



Energy Consumption in Peer-Assisted CDN

Mohamad Dikshie Fauzie
Graduate school of Media and Governance
Keio UniversityFujisawa, Japan
dikshie@sfc.wide.ad.jp

Achmad Husni Thamrin
Graduate School of Media and Governance
Keio UniversityFujisawa, Japan
husni@ai3.net

Jun Murai
Faculty of Environment and Information Studies
Keio UniversityFujisawa, Japan
jun@wide.ad.jp

Abstract: Many CDN companies utilize peer-to-peer to scale content delivery services and save bandwidth, but the energy consumption in a peer-assisted CDN is not well understood. This paper presents an analysis of energy consumption in a peer-assisted CDN system, where an ISP manages its own CDN and its users participate in a P2P network to assist content delivery, by looking into two types of services: live streaming and online storage. We developed a simple energy consumption model for peer-assisted CDN based on the empirical energy required for devices and the fact that data centers require cooling overhead and peers do not. For the live streaming service, we found that delegating workloads to peers can save up to 11% of the energy in the data center compared to the pure CDN architecture, while the savings in the total energy consumption of the system is less than 1%. For the peer-assisted online storage, we looked into several server bandwidth allocation strategies for peers of a peer-assisted online storage service and in the best case found that the energy savings of the data center and of the system are 21% and less than 2%, respectively.

Keywords: P2P, CDN, Energy, Power

I. INTRODUCTION

Streaming content, especially video, represents a significant fraction of the traffic volume on the Internet, and it has become a standard practice to deliver this type of content using Content Delivery Networks (CDNs) such as Akamai and Limelight for better scaling and quality of experience for the end users. For example, YouTube uses Google cache and MTV uses Akamai in their operations.

With the spread of broadband Internet access at a reasonable flat monthly rate, users are connected to the Internet 24 hours a day and they can download and share multimedia content. P2P (peer to peer) applications are also widely deployed. In China, P2P is very popular; we see many P2P applications from China such as PPLive, PPStream, UUSE, Xunlei, etc. [25]. Some news broadcasters also rely on P2P technology to deliver popular live events. For example, CNN uses the Octoshape solution that enables their broadcast to scale and offer good video quality as the number of users increases.

From the Internet provider point of view, the presence of so many always-on users suggests that it is possible to delegate a portion of computing, storage and networking tasks to the users, thus creating P2P networks where users can share files and multimedia content. Starting from file sharing protocols, P2P architectures have evolved toward video on demand and support for live events.

Broadband network access helps P2P applications to perform better. xDSL networks are deployed worldwide, and in some countries, such as Japan, even higher bandwidth fiber to the home (FTTH) already exceeds DSL in market

penetration. In the coming years, network operators throughout the world will massively deploy FTTH. As access bandwidth increases, P2P systems may become more efficient since a peer can contribute much more.

In Peer assisted CDN¹, users can download content from CDN nodes from or other users or peers. A user may cache the content after download to serve requests from other users. Due to the complexity of the behavior of peers, the process should be done in the home gateway user where the ISP can control it.

On the other hand, the data center where the CDN server is placed faces costs for powering the data center. The Uptime Institute, a global data center authority, surveyed 1100 data center owners and operators in 2012 and reported that 55% of organizations will increase their financial budget 10% over 2011 [1]. 30% of organizations were expected run out of data center capacity (power, cooling, space, and network) by the end of 2012 [1]. More than 50% of the organizations surveyed reported that saving energy² is a major priority. Even in the data centers using the state of art cooling technologies heat dissipation accounts for at least 20% and as much as 50% of the total power consumption [2]. The increases in energy cost and the demand due to growth of traffic urges the data center operators and owners to look for ways to reduce energy usage in the years to come. Although reducing energy consumption can effectively reduce overall cost, this will limit the capacity for growth and scalability of the service provisioning. For example: routers and servers spend most of their energy on the baseline activities such as running the fans, spinning disk, powering the backplane, and

¹In this paper we use Peer Assisted CDN and CDN-P2P interchangeably.

²As we are discussing steady-state operation, energy and power are in direct correspondence so we use the terms interchangeably.

powering the memory. Even in an idle state, modern systems can be consuming anything from 50% to 80% of the power consumed under maximum load [3][4].

Alternatively, the data center can be revamped by relocating some services to end-host computers or peers. Peers contribute their communication, storage, and computation resources to exchange data and provide services while the data center performs central administration and authentication as well as backend processing. A P2P network formed by peers offers flexibility and scalability in service delivery.

In this paper, we study the energy consumption of hybrid CDN-P2P in two use cases: live streaming and online storage services. It has been shown that CDN energy consumption is better than P2P architecture [5][6]. The questions are: with the opportunity to offload the CDN's workload to the peers, how much energy saving can the CDN server get and how large is the difference compared to a pure CDN architecture. If we can estimate the difference between a CDN architecture and a peer-assisted CDN combined with an estimate of peer power consumption, we can use this difference as a basis calculation for giving an incentive to users since peer assisted relies heavily on the user's uptime and upload rate. Furthermore, since the power consumption is reduced, the power requirement inside the data center can be reduced thus relaxing capacity planning.

The rest of this paper is organized as follows. Section II presents system description. Section III provides result and analysis. In section IV, we will present related work. Section V will conclude this paper.

II. SYSTEM DESCRIPTION

A. Live Streaming:

Fig. 1 shows an example model of a peer-assisted CDN for live streaming adapted from [9]. CDN servers deliver video contents from the content provider to end-users. The CDN usually is organized into several tiers usually to cope with the scale demand. Edge CDN servers are directly responsible for serving end users. The goal of the server side peer is for efficient data distribution with some measures to guard against some node failures and network delay.

The CDN overlay is largely tree-based. To provide greater reliability, a CDN node may allow retrieving the content either from other nodes. Edge CDN servers are responsible for serving end users.

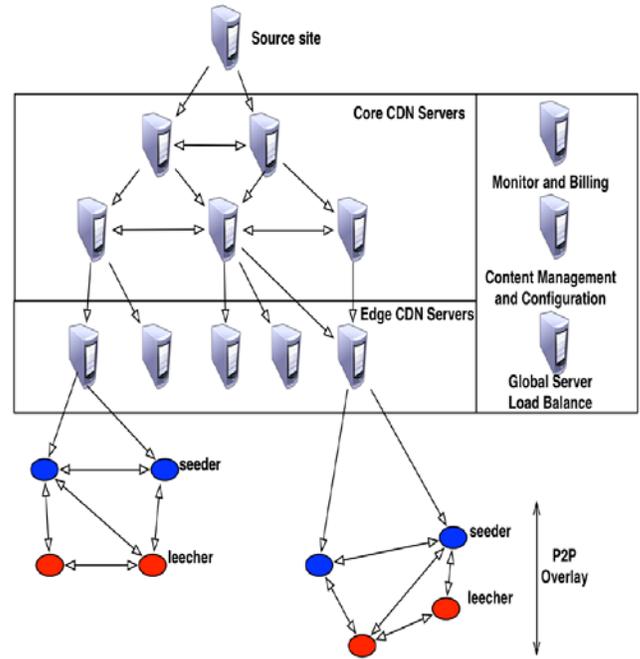


Figure 1. Example model of peer-assisted online storage architecture.

For this system, we introduce the concepts of seeder and leecher. A peer that is served by an edge CDN server is called a seeder while a peer that is served by seeders is called a leecher.

A peer obtains the URL from a content source. The global server load balancer finds a suitable edge CDN node for this peer. The peer is then redirected to the nearest edge CDN. The edge CDNs has decision logic that decides if a new peer should be served directly by the edge CDN or if it should be redirected to the P2P overlay.

In the P2P overlay, the stream is separated into several substreams according to the stream id and peers are organized in a tree-based overlay. A working peer-assisted CDN live streaming system is defined by parameters such as: (1) video bitrate, (2) the total number of peers, (3) the edge CDN servers bandwidth, and (4) peer upload bandwidth capacity and churn rates.

The maximum number of seeders is bounded by the CDN's capacity, while the maximum number of leechers is bounded by the number of seeders with a certain upload rate. Let us denote the number of the seeders by n_s , the number of leechers by n_l , the maximum bitrate supplied by seeders to leechers by ρ , and the video bitrate by r . The number of leechers that can be supported by seeders is:

$$n_l = n_s \cdot \rho \tag{1}$$

The number of seeders that support or upload content to leechers is:

$$n_s^u = n_l \cdot \frac{r}{\rho} \tag{2}$$

In peer-assisted live streaming, we introduce the utilization policy where the CDN server admits peers as seeders as long as the CDN utilization does not exceed 50%, which we defined as 50% of the capacity of a Gigabit Ethernet. When the utilization hits 50%, incoming peers are admitted as leechers, hence they receive the contents from seeders. When more peers join the system and the upload capacity of the seeders is exceeded, the

policy raises the utilization cap and the server admits the newly joined peers as seeders. We consider this policy to be better than adding a new server from the point of view of energy consumption.

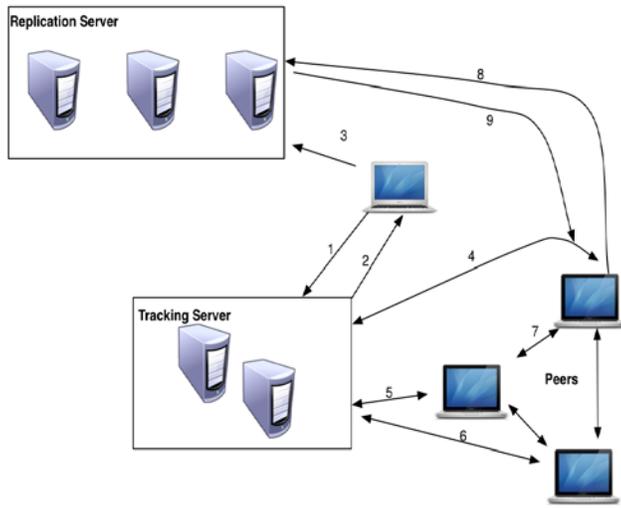


Figure 2. Example model of peer-assisted online storage architecture.

B. Peer-Assisted Online Storage:

Fig. 2 illustrates the architecture of peer-assisted online storage for a file hosting system (one-click hosting service with peer-assistance) and interactions among the main components [10]. In this system, each file provided by users is treated as a *swarm*. Each end user is treated as a peer.

In Fig.2, arrows 1,2,and 3 denote the interaction between a participating peer and tracking server and replication servers for uploading a new file. Arrows 4,5, and 6 denote the interaction between peers and the tracking server to maintain the peer topology. Arrow 7 denotes the sharing of the file and exchange of availability data among peers. Arrows 8 and 9 represent peer requests and server response.

The tracking server's function is to maintain swarm information and bootstraps peers. Replication servers working as dedicated content servers have a function for maintaining the availability of swarms when peers do not actively serve them alone.

We choose this peer-assisted online storage model because this model has been implemented widely in China, e.g. FS2You [29], and because one-click file hosting services are very popular right now [28]. Such services rely heavily on server farms inside the data center, thus energy cost becomes important [27]. In this model, since the server holds an important role in this system, we present a simple mathematical model of server bandwidth allocation strategies [11][12] as a basis for energy calculations, as follows:

Type-1 represents less popular files and type-2 represents popular files.

S_{i1} represents server bandwidth allocated to a type-1 file and S_{i2} represents server bandwidth allocated to a type-2 file.

S is the total server bandwidth.

S_{max1} is the maximum amount of server bandwidth that can be assigned to a file of type-1 and S_{max2} is the maximum

amount of server bandwidth that can be assigned to a file of type-2.

M_{t1} is the number of type-1 files and M_{t2} is the number of type-2 files.

μ is upload rate of a peer.

λ_{t1} is the arrival rate of new peers in type-1 file and λ_{t2} is the arrival rate of new peers in type-2 file.

$$\lambda = M_{t1}\lambda_{t1} + M_{t2}\lambda_{t2}$$

$$M = M_{t1} + M_{t2}$$

η_{t1} is the file sharing effectiveness. It is the fraction of the upload capacity of peers that is being utilized for type-1 file.

η_{t2} is the file sharing effectiveness. It is the fraction of the upload capacity of peers that is being utilized for type-2 file.

T_d is the average downloading time.

x_i is the average number of peers.

There are three server bandwidth allocation strategies: (1) lower bound of the average downloading time; (2) request driven strategy; (3) water leveling strategy. In the lower bound strategy, the server uses the bandwidth for type-1 files until S_{i1} reaches its maximum value, then the residual server bandwidth can be assigned to type-2 files. In the request driven strategy, the server serves every request from peers. The server bandwidth is equally divided among all the peers. Let's assume that the number of requests for a file to the server is proportional to the peer arrival rate of the file. Let's also assume that when the amount of server bandwidth assigned to one of the two types of files has reached its maximum value, the residual server bandwidth will be assigned to the other type of file. In the water leveling strategy, the server bandwidth is equalized across all the files by taking file popularity into consideration. The server serves the requests from peers according to a certain probability, which is inversely proportional to the peer arrival rate of the file. Let's assume that the number of requests for a file to the server is proportional to the peer arrival rate of the file, the server will serve the same number of requests for different files and therefore the server bandwidth is equally allocated across all the files. In order to be able to calculate our power consumption, we need to get the number of peers in the system that can be expressed as [12]:

$$\sum x_i = T_d \sum \lambda_i \quad (3)$$

Furthermore, we can calculate T_d :

$$T_d = \frac{1}{M_{t1}\lambda_{t1} + M_{t2}\lambda_{t2}} \left(\frac{M_{t1}f_{t1}\lambda_{t1}\eta_{t1} + M_{t2}f_{t2}\lambda_{t2}\eta_{t2}}{\mu\eta_{t1}\eta_{t2}} - \frac{S_{t1}(M_{t1}\eta_{t2} - M_{t2}\eta_{t1}) + SM_{t2}\eta_{t1}}{\mu\eta_{t1}\eta_{t2}} \right) \quad (4)$$

C. Energy Model:

In this paper, our goal is to provide a general view and fair comparison of the energy consumed by a pure CDN and a hybrid CDN-P2P architecture. To do so, we designed a series of models and performed an analysis. Our energy model is similar to the models used in [13]. The differences with [13] are, firstly, our baseline energy is not a function of bitrate flow. Our baseline energy is based on the minimum energy required to turn on the device without any traffic flowing through the device. Secondly, our overhead ratio is based on the Coefficient of Performance (COP) of the cooling cycle in data center, which we will explain at the end of this section.

Let $E_s, E_r,$ and E_p denote the energy consumption of a single request at each a CDN server, router, and peer respectively. Next, we define baseline energy consumption as the energy consumed to keep the device on, even when there is no traffic. Let $E_{s-base}, E_{r-base},$ and E_{p-base} denote the baseline energy consumption for CDN server, router, and peer respectively; and $E_{s-max}, E_{r-max}, E_{p-max}$ denote the power consumption of server, router, and peer when operating at the maximum capacity.

Next, we introduce work-induced energy as the energy consume per bit transferred. Let $\delta_s, \delta_r,$ and δ_p denote the work-induced energy consumed per additional bit transferred by each CDN server, router, and peer,

$$\delta_s = \frac{(E_{s-max} - E_{s-base})}{M_s} \quad (5)$$

$$\delta_r = \frac{(E_{r-max} - E_{r-base})}{M_r} \quad (6)$$

$$\delta_p = \frac{(E_{p-max} - E_{p-base})}{M_p} \quad (7)$$

Furthermore, we can get:

$$E_s = \delta_s B + E_{s-base} \quad (8)$$

$$E_r = d\delta_r B + E_{r-base} \quad (9)$$

$$E_p = \delta_p B + E_{p-base} \quad (10)$$

Where d is the number of hops and B is the size of file to be transferred in bits.

We now introduce the overhead for the server and routers. The only overhead that we will consider here is cooling power. Since servers and routers are placed in the data center, the data center needs to be provisioned with adequate cooling. This cooling overhead in the data center is quantified by the coefficient of performance (COP). The COP value itself has been empirically determined to be [14]:

$$COP(T) = 0.0068T^2 + 0.0008T + 0.458 \quad (11)$$

Where T is the temperature supplied by the cooling unit in C. Finally, the cooling cost can be calculated as [14]:

$$C = \frac{Q}{COP(T)} \quad (12)$$

Where Q is the amount of power consumed by the servers and hardware. We assume a uniform T at each cooling unit. Taking into account the cooling energy overhead, the total energy consumption is as follows:

$$E_t = E_s \left(1 + \frac{1}{COP(T)}\right) + E_r \left(1 + \frac{1}{COP(T)}\right) + E_p \quad (13)$$

We do not include the cooling overhead in the peer energy consumption because most of the peers in homes do not need a separate cooling supply.

III. RESULT AND ANALYSIS

A. Numerical Parameters:

The parameters used in this analysis were adapted from [4][8][12], and [13]. The parameters values are shown in Table I. We choose the numerical parameters from [4][8][12], and [13] because these parameters were gathered from empirical measurements.

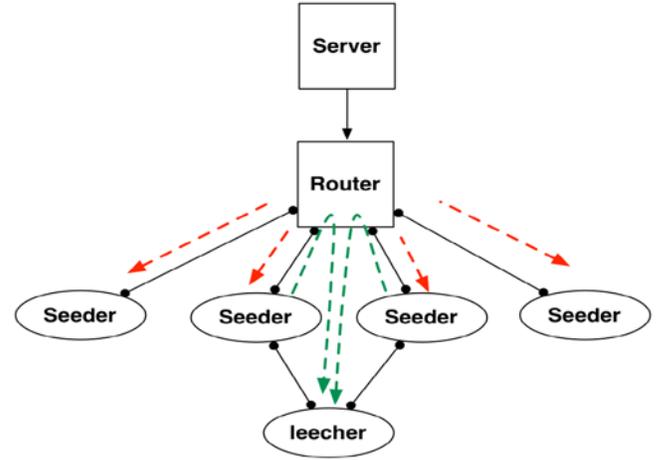


Figure 3. Simplified physical representation flow of data.

Table I. Numerical Parameters from [4][8][12][13]

Symbol	Description	Values
δ_s	Work induced at server per bit transferred	$5.2 \cdot 10^{-8}$ (J/b)
δ_r	Work induced at router per bit transferred	$8.0 \cdot 10^{-9}$ (J/b)
δ_p	Work induced at peer per bit transferred	$16 \cdot 10^{-9}$ (J/b)
E_{r-base}	Router baseline power consumption	750 watts
E_{s-base}	Server baseline power consumption	291 watts
E_{p-base}	Peer baseline power consumption	13.5 watts
r	Video bitrate in live streaming	1 Mbps
d_s	Number of hops	1
N_s^{st}	Upload rate of peers in live streaming	[0.25, 0.5, 0.75, 1] Mbps
N	Number of peers in live streaming	[100, ..., 1000]
λ_{t1}	Type-1 peer arrival rate to less popular files in online storage (Poisson process)	0.1
λ_{t2}	Type-1 peer arrival rate to popular files in online storage (Poisson process)	1
η_{t1}	File type-1 sharing effectiveness. The fraction of the upload capacity of peers that is being utilized in online storage	0.5
η_{t2}	File type-2 sharing effectiveness. The fraction of the upload capacity of peers that is being utilized in online storage	1

M_{t1}	Number of files in type-1 files or less popular files	10
M_{t2}	Number of files in type-2 files or popular files	1
$f_{t1} = f_{t2}$	File size in online storage	100 MB
μ	Upload rate of peers in online storage	0.5 Mbps
c	Downloading rate of peers in online storage	1Mbps
T	Air temperature supplied from cooling unit in data center	[20,25] correspond to COP values [3.194, 4.728]

B. Live Streaming Service:

In the live streaming case, each video stream flows from the CDN node to the network, in this case a router, and then arrives at the seeders. If seeders need to upload data to leechers it will flow through the router then arrive at the leechers. We show the logical network architecture of peer-assisted CDN for live streaming in Fig.1. Fig. 3 shows a simplified physical representation of the network. While in a logical network the peer can communicate directly with another peer, in a real physical world the communication between peers always passes through a router inside the data center.

Fig. 4 shows energy usage for CDN server, router, and the total energy consumption for the CDN scenario (without peer assist). We plot the energy consumption for CDN server, router, clients, and total energy for two COP coefficient values (T). All energy consumption components increase in value as the number of peers increases; and peers consume most of the energy. The effect of variations in T on total energy is small.

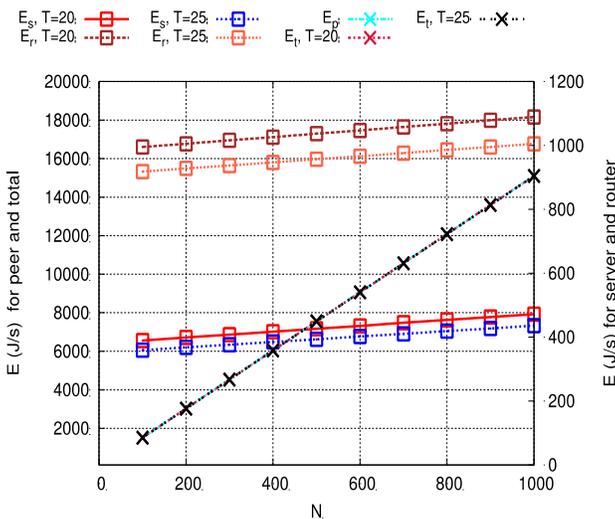


Figure 4. Power consumption for the server, router, peer, and total system for CDN architecture. Note that the server and router energy are plotted using the right hand scale, and the peer and total energy are plotted using the left hand scale.

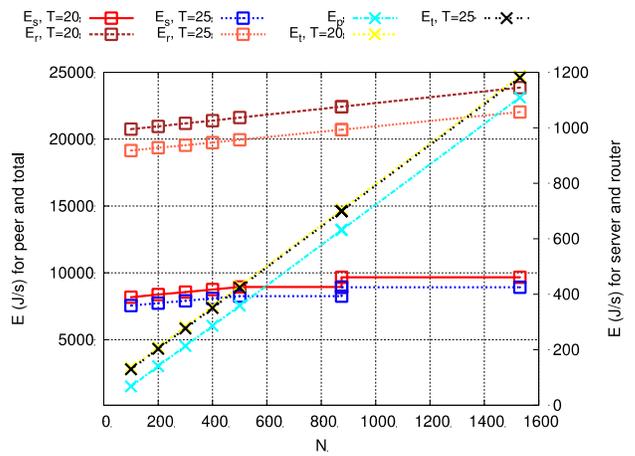


Figure 5. Power consumption for the server, router, peer and total system for peer-assisted CDN with $N_s^u = 0.75$. Note that the server and router energy are plotted using the right hand scale, and the peer and total energy are plotted using the left hand scale.

Fig. 5 shows the energy consumption of all components for the CDN-P2P scenario. We use the peer upload rate $N_s^u = 0.75$ in Fig. 5. We observe that there is almost no change in peers' power consumption compared to the pure CDN. The router consumes more power with a higher rate of increase because in CDN-P2P peers the traffic originated from the seeders passes through the router twice. The CDN server power consumption has small increases between $N=100$ and $N=500$, while there are sections with no power increase at $N > 500$, which is where the seeders are uploading contents to the leechers. The server power consumption remains flat as long as the upload rate does not exceed the defined peer upload rate.

Fig. 6 shows the energy savings of CDN-P2P compared to CDN architecture for CDN server with the utilization policy as explained in Sec. II. A, for $N_s^u = [0.25, 0.5, 0.75, 1]$.

Let's take $N_s^u = 0.75$ as an example. The first 500 nodes are served directly by the CDN server since the utilization of the CDN server is 50% or less. We consider 500 nodes to be 50% utilization because a video-rate of 1 Mbps will result in total network traffic of 0.5 Gbps, which is half of a Gigabit Ethernet interface. When more peers join the system, these peers will be treated as leechers as long as the upload ratio condition is fulfilled. The number of leechers that can be supported by seeders is 375. Therefore from $N=500$ to $N=875$, the CDN server does not need to increase utilization because the leechers can be supported by seeders, thus we see that the CDN server saves energy. In this phase, compared to CDN architecture, CDN-P2P energy saving is around 7.8%. Next, we have 875 total peers in the system, is apportioned into 500 peers as seeders and 375 as leechers. Since more peers joining the system, the CDN increases the utilization from 50% to 87.5% so all current 875 peers become seeders. In this phase, 875 seeders can support an additional 656 leechers. Therefore, from $N=875$ to $N=1531$ the CDN utilization is flat at 87.5% because 875 seeders can support 656 leechers thus we have energy savings around 11% compared to CDN architecture. Other values of N_s^u have same pattern as shown in Fig. 6.

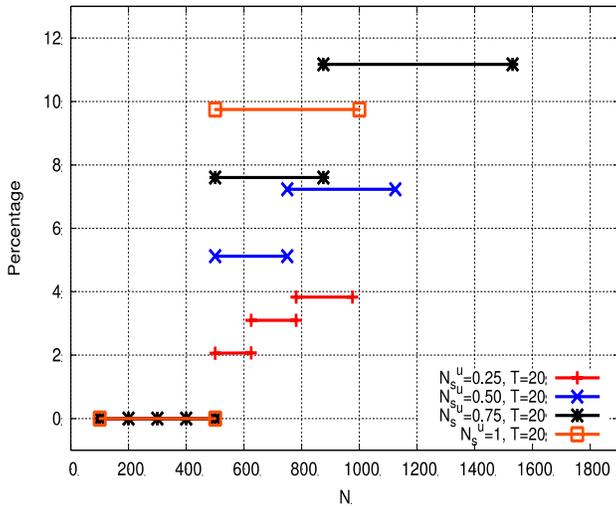


Figure 6. Savings in power consumption between CDN architecture and peer-assisted CDN for server with $N_s^u = [0.25, 0.5, 0.75, 1]$.

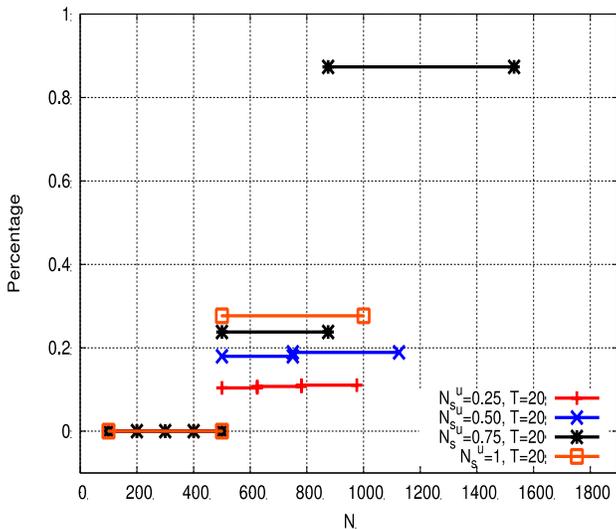


Figure 7. Savings in power consumption between CDN architecture and peer-assisted CDN for total system with $N_s^u = [0.25, 0.5, 0.75, 1]$.

Fig. 7 shows the total energy savings of CDN-P2P compared to CDN architecture for $N_s^u = [0.25, 0.5, 0.75, 1]$. As the savings only occurs in the CDN server, we see the same patterns as in Fig. 6. but with a much lower percentage of energy savings, which is 1%.

C. Online Storage Service:

To calculate the energy consumption in peer-assisted online storage, we must be able to determine the number of peers in the system. We get average downloading time (T_d) values by varying server bandwidth values (S) from 0 to 150 MBps using Eq. (4). After getting the average downloading time, we can get the average number of peers using Eq. (3). We found that the number of peers is inversely related to the server bandwidth. The number of peers is the horizontal axis in Fig. 8 and Fig. 9. The figures cover more number of peers compared to the live streaming service and we can look at the comparable number of peers in both cases when we want to do comparisons.

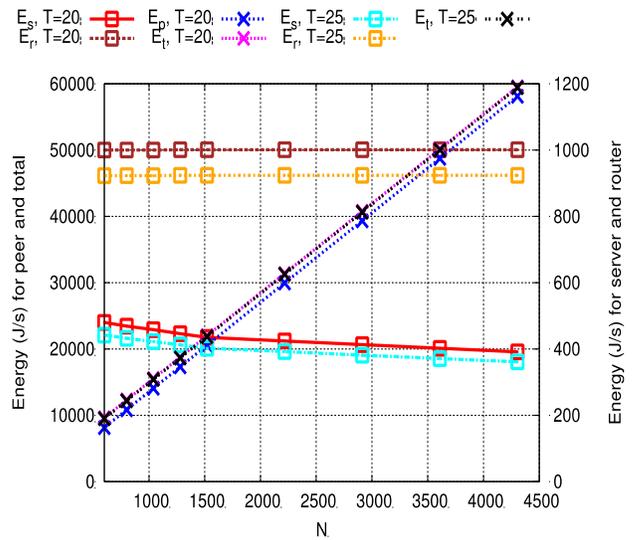


Figure 8. Power consumption for lower bound strategy. Note that the server and router energy are plotted using the right hand scale, and the peer and total energy are plotted using the left hand scale.

Fig. 8 shows the power consumption for the lower-bound strategy for the server, router, peers, and the total system. We found that increasing the number of peers decreases CDN server power consumption because the bandwidth usage of the CDN server decreases. The router power consumption is flat at around 1000J/s because the server bandwidth reduction is offset by the increasing number of peers. We also found that the other strategies request driven and water leveling, have the same pattern as the power consumption of the lower-bound strategy.

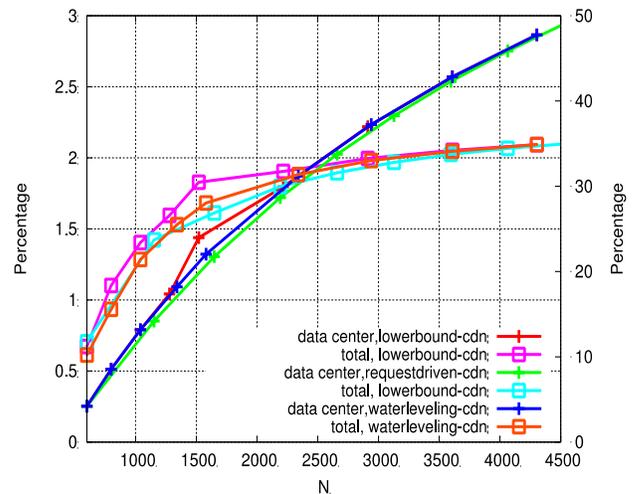


Figure 9. Savings in power consumption between each bandwidth allocation and CDN architecture for total and server. Note that the data center energy is plotted using the right hand scale, and the peer and total energy is plotted using the left hand scale.

Fig. 9 shows a comparison of the energy consumption between the request driven and the lower bound strategy, and between the water-leveling strategy and the lower bound strategy for different numbers of peers. Compared to the water-leveling strategy, the request driven strategy required more energy because the request driven strategy equalizes the server bandwidth across all the peers. The water-leveling strategy

equalizes server bandwidth across all the files by taking file popularity into consideration, thus minimizing downloading time. We mentioned before that the number of peers is inversely related to the server bandwidth, therefore for the same server bandwidth, we get different numbers of peers for each strategy. This implies that for the same number of peers, we get different server bandwidth. That is the reason for $1000 < N < 2500$ the power consumption diverges. In very limited server bandwidth (less than 45MBps) and sufficient server bandwidth (more than 120MBps) each strategy has the same downloading performance [12]. That is the reason for $N < 1000$ and $N > 2500$ we have the same the number of peers for same bandwidth. As shown in Fig. 9, for $N < 1000$ and $N > 2500$ the savings for each strategy is relatively the same.

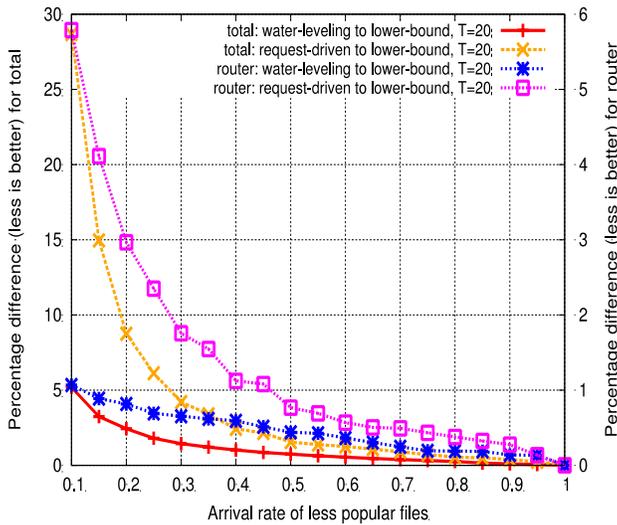


Figure 10. Power consumption for total (left axis) and router (right axis) under different server bandwidth allocation strategies when peer arrival rate of less popular varies. We use fixed server bandwidth $S = 50$ MBps.

File popularity has a strong correlation with downloading performance. We examine popularity by varying the peer arrival rate of less popular files while the server has fixed bandwidth. Specifically, we increase the type-1 files popularity from 0.1 to 1 and we choose a fixed server bandwidth $S=50$ MBps which is similar to FS2You. Fig. 10 shows the difference in total power consumption (left axis) and router power consumption (right axis) of that case. Since the server bandwidth is fixed, we only show the power consumption changes in the routers and the total. When the arrival rate is less than 0.5, the request driven strategy has worse downloading performance compared to the water leveling strategy. This implies more peers exist in the request driven strategy than the water leveling strategy. Therefore, the energy consumption of the request driven strategy is higher than the water leveling strategy. Generally for each strategy, increasing the peer arrival rates to less popular file makes the total energy consumption and router energy consumption increase because more peers are present in the system. Increasing the peer arrival rates to the less popular file makes both the request driven and the water leveling strategy energy consumption converge to a lower bound.

This is because more peers in the system improve P2P content availability, thus improving downloading performance that converges to the lower bound strategy.

IV. RELATED WORK

Content Distribution Networks with peer assist have been successfully deployed on the Internet, such as Akamai [15] and LiveSky [9.] The authors of [15] conclude from two real world traces that hybrid CDN-P2P can significantly reduce the cost of content distribution and can scale to cope with the exponential growth of Internet video content.

Yin et al. [9] described LiveSky as commercial operation of a peer-assisted CDN in China. LiveSky solved several challenges in the system design, such as dynamic resource scaling of P2P, low startup latency, ease of P2P integration with the existing CDN infrastructure, and network friendliness and upload fairness in the P2P operation. Measurement from LiveSky showed that LiveSky could save around 40% of the bandwidth [9]. The author in [16] and [17] proposed mechanisms to minimize CDN server bandwidth to make the content distribution cheap. They designed different peer prefetching policies for a video on demand system in surplus mode while ensuring user quality of experience. A similar analysis has been done in [18] for a live video streaming system where the authors proposed different limited peer contribution policies to reduce CDN bandwidth requirements. An ISP friendly rate allocation scheme for a hybrid CDN-P2P video on demand system is discussed in [19]. This technique tries to minimize CDN server bandwidth while reducing ISP unfriendly traffic and maximizing peer prefetching. Load on the CDN server has been shown to be reduced using this approach while reducing cross ISP traffic. The above studies were performed for video on demand or live video streaming. While video is the most popular content, the systems can be used for other type contents. Moreover while content-based services are growing, energy consumption of a content distribution system has not been analyzed.

First, related to CDN and energy usage, in a seminal work [20], Qureshi et al. show that if costs are based on electricity usage and if the power prices vary in real-time, global load balancing decisions can be made such that users are routed to locations with the cheapest power without significantly impacting user performance or bandwidth cost. In [21], the authors proposed utilizing batteries for CDN for reducing total supplied power and total power costs. The authors [21] also proposed battery-provisioning algorithms based on workload of CDN server. The result shows that batteries can provide up to 14% power savings.

The idea of utilizing an ISP controlled home gateway to provide computing and storage services and adopts managed peer-to-peer model is presented in [8]. Valancius et al. [8] show that such a distributing computing platform in NaDa (Nano Data Center) save at least 20-30% energy compare to traditional data centers.

The comparison between CDN architecture and peer-to-peer architecture is discussed in [5] and [6]. Both authors in [5] and [6] agree that CDN architecture saves more energy compared to a peer-to-peer architecture. Another study of energy efficiency in content delivery architectures is presented by Guan et al.

[22].They compare the energy efficiency of CDN architecture and CCN architecture. The authors conclude that CCN is more energy efficient in delivering popular content while the approach with optical bypass is more energy efficient in delivering infrequent accessed content.

To the best of our knowledge, the first study of energy consumption in peer-assisted CDN architecture is presented in [23]. Mandal et al. [23] mentioned that hybrid CDN-P2P systems could reduce energy in the optical core network around 20-40%. The authors only considered the energy consumption of hardware, especially the optical devices and do not include overhead inside the data center, CDN server energy consumption, and power consumed by peers. Our work is quite different, in that we take the number of peers and add the cooling power overhead of the data center for different applications purpose: live streaming and online storage.

V. CONCLUSION

In this paper, we compared the energy consumption between peer-assisted CDN and pure CDN for live streaming and online storage services both at the data center as well as in total. Employing peer-to-peer capability to assist a CDN is thought to lower the energy requirements at data centers, and we found that the maximum savings at the data centers are 11% and 21%, respectively for the live streaming and online storage services. These savings may change depending on the COP values used and should be better if a new generation of power proportional server were used [24].

One thing to note is that as the number of peers increases, the server's energy consumption increases for the live streaming and decreases for the online storage service due to the differences in the ways both services handle peers.

However, the server's energy consumption is swamped by the peers' energy consumption. Despite this difference in behavior in the two cases, when comparing Peer Assisted CDN to pure CDN, we found the total energy savings of less than 1%. Nevertheless, the total energy consumption is large, so that even a small percentage improvement results in valuable net reduction.

Several areas that we identified as the future work are: 1. The effect of peer's uptime variations; 2. More realistic file popularity models for the online storage service; and 3. How CDN providers or ISPs give incentives to the peers based on the understanding of the energy consumptions.

VI. REFERENCES

- [1] Uptime Institute, "the 2012 uptime institute data center industry survey," <http://uptimeinstitute.com/2012-survey-results>.
- [2] Google, "Google data center power usage," <http://www.google.com/corporate/datacenters/measuring.html>.
- [3] L. Barroso and U. Holzle, "The case for energy-proportional computing," *Computer*, vol. 40, no. 12, pp. 33–37, 2007.
- [4] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsang, and S. Wright, "Power awareness in network design and routing," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, 2008, pp. 457–465.
- [5] A. Feldmann, A. Gladisch, M. Kind, C. Lange, G. Smaragdakis, and F.- J. Westphal, "Energy trade-offs among content delivery architectures," in *Telecommunications Internet and Media Techno Economics (CTTE), 2010 9th Conference on. IEEE*, 2010, pp. 1–6.
- [6] J. Baliga, K. Hinton, and R. S. Tucker, "Energy consumption of the internet," in *Optical Internet, 2007 and the 2007 32nd Australian Conference on Optical Fibre Technology. COIN-ACOFT 2007. Joint International Conference on. IEEE*, 2007, pp. 1–3.
- [7] M. Dischinger, A. Haeberlen, K. P. Gummadi, and S. Saroiu, "Characterizing residential broadband networks," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC '07. New York, NY, USA: ACM, 2007, pp. 43–56. [Online]. Available: <http://doi.acm.org/10.1145/1298306.1298313>
- [8] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez, "Greening the internet with nano data centers," in *Proceedings of the 5th international conference on Emerging networking experiments and technologies*. ACM, 2009, pp. 37–48.
- [9] H. Yin, X. Liu, T. Zhan, V. Sekar, F. Qiu, C. Lin, H. Zhang, and B. Li, "Livesky: Enhancing cdn with p2p," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 6, no. 3, pp. 16:1–16:19, Aug. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1823746.1823750>
- [10] Y. Sun, F. Liu, B. Li, B. Li, and X. Zhang, "Fs2you: Peer-assisted semi-persistent online storage at a large scale," in *INFOCOM 2009, IEEE*, 2009, pp. 873–881.
- [11] S. Das, S. Tewari, and L. Kleinrock, "The case for servers in a peer-to-peer world," in *Communications, 2006. ICC '06. IEEE International Conference on*, vol. 1, 2006, pp. 331–336.
- [12] Y. Sun, F. Liu, B. Li, and B. Li, "Peer-assisted online storage and distribution: modeling and server strategies," in *Proceedings of the 18th international workshop on Network and operating systems support for digital audio and video*, ser. NOSSDAV '09. New York, NY, USA: ACM, 2009, pp. 13–18. [Online]. Available: <http://doi.acm.org/10.1145/1542245.1542249>
- [13] S. Nedeveschi, S. Ratnasamy, and J. Padhye, "Hot data centers vs. cool peers," in *Proceedings of the 2008 conference on Power aware computing and systems*, ser. HotPower'08. Berkeley, CA, USA: USENIX Association, 2008, pp. 8–8. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1855610.1855618>
- [14] J. Moore, J. Chase, P. Ranganathan, and R. Sharma, "Making scheduling "cool": temperature-aware workload placement in data centers," in *Proceedings of the annual conference on USENIX Annual Technical Conference*, 2005, pp. 5–5.

- [15] C. Huang, A. Wang, J. Li, and K. W. Ross, “Understanding hybrid cdn-p2p: why limelight needs its own red swoosh,” in *Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, ser. NOSSDAV '08. New York, NY, USA: ACM, 2008, pp. 75–80. [Online]. Available: <http://doi.acm.org/10.1145/1496046.1496064>
- [16] C. Huang, J. Li, and K. W. Ross, “Can internet video-on-demand be profitable?” *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 133–144, Aug. 2007. [Online]. Available: <http://doi.acm.org/10.1145/1282427.1282396>
- [17] C. Huang, J. Li, and K. Ross, “Peer-assisted vod: Making internet video distribution cheap,” in *Proc. of IPTPS*, vol. 7, 2007.
- [18] D. Xu, S. S. Kulkarni, C. Rosenberg, and H.-K. Chai, “Analysis of a cdn– p2p hybrid architecture for cost-effective streaming media distribution,” *Multimedia Systems*, vol. 11, no. 4, pp. 383–399, 2006.
- [19] J. Wang, C. Huang, and J. Li, “On isp-friendly rate allocation for peer-assisted vod,” in *Proceedings of the 16th ACM international conference on Multimedia*, ser. MM '08. New York, NY, USA: ACM, 2008, pp. 279–288. [Online]. Available: <http://doi.acm.org/10.1145/1459359.1459397>
- [20] A. Qureshi, R. Weber, H. Balakrishnan, J. Gutttag, and B. Maggs, “Cutting the electric bill for internet-scale systems,” *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 4, pp. 123–134, 2009.
- [21] D. S. Palasamudram, R. K. Sitaraman, B. Urgaonkar, and R. Urgaonkar, “Using batteries to reduce the power costs of internet-scale distributed networks,” in *Proceedings of the Third ACM Symposium on Cloud Computing*, ser. SoCC '12. New York, NY, USA: ACM, 2012, pp. 11:1– 11:14. [Online]. Available: <http://doi.acm.org/10.1145/2391229.2391240>
- [22] K. Guan, G. Atkinson, D. Kilper, and E. Gulsen, “On the energy efficiency of content delivery architectures,” in *Communications Workshops (ICC), 2011 IEEE International Conference on*, 2011, pp. 1–6.
- [23] U. Mandal, C. Lange, A. Gladisch, and B. Mukherjee, “Should isps adopt hybrid cdn-p2p in ip-over-wdm networks: An energy-efficiency perspective?” in *Advanced Networks and Telecommunications Systems (ANTS), 2012 IEEE International Conference on*, 2012, pp. 6–11.
- [24] A. Krioukov, P. Mohan, S. Alspaugh, L. Keys, D. Culler, and R. Katz, “Napsac: design and implementation of a power-proportional web cluster,” *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 1, pp. 102–108, Jan. 2011. [Online]. Available: <http://doi.acm.org/10.1145/1925861.1925878>
- [25] Long Vu, Indranil Gupta, Klara Nahrstedt, and Jin Liang. 2010. Understanding overlay characteristics of a large-scale peer-to-peer IPTV system. *ACM Trans. Multimedia Comput. Commun. Appl.* 6, 4, Article 31 (November 2010), 24 pages. DOI=10.1145/1865106.1865115 <http://doi.acm.org/10.1145/1865106.1865115>
- [26] httpperf. <http://www.hpl.hp.com/research/linux/httpperf/>.
- [27] <http://arstechnica.com/tech-policy/2012/03/isp-storing-25-petabytes-of-megaupload-data-costs-us-9000-a-day/>
- [28] Aniket Mahanti. 2011. Measurement and analysis of cyberlocker services. In *Proceedings of the 20th international conference companion on World wide web (WWW '11)*. ACM, New York, NY, USA, 373-378. DOI=10.1145/1963192.1963346 <http://doi.acm.org/10.1145/1963192.1963346>
- [29] FS2You web. <http://www.rayfile.com/en/>