

Improving deployability of peer-assisted CDN platform with incentive

Tatsuya Mori*, Noriaki Kamiyama*, Shigeaki Harada†, Haruhisa Hasegawa*, Ryoichi Kawahara*

* NTT Service Integration Laboratories, Midoricho 3–9–11, Musashino-shi, Tokyo, 180–8585 Japan
{mori.tatsuya, kamiyama.noriaki, hasegawa.haruhisa, kawahara.ryoichi}@lab.ntt.co.jp

† NTT West R&D Center, Shimaya 6–2–82, Konohana ward, Osaka-shi, Osaka, 554–0024 Japan
s.harada@rdc.west.ntt.co.jp

Abstract—As a promising solution to manage the huge workload of large-scale VoD services, managed peer-assisted CDN systems, such as *P4P* [25] has attracted attention. Although the approach works well in theory or in a controlled environment, to our best knowledge, there have been no general studies that address how actual peers can be incentivized in the wild Internet; thus, deployability of the system with respect to incentives to users has been an open issue. With this background in mind, we propose a new business model that aims to make peer-assisted approaches more feasible. The key idea of the model is that users *sell* their idle resources back to ISPs. In other words, ISPs can leverage resources of cooperative users by giving them *explicit* incentives, e.g., virtual currency. We show the high-level framework of designing optimal incentive amount to users. We also analyze how incentives and other external factors affect the efficiency of the system through simulation. Finally, we discuss other fundamental factors that are essential for the deployability of managed peer-assisted model. We believe that the new business model and the insights obtained through this work are useful for assessing the practical design and deployment of managed peer-assisted CDNs.

I. INTRODUCTION

A. Rising bandwidth scavenger and existing solutions

Recent growth in large-scale video sharing services such as YouTube [26] and MSN Video [18] have been significant. These services are estimated to facilitate hundreds of thousands of newly uploaded videos per day and support hundreds of millions of video views per day. The content available on these services are known as user-generated content (UGC) and are posted by end-users. Besides UGC, professionally generated video on demand (VoD) services, have also been successful. For example, in Feb 2009, a web video site Hulu [15] streamed about 18 million videos a day [10].

The great popularity of these VoD services has lead to a shift in Internet traffic mix. It is reported [8] that P2P traffic share dropped to 51% at the end of 2007, down from 60% the year before, and that the decline in this traffic share is due primarily to an increase in traffic from web-based video sharing services. It is also reported in 2008 [3] that statistics on one million broadband users in North America showed that HTTP traffic accounts for 46% of all broadband traffic while P2P applications now account for only 37%. We envision that this trend will potentially keep growing; thus, managing the huge volume of video traffic will continue to be a challenging task for both content providers and ISPs.

As a promising solution to address the huge workload on content providers, peer-assisted VoD systems, which leverage P2P technology, are now attracting attention. Recent studies on the world’s largest commercial VoD services, e.g., YouTube and MSN Video, concluded that peer-assisted solutions have a great advantage in offloading the workload of content delivery servers [5], [11], [14], [28]. Peer-assisted VoD delivery architectures have been widely studied [6], [20], [22]. In the commercial market, some content providers have been deploying peer-assisted video services [4].

The peer-assisted solutions are greatly successful in reducing the workload on content providers; however, they potentially cause an inevitable problem. That is, although P2P’s flexibility in choosing peers benefits the robustness and scalability of content distributing systems, if not carefully designed, it scatters traffic globally without consideration of the underlying network structure, resulting in an increase in cross-ISP traffic. Thus, peer-assisted solutions could impose *amplified* traffic on the external backbone links of ISPs.

To avoid this problem, it is crucial that ISPs localize P2P traffic. Managed peer-assisted approaches have been extensively studied in various ways [1], [2], [7], [16], [17], [25]. Xie et al. [25] recently proposed a fundamental framework, called P4P, that localizes P2P traffic by providing P2P applications with network information such as topology and link status. The goal of the P4P framework is to build interfaces to enable P2P applications and networks to communicate with each other and establish optimized peer selection. Choffnes et al. [7] demonstrated the effectiveness of biased peer selection, which leverages Content Delivery Network (CDN) redirection information, in reducing cross-ISP traffic. The problem of traffic localization is discussed in the broader problem space – the Application-Layer Traffic Optimization (ALTO), which addresses the way to provide the topology of the underlying network while allowing the requesting node to use such information to effectively reach the node that has content [17]. This problem is now being discussed in the IETF.

Although the *state-of-the-art* approach mentioned above, i.e., managed peer-assisted CDN (MP-CDN in short), works well in theory or in a controlled environment, to the best of our knowledge, there have been *no* general studies that address how actual peers can be incentivized in the wild Internet; thus, deployability of MP-CDN with respect to incentive has been

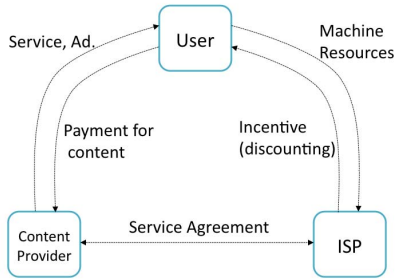


Fig. 1. Business model of the three players.

an open issue [21], [23]. The question we address is:

“What is the motivation for peer nodes to participate in the system and contribute their resources?”

Rodriguez et al. [23] pointed out that in addition to the current incentive mechanism in the layer of P2P applications, i.e., nodes are required to contribute to the system roughly as much data as they want to download, other form of incentives such as economic savings, should be explored in depth to establish feasible P2P CDN services. Rajahalme et al. [21] demonstrated that it is essential that caching incentives should exist to achieve *data-oriented* networking, where caching data plays a central role. We note that some applications, such as BBC iPlayer [4], embed uploading function implicitly, i.e., the resource of a user is agnostically leveraged in the system when the user is using the application. However, it is not clear that end-users will be willing to accept the implicit contribution in general.

B. Our solution: a new business model

To address the incentive problem in the MP-CDN framework, we propose a new business model for efficient and inexpensive delivery of large-scale traffic. The key idea of our model lies in the new incentive structure among the three key players, user, ISP, and content provider. Our objective is to encourage these players to work effectively by leveraging an incentive model that does *not* exist in the current Internet business model.

Figure 1 illustrates the business model among the three key players. An ISP can make use of user resources by giving them *explicit* incentives, e.g., giving virtual currency or discounting monthly fee, etc. An explicit incentive will clearly motivate users to contribute their resources as a part of the CDN system operated by an ISP. The ISP in turn can manage traffic flows by fully controlling the user resources with the knowledge of the exact information of peers and their files, network topology, routing matrix, and link congestion status.

We show how the model benefits the three players. First, end-users can get additional bonus by *selling* their idle resources. As a consequent, they can also enjoy the high quality of experience (QoE) and high availability of the content files that are cached on a large number of distributed peer nodes in the ISP. This idea may look like a business model of modern electrical grid system, where home owners install small power

generation systems such as solar panel or wind turbine, and they are allowed to sell the electricity back to the grid for the same price they would pay for it.

Second, leveraging user resources enables an ISP to establish efficient traffic management. Reducing cross domain traffic saves transit cost on the ISP’s external links. It also allows the ISP to control the timing of installing the new network line cards on the high-end boarder routers by cutting verbose cross-domain traffic. Traffic management also enables the ISP to avoid creating hot spots in its network and improve the QoE of the content delivery service. By offering high-quality and low-cost content delivery service, the ISP will be able to have new business opportunities with content providers with this model. Thus, the ISP can amortize the cost of incentive for users.

Finally, it enables a content provider to offload the workload on its origin servers because the ISP can localize many user requests within its network. Moreover, if a content provider intends to provide the services to users of the ISPs in a specific region, e.g., in a specific country, the content provider may be able to expect the large reduction in cost of using CDN platform because ISPs with MP-CDN need to prepare resources within their networks while conventional global CDN companies need to prepare the resources in all over the world.

Thus, the new business model can establish a “win-win-win” situation among the three players. Rather than imposing a *bandwidth cap* to control traffic from a bandwidth scavenger, which is a user or an application that sends out massive amount of traffic, as many ISPs are doing today, promoting the new incentive model will help evolve emerging Internet services and opportunities.

The remainder of the paper is structured as follows: In section II we show a model of MP-CDN we consider in this work. In section III, we show a high-level framework for designing optimal incentives while keeping a good balance of cost and revenue of an ISP. In section IV, we analyze the trade-off between system performance and cost for incentive through a simulation. we also study how other external factors affect the performance of MP-CDN. Section V discusses other potential factors that are important for the deployability of MP-CDN. Finally, section VI summarizes the work and discusses future work.

II. MODEL OF MP-CDN

In this section, we describe an overview of the basic MP-CDN system. Figure 2 illustrates the model of an MP-CDN. In the model, an ISP manages a server, which is called a PM (peer management) server. As we shall discuss later, the ISP can build multiple PM servers to avoid creating a single point of failure. A PM server plays a central role in the system and has the three fundamental functions, i.e., *cache management*, *accounting*, and *peer selection*.

cache management: As the central servers in centralized P2P networks, such as Napster [19] and eDonkey2000 [9], a PM server maintains a list of connected nodes and the files they keep in their storage. First, peer node A sends a data request

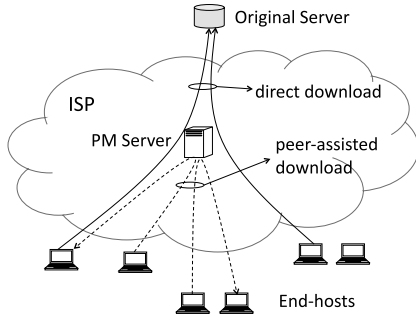


Fig. 2. Model of MP-CDN.

to the PM server. If PM server finds node B that has the file, the PM server will notify node A to grab data from node B. Otherwise, the request will be forwarded to the original server. Thus, the PM server provides a function of cache management, where the virtual cache server is built on top of storage devices of participating peer nodes. We note that achieving a good cache hit rate is the key success factor for the business model.

accounting: All peer nodes in an MP-CDN send the status of their connections to the PM server, e.g., completion of downloads/uploads. Thus, a PM server can keep track of the amount of contributions made by each peer. The incentive amount to each user can be determined based on the information about their contributions.

peer selection: Like the peer selection mechanism proposed by Xie et al. [25], when a PM server detects multiple nodes that have a content file requested by a peer, the PM server determines the peers following its own criteria; e.g., select the ones that can localize traffic the most. This mechanism allows the ISP to lower cross-domain traffic, which results in efficient use of the internal network.

III. DESIGNING OPTIMAL INCENTIVE

The primary goal of an ISP that deploys MP-CDN with incentive is to achieve a certain amount of cache hits, while spending the least for incentives. The MP-CDN will achieve good cache hit ratio if the ISP can increase the number of participating peers. On the other hand, the increase in the number of peers results in an increase in the cost of the incentives to peers. Thus, the two requirements are in a trade-off relationship.

In this section, we show a high-level framework of designing the optimal incentive amount to users under the trade-off. Through the framework, we attempt to demonstrate that an ISP is able to design optimal incentive under a certain condition. In the following analysis, we are interested in the macro property of the system rather than the micro behaviour of particular nodes or files. We assume that there is no cache capacity limitation and we are interested in an asymptotic but not transient status. More detailed analysis on caching property will be discussed in the next section.

Let N be the number of ISP users. The amount of total revenue for the ISP, R , is the summation of the ISP service and CDN service incomes. The former is from ISP users

for Internet access service while the latter is from a content publisher for outsourcing CDN service. Here, we assume that both incomes are given as a function of N ; i.e., we have $R = f(N)$.

We also assume that the cost for the ISP, C , is the summation of the bandwidth cost, C_b , cost of user incentives C_i , and remaining costs, C_r ; i.e., $C = C_b + C_i + C_r$. The ISP needs to ensure that the revenue is larger than cost to make the service sustainable; i.e.,

$$R \geq C_b + C_i + C_r. \quad (1)$$

Let $n(x)$ be the number of peer nodes who are willing to provide their resources as part of the P2P platform with the incentive of x . In general, $n(x)$ may be modeled as an increasing function of x . For simplicity, we assume that the ISP can use all $n(x)$ nodes as the CDN platform. That is, all the $n(x)$ nodes are available 24/7.

For simplicity, we assume that all traffic is generated by the CDN service. Then, the cost C_b is given as a function of b , which is the amount of bytes generated by the CDN service. The cost of incentives to users is the product of x and $n(x)$. To summarize, we have $C_b = g(b)$ and $C_i = xn(x)$.

The total amount of traffic in bytes carried over the CDN service, b , is given as the function of x , $b = (N - n(x) + 1)v_u$, where v_u is the average traffic volume carried to a user over a time period. Note that the amount of traffic decreased by peer nodes; i.e., once a peer node downloads a content, other $n(x) - 1$ peers download the content within the P2P network. This is localization of traffic. Note also that we assume the peer nodes are fixed in the time period.

From Eq. 1, we have the following inequality:

$$f(N) \geq g(b) + xn(x) + C_r. \quad (2)$$

Our goal is to find x that minimizes C under the condition of Eq. 2. Next, we explore the model with the simplest form of $g(b)$ and $n(x)$, i.e., $g(b) = \theta b$ and $n(x) = px^1$. The first Eq. assumes that the cost of bandwidth is proportional to the amount of bytes carried over the external link, and θ is the cost required to carry a byte. The second Eq. assumes that the number of cooperative users is proportional to the amount of incentives. We note the $n(x)$ should not exceed N , i.e., $x \leq N/p$. Eq. 2 becomes the simple quadratic inequality: $px^2 - pc_u x + (N + 1)c_u + C_r - f(N) \leq 0$, where $c_u = \theta v_u$ is the average bandwidth cost per user. Since the left term of the inequality is concaved upward, the solutions of the inequality exist only if $f(N) \geq -pc_u^2/4 + (N + 1)c_u + C_r$.

Under the above condition, the optimum incentive amount $x = x^*$ that minimizes the total cost is given as $x^* = c_u/2$. Thus, in this model, the optimum amount of incentive to a user is given by half of the average bandwidth cost per user.

We note that our objective here is to demonstrate a high-level framework of extracting optimal incentive amount under

¹We note that these assumptions are very naive. For example, the traffic cost model of an ISP can be stepwise functions instead of the linear function shown here. However, our goal is to show a high-level view of the economic structure using a simple model. We will discuss more sophisticated model for the acceptance of incentive in the next section.

a certain condition, rather than extracting a closed form of the optimum incentive amount in a specific condition. Once we fix the model of the system and other external factors, such as expected revenue and cost of the service, we can derive the optimal incentive in a closed form or numerically with the framework shown in this section.

IV. ANALYSIS OF CACHE PERFORMANCE

In this section, we study how incentives and other external factors affect the cache performance through a simulation. As we described in the previous section, getting good trade-off between incentive and cache performance, i.e., lowering cost of carrying traffic, is a key success factor for feasible deployability of MP-CDN.

We first describe the details of the MP-CDN model we use in our simulation. We then show the intrinsic trade-off between cache performance and cost of incentives. Finally, we study how the trade-off is affected by external factors such as the number of nodes, the number of distinct content files, and storage capacity for each node.

A. Model

Assumptions: We use the minimal model in the simulation (see Fig. 2). To focus our attention on the analysis of cache performance, which plays a crucial role in establishing efficient network utilization, let's assume that there are no network resource constraints, e.g. bandwidth, on the internal network entities. Thus, the peer selection in an ISP is performed based on the content possession information rather than the underlying network information. The objective of a PM server is simply to increase cache hits with least cost of incentives. We also make the following assumptions in the model:

- The properties of nodes are identical.
- Each node selects a content file according to the stretched exponential distribution [12]. (see below for more detail)
- The nodes are incentivized according to the univariate logit model. Only incentivized nodes will contribute to uploading their files. (see below for more detail)
- The incentive amount to users is fixed, regardless of the amount of contribution from each user.
- The nodes keep limited number of content files with the least frequently used (LFU) cache algorithm, which is managed globally by the PM server; i.e., the frequency information is provided by the PM server.
- The nodes do not send requests for files they have already downloaded and not removed.
- The number of nodes and content files are fixed, i.e., there are no churns in the simulation.
- The size of content files is fixed.

Although we could use more realistic and complicated model here, e.g., heterogeneous setting of node properties, we first focus our attention on the macro relationship between the incentive and the cache performance on the system. In our future work, we will study how heterogeneity of nodes could affect the workload on each user etc.

File access pattern: For the file selection model, we adopt the stretched exponential distribution, which can well capture the

access pattern of various modern media services such as web, YouTube, VoD, and IPTV, that benefit CDN architecture [12]. The probability distribution of the stretched exponential distribution is given as $\Pr[X < x] = 1 - \exp(-(x/x_0)^c)$, where c is the shape parameter and x_0 is the scale parameter. The smaller the shape parameter, c , the more biased the accesses to the most popular content files. The probability that i 'th media is selected, p_i , is given as $p_i = y_i / \sum_i y_i$, where $y_i = (-a \log(i) + b)^{1/c}$, $a = x_0^c$, and $b = y_1^c$. Following the statistics of CTVoD-04 discussed by Guo et al. [12], which is derived from a study on a large-scale VoD system deployed by China Telecom [27], we adopt the parameters of $a = 12.96$, $b = 118.9$, and $c = 0.4$.

Acceptance of incentive: For the incentive models, we used the logit model. The logit model is a well known approach in various field for modeling the probability of a customer accepting a new product or technique [24]. The advantage of adopting the logit model lies in its simplicity. Once we have the survey data on users, we can easily perform preference regression on the model with explicit measures of confidence level and significance.

Given incentive of x , a node becomes cooperative to the system with the probability of $p(x)$, such that $p(x) = 1/(1 + e^{-(\beta_1 + \beta_2 x)})$, where $\beta_1 = -\log(1/\xi - 1)$, $\beta_2 = \log(1/\xi - 1)/\alpha$, and ξ is the probability that a node is cooperative without having incentive (i.e., $p(0) = \xi$). α is the mean of x required to make nodes cooperative. In this work, we use $\xi = 0.01$ and $\alpha = 10$.

Simulation setup: We performed 10 experiments with independent random seeds for each parameter setting. Let N be the number of end-hosts, m be the number of content files, and S be the storage capacity for each node with respect to the number of files (note that the file size is fixed in this work), respectively. The end-hosts independently generate requests with exponentially distributed inter-arrival time. The mean inter-arrival time is set to 72. Note that a unit of simulation time is roughly assumed to be five minutes in real time; i.e., the mean inter-arrival time is set to 6 hours in real time. Similarly, the time to finish downloading/uploading a file is set to 6, which is 30 minutes in real time. Note that we assumed that the size of content files is fixed and there is no bandwidth bottleneck in the network. To see the asymptotic property of the system, the simulation time was set to 10000, which roughly corresponds to a month in real time. Figure 3 shows an example of a simulated path.

B. Trade-off between cache performance and cost of incentive

We studied the trade-off between the cache performance and cost of incentives. Figure 4 shows the results. The parameters were set to $N = 100$, $m = 1,000$, $t = 10,000$, and $S = 5$. Figure 4 (a) shows the number of incentivized users, $N(x)$, with the incentive per user, x . Obviously, the shape of $N(x)$ comes from the sigmoidal function of the logit model. Accordingly, the cache hit ratio shown in Fig. 4 (b), which is defined as the number of cached downloads divided by the total number of downloads, behaves in a similar manner. That is, the increase in the number of incentivized peers corresponds to the increase in the storage space that can be used for caching.




Fig. 3. Asymptotic property of system behaviour. Parameters are set to: $N = 100$, $m = 1000$, and $S = 5$. Top: Number of total downloads; Middle: Number of cached downloads; Bottom: Cache hit ratio, i.e., number of cached downloads divided from total number of downloads.

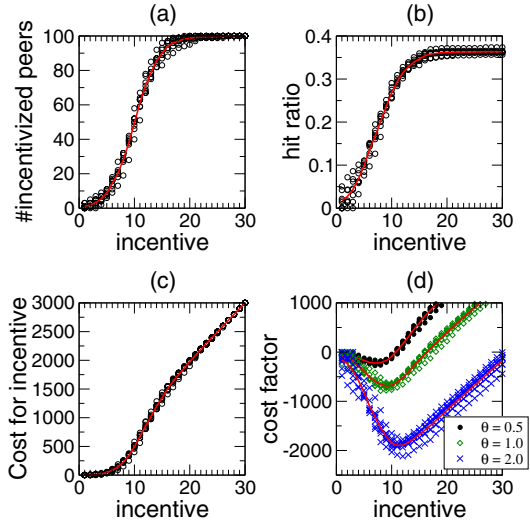


Fig. 4. Effect of incentive: (a) Number of incentivized peers, (b) Cache hit ratio, (c) Total cost of incentive, and (d) Cost factor. Circles indicate each simulation run (10 for each parameter setting) while lines indicate the mean values.

Although the behaviour of $N(x)$ and the cache hit ratio looks similar, we note that the relationship between the two is not trivial. The reason is as follows. Each node independently requests files and manages the cache space with the LFU algorithm, following the perfect file access frequency information provided by the PM server. Thus, the globally managed cache can keep *duplicated* files in the aggregated cache space. Therefore, the established cache performance by the MP-CDN model is different from the one for the LFU cache deployed in a single server, for which case, the cache does not keep the duplicated files and the theoretical cache performance can be calculated from the file access pattern and the available storage space. We also notice that after all the peers are incentivized, i.e., when $x \geq 20$, the cache performance will no longer be improved.




Fig. 5. Effect of the number of total nodes in the system: Left: Cache hit ratio; Right: total cost of incentive. Circles represent each simulation run (10 for each parameter setting) with each line representing the mean values.

Next, Fig. 4 (c) shows the total cost of incentive, C_i , i.e., $C_i = xN(x)$. The total cost increases non-linearly when $x < 20$ and increases linearly otherwise. Figure 4 (b) and (c) show the trade-off between the cache performance and cost of incentive. To see this trade-off more clearly, Fig. 4 (d) shows how the cost factor is affected by incentive. The cost factor C_f is defined as $C_f = C_i - \theta T_c$, where T_c is the amount of traffic saved by the cache, θ is the weight parameter, and T_c is given by the number of cached files in the system. As we showed in section III, these factors are essential for the balance of total revenue and cost. We can see that for each weight, there exist an amount of incentive that minimize cost factor C_f .

C. External factors and cache performance

This section studies how other external factors such as the number of nodes, the number of distinct content files, and storage capacity for each node affect the cache performance. The total cost of incentive under the conditions is also explored.

Number of participating nodes: Figure 5 shows the mean cache hit ratios under various number of nodes, N ($0 \leq N \leq 10^4$), in the system. The amount of incentive per user is set to $x = 10$, the number of content files is set to $m = 1000$, and storage capacity per node is set to $S = 5$. Both cache hit ratio and cost of incentive increase as the number of participating nodes gets higher. Obviously, total cost of incentive is proportional to the number of nodes. In contrast, the cache hit ratio gets higher rapidly when the number of nodes is fairly small, say, $N \leq 1000$. For example, for the number of nodes $N = 500$ and $N = 1000$, the mean cache hit ratios are 0.53 and 0.67, and total cost of incentive are 2459 and 4957, respectively. In this case, 14% of improvement in cache hit ratio needs to double the cost of incentive. These properties play an important role for an ISP in balancing the scale of the users (nodes) and total cost of incentive. We note that an ISP is able to control the number of *incentivized* nodes by adjusting the amount of incentive per user as well.

Number of content files: Figure 6 (a) shows the mean cache hit ratios under various number of content files, m ($1 \leq m \leq 10^4$). Here, the amount of incentive per user is set to $x = 10$, the number of nodes is set to $N = 100$, and storage capacity per node is set to $S = 5$.




Fig. 6. Effect on cache performance: the number of distinct content files in original server (left) and the amount of storage capacity, S per user (right). $S = 0$ stands for the infinite amount of storage space. Circles indicate each simulation run (10 for each parameter setting) with each line representing the mean values.

Note that the total cost of incentive is independent of m . As we can see, if original server has a lot of distinct content files, the cache hit ratio gets worse under a fixed amount of storage space per user, i.e., $S = 5$ in this case. That is, when there are a large number of distinct content files, we have a large number of frequently accessed files as well. Therefore, if the total amount of storage capacity is smaller than the number of these popular files, a probability of cache miss increases. Thus, to achieve high cache performance under a given storage amount, it is an effective strategy to keep the number of available content files fairly low within an interval. We also notice that cache performance gets bit worse when m is very small, say, $m = 1$. This comes from the fact that it is assumed that nodes do not request files already downloaded by themselves before. Thus, the number of total downloads gets smaller and the probability that a request hits the cache becomes smaller. We note that the degradation is negligible and we do not need to use MP-CDN when m is quite small.

Storage capacity of users: Figure 6 (b) shows the mean cache hit ratios under various storage capacity per user, S ($0 \leq S \leq 30$), where we denote infinite capacity as $S = 0$. The amount of incentive per user is set to $x = 10$, the number of nodes is set to $N = 100$, and the number of content files is set to $m = 1000$. As is expected, the cache performance gets better as the increase in S . The good news here is that the cache performance with infinite storage space can be established with fairly small storage space, say $S = 20$. Assuming that the size of each content file is 2 GB (for example, the size of a 2-hours movie file encoded at the mean bit rates of 2,000 Kbps with a frame size of 640×480 pixels can be 2 GB), peer nodes need to reserve storage capacity of 40 GB, which may be affordable amount of resource on a commodity PC today. This good property reflects the potential power of peer-assisted CDN, i.e., even though the resource of each node is not large, the aggregated resource can become very large.

V. DETAILED DISCUSSION

Finally, we discuss other potential factors that could significantly affect the deployability of MP-CDN services.

A. Scalability

An MP-CDN establishes a content delivery mechanism that is efficient in terms of reducing workload on servers and network traffic across ISPs. Participating users benefit from high performance, availability, and rewards from the ISP, while the ISP benefits by cutting the cost of network traffic. Content providers also benefit from the lowering of server workload. What is the compensation for these benefits? The trick is in the role of the PM server, which acts as an *oracle* in the network. That is, the PM server keeps track of nodes, content files, and underlying network information such as network topology, routing information, and link statistics. The use of a centralized PM server that indexes all content might introduce a single point of failure. Therefore, an MP-CDN success lies in the scalability of the PM server.

ISPs need to deploy a PM server as essential infrastructures, such as DNS cache servers or web servers, which also have a risk of being a single point of failure if they are managed inappropriately. There are a large number techniques that establish the high scalability of large-scale servers: load balancing, high-performance computing, and DHT. We will need to leverage some of these approaches to make MP-CDN scalable. Optimizing database design, processing, and configuration are also challenging and crucial tasks.

We finally note that the system is especially well suited to supporting downloads of large files, so that the number of downloads supported by PM server may not be that large although the total number of bytes delivered by each peer node may be very large.

B. Underlying network infrastructure

One assumption taken for granted in this paper is that transferring of data between peers within an ISP is a traffic pattern that the ISP's infrastructure is designed to efficiently support. It may well be that once an ISP goes through the effort of establishing PM servers, it may place less stress on the network to deliver all content from these servers to end users rather than from peer to peer. However, if a network has been designed to provide more download capacity to end users than upload capacity, e.g., CATV access network, then the uploading of files by peers might cause congestion in the network. We believe that as far as an ISP can take control over selecting peers with the *exact* knowledge of its network property and status, traffic optimization should work effectively, given an underlying network infrastructure. For example, we can assign pair of peers so that traffic flows do not create hot spots in the network. We can also adaptively utilize caching servers if we see congestion on upload links. Moreover, as Fiber to the premises/home (FTTP/FTTH) services become popular in near future, the problem of asymmetric network resources will be naturally solved.

C. Privacy

As with many other web caching systems, an MP-CDN operator has to make sure to protect the privacy of users. In general, network administrative organizations have a clear

privacy policy for the use of web caching systems. The contents and logs on the cache servers will be properly managed, following the privacy policy. Here is a question we want to ask: How can we make sure that privacy is protected on the end-nodes? Unlike network operators, end-users do not have the mapping of IP addresses and actual user identities in general. Thus, the privacy problem should be somewhat relaxed at this level. However, continued accesses from a fixed node could expose intrinsic features that can be associated with a specific user. To cope with this potential problem, an MP-CDN can use probability to introduce weak randomness in the peer selection mechanism, i.e., given a set of nodes with the same network cost, a PM server will randomly pick one from the candidates so that privacy is kept among end-users.

D. Toward more efficient content delivering

Like many other P2P solutions, an MP-CDN can also benefit from the strategy of splitting large files into chunks, allowing nodes to retrieve different chunks from different nodes. When a PM server receives a request of a content file from node A, the PM server can reply to node A that node A can download pieces of the content file from a set of multiple nodes. Although this strategy of pipelining will potentially accelerate the efficiency of downloads, it could also significantly increase the number of transactions between nodes. As Harma and Felber [13] suggested, the trade-off between the number of transactions and the size of chunks should be carefully designed based upon the required goal and resources. We will leave this as our future work.

VI. CONCLUSION AND FUTURE WORK

We presented a new business model that aims to make an MP-CDN more feasible. The key idea is to allow ISPs to provide end-users with *direct* incentives, which may encourage the users to contribute their resources. The business model enables a “win-win-win” situation among the three key players, user, ISP, and content provider. That is, the new approach relaxes the intrinsic constraints of the current Internet business model. We believe that the new model will shed new light on emerging Internet services and opportunities not only in business but also in technical fields. We presented a high-level framework for designing an optimal incentive amount for users. We also researched into the intrinsic trade-off between the performance of system and cost for incentive, and the other external factors on the cache performance through a simulation study. Finally, we discussed other potential factors that are essential for deploying managed peer-assisted CDNs.

Our future work includes: 1. enhancing the analysis for more realistic cases, i.e., heterogeneity of peer characteristics and internal network topology, having constraints on network resources, and more realistic model of traffic cost of an ISP, e.g., the stepwise functions, 2. formulating cache performance of an MP-CDN mathematically, and, 3. studying the behaviour of actual users and system performance, using actual implementation in a testbed network.

REFERENCES

- [1] V. Aggarwal, O. Akonjang, and A. Feldmann. Improving user and isp experience through isp-aided p2p locality. In *11th IEEE Global Internet Symposium 2008*, Apr. 2008.
- [2] V. Aggarwal, A. Feldmann, and C. Scheideler. Can ISPS and P2P users cooperate for improved performance? *ACM SIGCOMM Computer Communication Review*, 37(3), July 2007.
- [3] N. Anderson. The YouTube effect: HTTP traffic now eclipses P2P. <http://arstechnica.com/news.ars/post/20070619-the-youtube-effect-http-traffic-now-eclipses-p2p.html>, June 19, 2008.
- [4] BBC iPlayer. <http://www.bbc.co.uk/iplayer/>.
- [5] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon. I Tube, You Tube, Everybody Tubes: Analyzing the World’s Largest User Generated Content Video System. In *ACM IMC 2007*, October 2007.
- [6] B. Cheng, X. Liu, Z. Zhang, and H. Jin. A Measurement Study of a Peer-to-Peer Video-on-Demand System. In *International Peer to Peer Symposium (IPTPS)*, Feb. 2007.
- [7] D. R. Choffnes and F. E. Bustamante. Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in P2P Systems. In *ACM SIGCOMM 2008*, Aug. 2008, to appear.
- [8] Cisco Systems, Inc. Cisco Visual Networking Index - Forecast and Methodology, 2007-2012. http://newsroom.cisco.com/dlls/2008/ekits/Cisco_Visual_Networking_Index_061608.pdf, June 2008.
- [9] eDonkey2000. <http://en.wikipedia.org/wiki/EDonkey2000>.
- [10] D. Frommer. Hulu Traffic Still Up Big After Super Bowl Spike, Mar 2, 2008. <http://www.businessinsider.com/hulu-traffic-still-up-big-after-super-bowl-spike-2009-3>.
- [11] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. Youtube traffic characterization: a view from the edge. In *ACM IMC 2007*, pages 15–28. ACM, October 2007.
- [12] L. Guo, E. Tan, S. Chen, Z. Xiao, and X. ong Zhang. The stretched exponential distribution of internet media access patterns. In *27th ACM Symposium on Principles of Distributed Computing (PODC 2008)*, Aug. 2008.
- [13] A. A. Hamra and P. A. Felber. Design choices for content distribution in P2P networks. 35(5):29–40, 2005.
- [14] C. Huang, J. Li, and K. W. Ross. Can internet video-on-demand be profitable? In *ACM SIGCOMM 2007*, pages 133–144, Aug. 2007.
- [15] Hulu. <http://www.hulu.com>.
- [16] T. Karagiannis, P. Rodriguez, and K. Papagiannaki. Should internet service providers fear peer-assisted content distribution? In *ACM IMC 2005*, October 2005.
- [17] E. Marocco, V. K. Gurbani, V. Hilt, I. Rimac, and M. Tomsu. Peer-to-peer infrastructure: A survey of research on the application-layer traffic optimization problem and the need for layer cooperation. In *IETF Peer-to-peer infrastructure workshop*, May 2008.
- [18] MSN Video. <http://video.msn.com>.
- [19] Napster. <http://en.wikipedia.org/wiki/Napster>.
- [20] P2P-Next. <http://www.p2p-next.org/>.
- [21] J. Rajahalme, M. Sarela, P. Nikander, and S. Tarkoma. Incentive-compatible caching and peering in data-oriented networks. In *ACM CoNEXT ReArch*, Dec 2008.
- [22] K. Ramakrishnan. Architecture for IPTV Distribution: Cooperative P2P and Multicast. In *ACM SIGCOMM 2007 Peer-to-Peer Streaming and IP-TV Workshop*, Aug. 2007.
- [23] P. Rodriguez, S.-M. Tan, and C. Gkantsidis. On the feasibility of commercial, legal p2p content distribution. *SIGCOMM Comput. Commun. Rev.*, 36(1):75–78, 2006.
- [24] C. J. S. *Logit Models from Economics and Other Fields*. Cambridge University Press, 2003.
- [25] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz. P4P: Provider Portal for Applications. In *ACM SIGCOMM 2008*, Aug. 2008, to appear.
- [26] YouTube. <http://www.youtube.com>.
- [27] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng. Understanding user behavior in large-scale video-on-demand systems. *SIGOPS Oper. Syst. Rev.*, 40(4):333–344, 2006.
- [28] M. Zink, K. Suh, Y. Gu, and J. Kurose. Watch Global, Cache Local: YouTube Network Traces at a Campus Network - Measurements and Implications. In *IEEE MMCN 2008*, 2008.