (in terms of payoffs) NE points, *i.e.*, the (N, C) and (C, N) strategy profiles. It is well known that mixed-strategies profiles and expected payoffs π_i can be derived as follows. Suppose player 2 chooses C with probability α, then $\mathbb{E}[\pi_1(C, \alpha)] = \frac{4-3\alpha}{4k_1}$ and $\mathbb{E}[\pi_1(N, \alpha)] = \frac{\alpha}{k_2}$. Hence, $\alpha = \frac{4k_2}{4k_1+3k_2}$. Due to the symmetry of the game, player 1 chooses C with probability $\beta = \alpha$. Considering the joint mixing probabilities, the expected payoff for both players is $\mathbb{E}[\pi_i^*] = \frac{4}{4k_1+3k_2}\, \forall i \in 1, 2$.

It is worth noting that in this anti-coordination game the mixed strategy NE is *inefficient*. Indeed, when players can correlate their strategies based on the result of an observable randomizing device (*i.e.*, a *correlated equilibrium* can be achieved), the expected payoff is $\mathbb{E}[\hat{\pi}_i] = \frac{k_1+k_2}{2k_1k_2}\, \forall i \in 1, 2$. We observe that $\mathbb{E}[\hat{\pi}_i]$, which corresponds to the social optimum, is *strictly larger* than $\mathbb{E}[\pi_i^*]$. This clearly suggests that some correlation among the nodes' actions should be introduced in order to improve system performance.

On-going work : The results above can be extended to an n-player setting, which will be treated in detail in an extended version of this work. Our current research aims at putting into practice our theoretic findings, following two complementary directions.

On the one hand, we note that, in the n-player caching game, a player can compute its best response to other players' strategies if it is aware of the current number x of caches in the network. Since in practice global knowledge cannot be assumed, we are investigating how far from efficiency our system settles when nodes compute an estimate $\hat{x}$ of the current number of caches in the network. Such an estimate can be obtained either through random sampling techniques based on gossiping, or by exploiting *local measurements* of the number of queries received by each node caching the content. An open question is how sensitive to estimation errors the achieved equilibrium is.

On the other hand, we observe that an external randomization device can help in improving efficiency, but correlated equilibrium is impractical when players' actions are not simultaneous, i.e., in a *asynchronous setting*. To address this issue, we allow communication between players through *signalling*. Simply stated, signalling replaces the external randomization device cited above and is used by a player to notify its strategy to others. The use of signalling however implies to take into account neighboring relations among players, as dictated by the underlying communication graph defined by the network topology.

Conclusions : We proposed a novel model for the caching problem in a heterogeneous network under a "flash-crowd" scenario. We provided the expression for the social cost and defined a two-player game to obtain insights into the design of efficient caching strategies. Based on the theoretical findings, our current work focuses on the design of strategies to be implemented in a practical network setting.

References

[1] B.G. Chun and K. Chaudhuri and H. Wee and M. Barreno and C.H. Papadimitriou and J. Kubiatowicz, "Selfish caching in distributed systems: a game-theoretic analysis", *ACM PODC*, 2004.

[2] W.L. Tan, F. Lam and W.C. Lau, "An Empirical Study on the Capacity and Performance of 3G Networks", *IEEE Transactions on Mobile Computing*, 2007

[3] P. Gupta and P.R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, 2000.

[3] This is the case that happens in practice, e.g., with the values of the system parameters used to compute the minimum cost.