On the CDN Pricing Game

Yang Song
University of Massachusetts Amherst
ysong@ecs.umass.edu

Arun Venkataramani
University of Massachusetts Amherst
arun@cs.umass.edu

Lixin Gao
University of Massachusetts Amherst
lgao@ecs.umass.edu

***Abstract*—Content Delivery Networks (CDNs) serve a large fraction of Internet traffic today improving user-perceived response time and availability of content. With tens of CDNs competing for content producers, it is important to understand the game played by these CDNs and whether the game is sustainable in the long term. In this paper, we formulate a game-theoretic model to analyze price competition among CDNs. Under this model, we propose an optimal strategy employed by two-CDN games. The strategy is incentive-compatible since any CDN that deviates from the strategy ends up with a lower utility. The strategy is also efficient since it produces a total utility that is at least two thirds of the social optimal utility. We formally derive the sufficient conditions for such a strategy to exist, and empirically show that there exists an optimal strategy for the games with more than two CDNs.**

I. Introduction

A content delivery network (CDN) is a large distributed system that caches content at multiple locations in the Internet. When a client makes a request, a CDN applies an algorithm to choose the best server (usually the nearest one) to serve the content. CDNs aim to enhance the user-perceived experience by reducing the delay and ensuring predictable performance. Today, a significant fraction of traffic on the Internet is delivered by CDNs.

Tens of CDNs are currently competing with each other to attract business from content producers. However, the impact of this competition on shaping the CDN market is poorly understood. For example, it is known that the unit price of CDN services has been dropping for at least five consecutive years [1], but it is unclear whether the dropping prices may result in a price war wherein no CDN can remain profitable. Recent events such as Level-3 Communication charging a lower price than Akamai Technologies to attract Netflix's video content drew a lot of attention [3, 4]. In the light of such events, it is important to develop a rigorous technical foundation to understand and analyze pricing strategies practiced by CDNs.

In this paper, we propose a game-theoretic model to analyze the competition among CDNs. The competition is formulated as a repeated game where CDNs make alternative moves. We first show that, if a CDN determines its price based on the current utility only, then the competition can lead to price wars. Then we demonstrate that if a CDN incorporates the future payoff-relevant information, then the game can avoid price wars and reach Markov perfect equilibrium under which every CDN achieves the highest utility. We prove that there exists a strategy that leads to Markov perfect equilibrium for two-CDN games. More importantly, we show that the strategy is both incentive-compatible and efficient. That is, any CDN that deviates from the strategy has a lower utility, and the total utility under the strategy is at least two thirds of the social optimal utility. We formally derive sufficient conditions for such a strategy to exist for two-CDN games, and show that these conditions can be easily satisfied by real world scenarios. At last, we simulate the games with more than two CDNs, and show that such a strategy can exist for games with more than two CDNs as well.

A summary of our contributions are as follows. (1) Propose a strategy so that any CDN deviating from it produces a lower utility; (2) Derive sufficient conditions for such a strategy to exist; (3) Prove that the total utility under the strategy is at least two third of the social optimal utility.

The paper is organized as follows. Section II presents the game-theoretic model. Section III introduces how content producers select CDNs. In Section IV, we describe CDN pricing games. Section V shows the strategies that lead to Markov perfect equilibrium in two-CDN games. In Section VI and Section VII, we present the sufficient conditions for such strategies to exist and the price of anarchy of the strategies. Section VIII simulates games with more than two CDNs. Section IX is the related works, and Section X concludes the paper.

II. Model

Our model focuses on multiple CDNs competing over a single content producer, indicating that each content producer makes choices independently of others. We denote a content producer as S. Assume S provides service to N end-users. The N end-users make up the entire universe, U, of users considered in the CDN pricing game. S may use one or multiple CDNs depending on their respective prices and coverage.

The service quality provided by a CDN is determined by many factors, including: the distance between the targeted user and nearby CDN servers, cache size at the CDN servers, congestion level, server selection and load balancing algorithms, traffic demand patterns, network conditions. To make our model easy to analyze and practical, we only consider the combined effect of all of these factors and simply partition the end-users into two categories with respect to each CDN: an end-user can either be a *high-quality* user, i.e., one that receives improved service quality through the CDN, or a *low-quality* user that does not benefit or benefit little from the CDN, e.g., users in China are likely to see little improvement in service quality by going through a CDN that has no cache

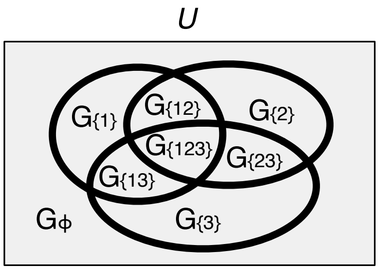

Fig. 1. Coverage overlap between S and CDNs

presence in Asia. Although this simplified partitioning of users into just two categories ignores the details mentioned before, it is nevertheless useful to gain insights into CDN games.

As we will see later, the set of the high-quality category users play an important role in attracting content. We refer to the set of users as the *coverage* of a CDN. Clearly, coverage can be different across different CDNs because of typical differences in their service locations and their architecture. We further assume that there is no significant difference in service among CDNs for the same category. That is, if an end-user belongs to the high (low) quality category of all CDNs, then it experiences the same service quality through all the CDNs. We denote a high-quality user's service improvement as η_1 per unit traffic, and a low quality category user's as η_2. Parameter η_1 represents the benefit of CDN service, such as reduction in end-user perceived latency and increased reliability.

The entire universe of users is partitioned into regions by CDNs' coverage. Let C_i denote CDN i, where $i = 1, 2, ..., M$. We use G_I to represent the subset of S's end-users that only receive high-quality service from C_is, where $i \in I$. For example, $G_{\{1\}}$ represents the end-users that are covered by C_1 *only*, and $G_{\{1,2\}}$ represents the end-users that are covered by both C_1 and C_2. Note that, $G_{\{1\}}$ and $G_{\{1,2\}}$ are exclusive sets that have no overlap. According to the partition, the universe U of users is divided into K groups, where K is the number of all possible combinations of CDNs. Figure 1 shows the partitions by three CDNs. The circles in the figure represent the coverage of different CDNs. We can see that, the coverage is divided into $K = 8$ parts, and $G_{\{\cdot\}}$ represents the part that is covered by C_is, where $i \in \{\cdot\}$, and G_ϕ represents the part that is not covered by any CDN.

III. A CONTENT PRODUCER'S CDN SELECTION

This section demonstrates how content producer S selects CDNs. We define S's utility as the benefit obtained from CDN service minus the price it pays. Assume the benefit from CDN service is the traffic volume multiplied by the improvement in service quality achieved for that traffic. We further assume that the price paid by S to a CDN is the total volume served by the CDN times the unit price charged by the CDN. For example, if S selects C_i for the end-users in G_I, then the utility of S gained from G_I is $(\eta_i^{G_I} - p_i)v_{G_I}$, where p_i is the price charged by C_i, and v_{G_I} is the total volume sent to G_I. If $i \in I$, then $\eta_i^{G_I} = \eta_1$. Otherwise, $\eta_i^{G_I} = \eta_2$.

A content producer S selects CDNs to maximize its utility. The selection depends on CDNs' prices and coverage. S chooses C_i for G_I if and only if S prefers it over all the alternatives. That is, if for $q = 1, 2, ..., M$ and $q \neq i$,

$$\eta_i^{G_I} - p_i > \eta_q^{G_I} - p_q \tag{1}$$

Thus, the total utility of S is:

$$\pi_s = \sum_{G_I} \max_i\{(\eta_i^{G_I} - p_i)v_{G_I}\} \tag{2}$$

If several CDNs provide the same utility, then S divides traffic evenly among them.

We define a matrix E to denote S's CDN selection. E is a $M \times K$ matrix, where rows represent CDNs and columns represent end-user groups in lexicographical order. E_{ij} represents the percentage of group j's traffic that is delivered by C_i. For example, suppose G_j represents the group belonging to the j-th column, and then the order of the groups in a two-CDN game are $G_1 = G_\phi, G_2 = G_{\{1\}}, G_3 = G_{\{1,2\}}, G_4 = G_{\{2\}}$. In this case, $E_{11} = 0.5$ means half of G_ϕ's traffic is delivered by C_1. With the definition of E, the utility of S can be written as:

$$\pi_s = \sum_{j=1}^{K} \sum_{i=1}^{M} [E_{ij}(\eta_i^{G_j} - p_i)v_{G_j}] \tag{3}$$

IV. CDN PRICING COMPETITION

A. CDN Pricing game

The CDN pricing game is defined as a repeated game where price decisions are made sequentially by individual CDNs. Assume that there is a sequence that starts with C_1 and ends with C_M. At the beginning of time period $kM + i$, where $k = 0, 1, 2, ...\infty$, C_i selects a price, and that price is stable for the rest of the period. All other CDNs cannot change their price during the period when C_i is about to move. After a CDN decides price at the beginning of each period, S makes a CDN selection according to Equation 3 and pays corresponding CDNs for the service.

The utility of a CDN is its income (the price paid by S) minus the cost to provide the CDN service. The utility of CDN C_i can be written as $\pi_{c_i} = (p_i - c)v_{c_i}$, where c is the cost per traffic unit and it is the same for all CDNs, and v_{c_i} is the total traffic volume served using C_i. Now, based on S's CDN selection in Equation 3, we can write C_i's utility during time interval t as:

$$\pi_i^t(p_1^t, p_2^t, ..., p_i^t, ..., p_M^t) = (p_i^t - c)(\sum_{j=1}^{K} E_{ij}v_{G_j}) \tag{4}$$

where p_i^t is the price of C_i at interval t.

In CDN pricing game, if all CDNs ignore the impact of the current action on the future, and focus on maximizing their *current* utility, then the actions may lead to price wars. If all CDNs plan ahead and maximize utility over time, then price wars can be avoided. We name the former strategy the *non-predictive strategy*, and the latter the *predictive strategy*. In Section IV-B, we focus on the non-predictive strategy. The predictive strategy will be discussed in Section IV-C.

B. The non-predictive strategy that is prone to price wars

We define price wars as a series of actions that lead to a state where there exist at least two CDNs that have zero utility and neither of the CDNs can increase its utility by changing its price.

Under the non-predictive strategy, any CDN that is about to move in the current period aims to maximize it current utility. Suppose C_i is about to move during period t. Then C_i will choose a price p_i^t to maximize $\pi_i^t(p_1^t, ..., p_i^t, ..., p_N^t)$. In order to formally prove the existence of price wars, we assume that if the price change $\Delta p \rightarrow 0$, then $\pi_i^t(p_1^t, ..., p_i^t, ..., p_N^t) - \pi_i^t(p_1^t, ..., p_i^t + \Delta p, ..., p_N^t) \rightarrow 0$. We have the following theorems.

***Theorem** 1:* In a multiple-CDN pricing game, if there exist at least two CDNs that have the same coverage and apply the non-predictive strategy, then the pricing game must lead to price wars.

The proof of the theorem can be found in [13].

The vulnerability to price wars naturally raises the following questions: if CDNs think ahead of time, will the price wars be avoided? Is there a condition that guarantees to avoid the price wars? What strategy to use given that all of the CDNs must act selfishly? Lastly, will the strategy be efficient and produce a utility that is near to the social optimal utility? The following content of the paper addresses these questions.

C. Search for predictive strategies

A key insight to avoid price wars is that if CDNs can predict price wars, then they may apply a strategy to avoid it. To allow CDNs to make decisions not only based on current utility but also on the future, we change the objective function that a CDN maximizes to be its utility over time. Any CDN to move in the current period knows that its rivals' future actions will be based on its current action. Thus, a CDN's strategy must be adjusted accordingly. In this section, we discuss how to derive such predictive strategies.

We introduce how to derive predictive strategies for two-CDN games, namely, we focus on a game where two CDNs C_1 and C_2 set their prices iteratively to maximize their own utility over time:

$$I_1 = \sum_{t=0}^{\infty} \delta^t \pi_1^t(p_1^t, p_2^t), \text{ and } I_2 = \sum_{t=0}^{\infty} \delta^t \pi_2^t(p_1^t, p_2^t) \quad (5)$$

, where $\delta < 1$ is a discount factor. The discount factor indicates the value reduction for the future utility since the future utility cannot be obtained immediately.

We use Markov perfect equilibrium theory [10, 14] to solve the two-CDN game. Assume that each CDN's next price is only determined by the last price of the other CDN, but not by the earlier histories of the prices. Then, the predictive strategy can be defined by a pair of reaction functions $(R_1(\cdot), R_2(\cdot))$ as follows.

$$p_1^{2\tau+1} = R_1^{2\tau+1}(p_2^{2\tau}), \quad p_2^{2\tau+2} = R_2^{2\tau+2}(p_1^{2\tau+1}) \quad (6)$$

where $\tau = 0, 1, 2, ...$

A pair of reaction functions (R_1, R_2) is the *optimal predictive strategy* if and only if R_1 maximizes C_1's utility over time given C_2's strategy R_2, and R_2 maximizes C_2's utility over time given C_1's strategy R_1. Note that once we specify a reaction function R_1, by optimizing the utility of C_2, we implicitly define the corresponding R_2 and vice versa.

Now, we look back at the function for utility over time in Equation 5. We can substitute the prices in the equation with the reaction functions. Suppose the current time period is t. Then maximizing I_1 at time t is equivalent to choosing a price p_1^t so as to maximize I_1. Because p_2^{t-1} is known at time t, I_1 can be viewed as a function of p_2^{t-1}: $I_1(p_2^{t-1})$. To maximize $I_1(p_2^{t-1})$ is equivalent to choose p_1^t to maximize $\{\pi_1^t(p_1^t, p_2^{t-1}) + \delta\pi_1^t(p_1^t, p_2^{t+1}) + \delta^2\pi_1^t(p_1^{t+2}, p_2^{t+1}) + \delta^3\pi_1^t(p_1^{t+2}, p_2^{t+3}) + ...\} = \{\pi_1^t(p_1^t, p_2^{t-1}) + \delta\pi_1^t(p_1^t, R_2(p_1^t)) + \delta^2\pi_1^t(R_1(R_2(p_1^t)), R_2(p_1^t)) + \delta^3\pi_1^t(R_1(R_2(p_1^t)), R_2(R_1(R_2(p_1^t)))) + ...\}$. Note that starting at period $t+2$, the game will repeat the same decision process as period t. Thus, we can substitute the utility after $t+2$ with $I_1(p_2^{t+1})$. That is because, at time $t+2$, p_2^{t+1} is the price that C_1's price will be based on. Thus, we have:

$$\begin{aligned} \max\ I_1(p_2^{t-1}) &= \max_{p_1^t}\{\pi_1^t(p_1^t, p_2^{t-1}) + \delta\pi_1^t(p_1^t, R_2(p_1^t)) \\ &\quad + \delta^2 I_1(R_2(p_1^t))\} \end{aligned} \quad (7)$$

Equation 7 indicates that, given C_2's price at time $t-1$, to maximize the utility over time at time t, C_1 has to find a price p_1^t knowing that C_2 will make a decision based on p_1^t at time $t+1$. By solving Equation 7, we can map every possible p_2^{t-1} to a corresponding price p_1^t, and by definition, the result is reaction function R_1.

Similarly, we can derive the utility function that C_2 aims to maximize at time k.

$$\begin{aligned} \max\ I_2(p_1^{k-1}) &= \max_{p_2^k}\{\pi_2^t(p_1^{k-1}, p_2^k) + \delta\pi_2^k(R_1(p_2^k), p_2^k) \\ &\quad + \delta^2 I_2(R_1(p_2^k))\} \end{aligned} \quad (8)$$

The resulting mapping function from p_1^{k-1} to p_2^k is, by definition, R_2.

Note that any strategy satisfies Equation 7 and 8 is the optimal predictive strategy for the two-CDN pricing game. We name the optimal predictive strategy the *Markov strategy*. According to the definition, if CDNs implement a Markov strategy, then their utility is maximized. In the following content, we show that, there exists a Markov strategy for CDN pricing games.

Under Markov strategies, if prices eventually become stable for all the players, then the prices are called a *focal price*[1].

[1]If prices oscillate within a same price cycle, then we call the price cycle the *Edgeworth cycle*. We only focus on focal prices in this paper

V. Markov strategies for two-CDN games

In this section, we propose Markov strategies that achieve the highest utility for each CDN. The derivation is omitted due to the space limitation (see [13] for details). Markov strategies exist under two types of coverage. The first one is that CDNs have the same size of coverage with arbitrary overlaps, and the second type is that one CDN's coverage is subsumed by the other CDN's. We choose these two types because they are the most representative coverage types in the current CDN market.

Suppose there are two CDNs, C_1 and C_2, and their prices are p_1 and p_2. In this case, the end-users can be divided into four groups: $G_{\{1\}}, G_{\{2\}}, G_{\{1,2\}}, G_{\{\phi\}}$. The traffic volume sent to the four groups are denoted by v_1, v_2, v_3 and v_4 respectively.

A. Markov strategies with same size coverage

In this subsection, we first introduce the Markov strategy with same size coverage ($v_1 = v_2$). Then we demonstrate the incentive of the strategy. Lastly, we show how to apply the strategy in practice.

For the same size coverage CDN pricing game, the Markov strategy is given by:

$$R_i(p) = \begin{cases} p_f & \text{for } p \geq p_f \\ p_l & \text{for } p_f > p > p_l \\ p_f & \text{for } p \leq p_l \end{cases} \tag{9}$$

where, $i = 1, 2$ and p_f is the focal price for both CDNs, and $p_l < p_f$. How to decide p_f and p_l is discussed in Section V-A2.

According to the strategy, if the rival CDN charges a price that is greater than or equal to focal price p_f, or smaller than p_l, then a CDN needs to charge the focal price p_f. If the rival CDN's price is between p_f and p_l, then a CDN needs to charge p_l.

1) Incentive of the same size coverage Markov strategy: In this paragraph, we explain why the strategy provides the optimal utility for an individual CDN. The formal proof can be found in [13]. First, the strategy prevents CDNs from charging a low price. According to the strategy, if any CDN charges a price lower than p_l, then the other CDN is required to charge the focal price p_f. Eventually, the CDN who charges at p_l or lower will change to the focal price p_f because the focal price is more profitable. The strategy also makes CDNs unable to derail from the focal price. When a CDN deviates from the focal price by charging a price between p_f and p_l, the other CDN is required to charge a low price (p_l) in order to force the deviating CDN to return to the focal price. When a CDN deviates from the focal price by charging a price higher than p_f, less traffic will select this CDN. By staying at the focal price, the other CDN can force the deviating CDN to go back to the more profitable focal price. Therefore, the same size Markov perfect strategy is motivated by incentives, and strategically forces the two CDNs to stay at the focal price, which yields the optimal utility.

2) Apply Markov strategies in practice: The next question is, how to decide the values of p_f and p_l. According to the proof in [13], p_f and p_l satisfy the following two functions: $\frac{\eta_1(2v_1+v_3)+\eta_2 v_4-c(2v_1+v_3+v_4)}{2v_1+\frac{5}{4}(v_3+v_4)} < p_f - c < \frac{\eta_1(2v_1+v_3)+\eta_2 v_4-c(2v_1+v_3+v_4)}{2v_1+v_3+v_4}$ and $(1 + \frac{3}{2\beta+1})(p_l - k - c) < p_f - c < (1+\frac{3}{2\beta+1})(p_l - c)$, where k is the unit price on the price grid, and $\beta = \frac{2v_1}{v_3+v_4}$. The condition of p_f shows the range of the focal price. In the first function, $\eta_1(2v_1 + v_3) + \eta_2 v_4$ represents the total benefit gained by using a CDN service and $c(2v_1 + v_3 + v_4)$ represents the total cost to maintain the CDN service. Thus, the right side of the first function is equal to the average social utility gained by per unit traffic. We define the average social utility as u_p. The left side of the first function is $\frac{\beta+1}{\beta+\frac{5}{4}} u_p$. According to the function, any price between u_p and $\frac{\beta+1}{\beta+\frac{5}{4}} u_p$ can be a focal price. Thus, if one CDN charges at any of the prices within the range, the other should follow the same price. After p_f is decided, we can set p_l to satisfy $p_l - c \approx \frac{2\beta+1}{2\beta+4}(p_f - c)$.

B. Markov strategies with subsumed coverage

For the subsumed coverage CDN games ($v_2 = 0$), the Markov strategy is as follows.

$$R_1(p_2) = \begin{cases} p_2 - k & \text{for } p_2 > p_2^f \\ p_1^f & \text{for } p_2^f \geq p_2 > p_1^f - (\eta_1 - \eta_2) \\ p_2 + (\eta_1 - \eta_2) - k & \text{for } p_2 \leq p_1^f - (\eta_1 - \eta_2) \end{cases} \tag{10}$$

$$R_2(p_1) = \begin{cases} p_1 - (\eta_1 - \eta_2) - k & \text{for } p_1 > p_1^f \\ p_2^f & \text{for } p_1^f \geq p_1 > p_2^f \\ p_1 - k & \text{for } p_2^f \geq p_1 > p_l \\ p_2^f & \text{for } p_l \geq p_1 \end{cases}$$

where k is the unit price on the price grid, and p_1^f, p_2^f are the focal price of C_1 and C_2, which satisfy $(p_2^f - k - c)(1+\beta) < (p_1^f - c)\beta < (p_2^f - c)(1+\beta)$ and $(p_1^f - k - c - (\eta_1 - \eta_2))(\frac{\beta}{2}+2) < 2(p_2^f - c) < (p_1^f - c - (\eta_1 - \eta_2))(\frac{\beta}{2}+2)$ and $p_2^f(v_1 + v_3 + v_4) > \eta_1 v_3 + \eta_2(v_1 + v_4)$. The application analysis which is similar to the same size coverage games can be found in [13]. It is interesting to see that, under the Markov strategy in Equation 10, the larger CDN always charges a higher price.

VI. Existence of Markov strategies

In this section, we discuss the sufficient conditions for Markov strategies to exist under the same size coverage and subsumed coverage. We demonstrate that the sufficient conditions can be satisfied by real world examples we tested.

A. Existence of Markov strategies with same size coverage

In this section, we introduce the sufficient conditions for a Markov strategy to exist if two CDNs have the same size coverage, i.e., $v_1 = v_2$. The conditions are as follows.

$$\text{Condition 1: } 0 \leq \beta < 2 \tag{11}$$

$$\begin{aligned} \text{Condition 2: } & (2v_3 + 5v_4)\eta_1 + 3c(v_3 + v_4) \\ & > (39v_3 + 46v_4)\eta_2 \end{aligned} \tag{12}$$

where, $\beta = \frac{2v_1}{v_3+v_4}$. It is straightforward to see that, Condition 1 requires the overlapped coverage of C_1 and C_2 to be large enough, and Condition 2 requires η_1 to be larger than η_2. We have the following theorem.

***Theorem** 2:* For a given pair of CDNs, if $v_1 = v_2$ and the conditions in Equations 11 and 12 are satisfied, then there must exist a Markov strategy.

The proof of the theorem is deferred to [13].

If two CDNs have the same coverage, then the sufficient condition can be simplified as:

$$\eta_1 v_3 + \eta_2 v_4 < \frac{5}{3}(\eta_1 - \eta_2)(v_3 + v_4) + c(v_3 + v_4) \qquad (13)$$

***Corollary** 1:* For a given pair of CDNs, if the two CDNs have the same coverage, and Equation 13 is true, then there must exist a Markov strategy.

This corollary shows that using the strategy we proposed, two same coverage CDNs can maximize their over time utility instead of having a price war as we have shown in Section IV-B.

1) An illustrative example: The sufficient conditions are easy to be satisfied in real applications. Suppose there is a content provider seeking CDN service from two CDNs: Level 3 (CDN1) and Amazon (CDN2). We assume that each CDN can provide high service quality for the users that are within 500 miles from its server locations, and that the number of Internet users within each region is equal to the population in the region multiplied by an Internet user ratio provided by www.internetworldstats.com. If the content provider sends the same amount of traffic v to each Internet user in the United States, then $v_1 = 1.3 \times 10^7 v, v_2 = 2.4 \times 10^5 v, v_3 = 1.2 \times 10^8 v$, and $v_4 = 1.1 \times 10^7 v$. Since v_1 and v_2 are much smaller than v_3, the scenario can be viewed as the same coverage case. If the two CDNs provide sufficient service improvement, more precisely, if $\eta_1 \geq 2.5\eta_2$, then the condition in Equation 13 is satisfied. By implementing the strategy in Equation 9, Level 3 and Amazon charge the same price p_f, and obtain the highest utility by maintaining the equilibrium. How Level 3 and Amazon select price p_f can be found in Section V-A2.

B. Existence of Markov strategies with subsumed coverage

In this section, we provide the sufficient conditions for a Markov strategy to exist if two CDNs have subsumed coverage (That is, $v_2 = 0$). Define $\beta = \frac{2v_1}{v_3+v_4}$, and the sufficient conditions are:

$$\text{Condition 1: } \beta > 1 \qquad (14)$$

$$\text{Condition 2: } (5v_1 - v_3 + 5v_4)\eta_1 + 6c(v_1 + v_3 + v_4) > (11v_1 + 5v_3 + 11v_4)\eta_2 \qquad (15)$$

***Theorem** 3:* For a given pair of CDNs, if one CDN's coverage is subsumed in the other CDN's coverage, and the conditions in Equations 14 and 15 are satisfied, then there must exist a Markov strategy.

1) An illustrative example: Suppose a content provider sends the same amount of traffic v to every user in the Internet. Then, under the same assumption in Section VI-A1, $v_1 = 3.0 \times 10^8 v, v_2 = 1.7 \times 10^7 v, v_3 = 2.3 \times 10^8 v, v_4 = 5 \times 10^7 v$. Since v_2 is much smaller than v_1, the scenario can be viewed as the subsumed coverage case. Because $\beta = 2.1$, Condition 1 in Equation 14 is satisfied. If the two CDNs provide sufficient service improvement, more precisely, if $\eta_1 \geq 2.7\eta_2$, then Condition 2 in Equation 15 is also satisfied. By implementing the strategy in Equation 10, Level 3 and Amazon can maintain a Markov perfect equilibrium for a global content producer, and the price of Level 3 should be 1.47 times of Amazon's price.

VII. PRICE OF ANARCHY

We define social optimal utility as the highest utility obtained by all of the CDNs if they optimize the summation of their total utilities instead of their individual utilities. Social optimal utility is the highest utility that can be reached if all CDNs act selflessly. The closer the total utility to the social optimal, the more efficient the strategy. The following theorem shows both the same size coverage and the subsumed coverage Markov strategy are efficient.

***Theorem** 4:* With the same size coverage and the subsumed coverage two-CDN pricing games, the utility obtained under the Markov strategies in Equation 9 and Equation 10 is at least two thirds of the social optimal utility.

The proof is in [13]. We can see that, Markov strategies provide an efficient way to compete. Without the strategies, the game can end up with price wars where the social utility is zero.

VIII. GENERALIZATION TO MULTIPLE CDNS

In previous sections, we proposed Markov strategies of two-CDN games, and presented the sufficient conditions for such strategies to exist. Next, we discuss how to generalize the results to more than two CDNs games. [5, 11, 12] proposed several methods to compute Markov perfect equilibrium with N players' game numerically. In this section, we use these methods to simulate CDN pricing game with three CDNs. We find that for the three-CDN games, the Markov strategy that leads to the equilibrium has similar rationale as two-CDN games.

The result for two sample experiments is shown in Figure 2. The figure illustrates the price that is chosen by different CDNs during each period of time until it converges to the focal price. In both experiments, we assume one CDN's coverage subsumes the other two CDNs', and v_s is the traffic volume sent to group s. In the first experiment, the two smaller CDNs have overlapped coverage, and in the second experiment, the two smaller CDNs have the same coverage. We can see that, the price charged by a CDN is proportional to its size, and this outcome is consistent with the prices in the real world CDN market. The strategy employed by the CDNs demonstrates the properties that we have discovered in the two-CDN case. That is, any price that is not equal to the focal price results in a

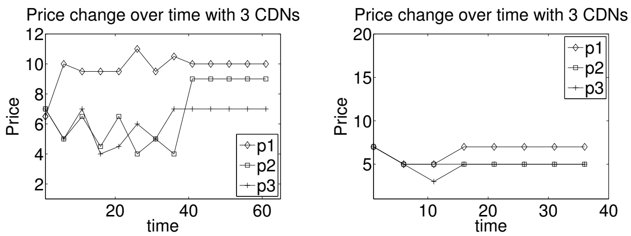

(a) Example 1.

(b) Example 2.

Fig. 2. Two examples that reach focal price with three CDNs. Example 1 shows two smaller CDNs have overlapped coverage ($v_{\{1\}} = 2$, $v_{\{1,2\}} = 3$, $v_{\{1,3\}} = 2$, $v_{\{1,2,3\}} = 3$, $v_{\phi} = 4$). Example 2 shows two smaller CDNs have the same coverage ($v_{\{1\}} = 4$, $v_{\{1,2,3\}} = 4$, $v_{\phi} = 2$).

lower utility, and by staying at the focal price, a CDN can force other rival CDNs to return to the focal price. This can be a general conclusion for the price games with more than two CDNs, and we leave the formal proof and experiments in the future work.

IX. Related works

CDNs play a critical role in the Internet supply chain, which consists of content producers, Internet service providers (ISPs), and access networks. CDNs not only provide significant benefit for both content producers and end users by reducing delay and improving robustness, but also achieve low infrastructure margin costs by aggregating content across multiple sites. Due to the significant value of CDNs, it is important to ensure that customers continue to derive benefit and CDNs simultaneously maintain profits and have incentives to manage their infrastructure. Pricing schemes investigate how to balance the two aspects by pricing the service.

Early works about CDN pricing focused on bulk pricing strategies. In [2, 7, 8], it is suggested that pricing functions should consider the burst traffic and provide volume discounts to content providers. A series of papers discussed CDN's pricing schemes in the Internet supply chain. A pricing method is presented in [6]. The method distributes benefits among different components in the supply chain. [4] addressed the questions about how CDNs should charge content providers and be charged by ISPs. As the CDN price competition becomes intense, people have already observed price wars [9]. However, there is no guideline or comprehensive studies about the best pricing strategy under CDN competition. Our paper aims to solve this problem. We propose a strategy which maximizes an individual CDN's utility. More importantly, the strategy is efficient in the sense that the total utility achieved is at least two thirds of the social optimal utility.

X. Conclusion

In this paper, we proposed Markov strategies so that each CDN can achieve the highest utility. The price competition is modeled as a repeated game and analyzed using Markov perfect equilibrium theories. We formally derived the sufficient conditions for Markov strategies to exist in two-CDN games, and we proved that the total utility is at least two thirds of the social optimal utility. Lastly, we showed that the pricing game can be generalized to more than two CDNs.

References

[1] *CDN prices.* `www.cdnpricing.com`.

[2] R. Buyya, M. Pathan, and A. Vakali. *Content Delivery Networks*. Springer Publishing Company, Incorporated, 1st edition, 2008.

[3] J. Chuang. Loci of competition for future internet architectures. *Communications Magazine*, 49(7):pp. 38–43, July 2011.

[4] D. Clark and S. Bauer. *Interconnection in the Internet: the policy challenge*. the 39th Research Conference on Communication, Information and Internet Policy, 2011.

[5] R. Ericson and A. Pakes. Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*, 62(1):pp. 53–82, 1995.

[6] P. Hande, M. Chiang, A. R. Calderbank, and S. Rangan. Network pricing and rate allocation with content provider participation. In *INFOCOM*, pages 990–998. IEEE, 2009.

[7] K. Hosanagar, J. Chuang, R. Krishnan, and M. D. Smith. Service adoption and pricing of content delivery network (cdn) services. *Manage. Sci.*, 54(9):1579–1593, Sept. 2008.

[8] K. Hosanagar, R. Krishnan, M. Smith, and J. Chuang. Optimal pricing of content delivery network (cdn) services. In *Proceedings of the Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04) - Track 7 - Volume 7*, HICSS '04, pages 70205.1–, Washington, DC, USA, 2004. IEEE Computer Society.

[9] O. Malik. *Akamai and the CDN Price Wars*. `http://gigaom.com/2007/08/06/cdn-price-wars/`.

[10] E. Maskin and J. Tirole. Markov perfect equilibrium, i: Observable actions. Harvard Institute of Economic Research Working Papers 1799, Harvard - Institute of Economic Research, 1997.

[11] A. Pakes. *IO Class Notes: Multiple Agent Dynamics; An Intro to Markov Perfect Equilibria*.

[12] A. Pakes and P. McGuire. Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *The RAND Journal of Economics*, 25(4):pp. 555–589, 1994.

[13] Y. Song, A. Venkataramani, and L. Gao. *On CDN selection game*. `https://sites.google.com/site/cdngametechreport/cdn`.

[14] J. Tirole and E. Maskin. A Theory of Dynamic Oligopoly I: Overview with Quantity Competition and Large Fixed Costs. *Econometrica*, 56:549–569, 1988.