IEEE INFOCOM 2011 Workshop on Green Communications and Networking

# Evaluating the Energy Benefit of Dynamic Optical Bypass for Content Delivery

Kyle Guan, Daniel C. Kilper, and Gary Atkinson
Bell Labs, Alcatel-Lucent, 791 Holmdel Road, Holmdel, NJ 07733

***Abstract*—We evaluate the energy efficiency of dynamic optical bypass for decentralized content delivery networks (CDNs). We build energy models based on the energy consumption of current network equipment and devices and analyze the energy tradeoff among key networking resources. Our results show that, due to the under-utilization associated with signaling and reconfiguration overhead, a CDN with dynamic optical bypass achieves the largest savings in delivering very large files (100 Gb (gigabit) and above) with high download rate (a few hundred downloads per hour). We also derive a threshold for a file size, which is approximately the bandwidth-overhead product scaled by the ratio between the power density of WDM equipment and routers: $C_\lambda T^\circ p_d^{wdm}/p_d^{r}$[1]. We show that, only for the delivery of content of sizes larger than this size threshold, a CDN with dynamic optical bypass is more energy efficient than CDN without bypass.**

**Index Terms— Content distribution network (CDN), dynamic optical bypass, energy efficiency**

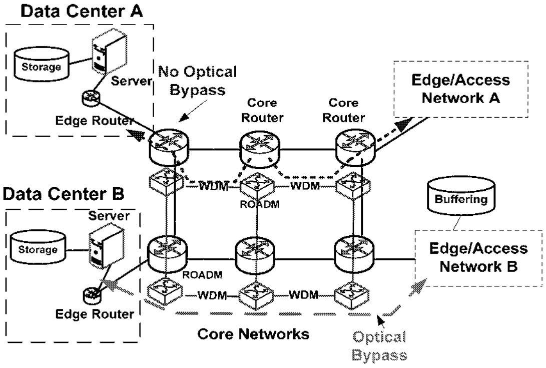

Fig. 1. Content delivery architectures: decentralized CDN with or without optical bypass.

## I. INTRODUCTION

Content service providers are rapidly expanding data centers to meet the growing demands for both user generated and multi-media content. Concerns over the energy requirements for these data centers and the associated network equipment have drawn attention recently [1]. In the coming decade, the improvement in networking energy efficiency is anticipated to lag the traffic growth [2]. The energy consumption is predicted to increase substantially unless the efficiency of the network, both in terms of equipment and the associated architectures, improves proportionately. As a continuation of our current research on energy efficient CDN [3], this paper focuses on evaluating and comparing the potential energy savings of different transport architectures for content distribution and dissemination.

In traditional content distribution networks (CDNs), content is delivered to end users through host servers that are centrally managed in a few data centers. With the growing demand of large multimedia content, the energy consumption becomes problematic due to factors such as over-provisioning (to satisfy the peak traffic) and heat dissipation requirements [1]. As all content is stored in a few fixed locations regardless of the relative download rate, delivering frequently accessed content from a few centralized locations increases bandwidth-mileage and often incurs unnecessary yet significant transport energy cost [4]. Decentralized CDNs [1] [4], through the deployment of micro or nano data centers throughout the network, push the content closer to the network edge and end users. As such, transit traffic and consequently the associated energy consumption are reduced.

Optical bypass is a physical layer approach in managing transport energy [3]-[7]. Static optical bypass is commonly used in reconfigurable optical add-drop multiplexer (ROADM) based networks to transparently patch signals through a node. In dynamic optical bypass, a transparent optical connection is provisioned on demand, thus circumventing intermediate core routers, as shown in Fig. 1. The dynamic capability is needed if a dedicated transparent path would be impractical. Since bypassing transit traffic is more energy efficient than processing them at the core routers, dynamic optical bypass can potentially yield significant energy savings, especially for delivering content over a large number of network hops. However, dynamic optical bypass faces technical challenges in setting up an end-to-end wavelength route rapidly over long-haul distances. As the overhead required by both signaling and network element reconfiguration increases, the utilization of the capacity in a wavelength is reduced. This, in turn, inadvertently increases the effective energy consumption per download. As such, these issues warrant a closer examination of optical bypass as an energy saving transport measure for content delivery.

In this paper, we evaluate the potential benefit of dynamic optical bypass used in decentralized CDN architectures. In comparison to previous work [3]-[7], our approach takes into consideration the following issues:

- We address the impact of technical challenges in setting

---

[1] $C_\lambda$, $T^\circ$, $p_d^{wdm}$, and $p_d^{r}$ denote the wavelength capacity, overhead, power density of WDM equipment and routers, respectively.

978-1-4244-9920-5/11/$26.00 ©2011 IEEE

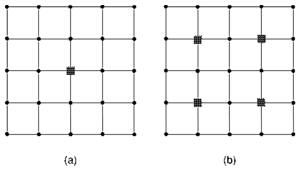

Fig. 2. The placement of copies on a 5 × 5 grid topology. (a) The optimal placement of 1 copy; (b) the optimal placement of 4 copies.

up dynamic optical bypass. Under-utilization due to signaling and reconfiguration overhead is included in the transaction based energy consumption models.

- We consider the delivery of content that spans several orders of magnitude in size and has a heavy-tailed distribution [8] (in size).
- We study the impact of dynamic optical bypass on the energy tradeoff among key networking resources. To this end, we provide an analytical framework to evaluate how networking, content, and equipment parameters impact the tradeoff between the content transport and storage energy.

We derive a threshold for file size, which is approximately the bandwidth-overhead product scaled by the ratio between the power density of WDM equipment and routers. As a rule of thumb, we show that a CDN with dynamic optical bypass is more energy efficient only for the delivery of content of sizes larger than this threshold. Our results also show that dynamic optical bypass achieves the largest per-bit energy savings in delivering large-sized content (e.g., 100 Gb above) with a high download rate (e.g., 100 downloads per hour). In today's Internet, the support of multimedia as well as massive archives of medical and scientific data has led to a significant and continuing increase in the content size as well as the tail weight of the size distribution [9]. With the maturity of the technology, a decentralized CDN using dynamic optical bypass is likely to gain more energy savings in the future.

The rest of the paper is organized as follows. In Section II, we provide the network and content size distribution models. In Section III, we set up energy models for different content delivery architectures. Optimization formulations and analyses that evaluate energy efficiency of different architectures are presented in Section IV. In Section V, we present the results of case studies and discuss their implications.

## II. NETWORK AND CONTENT SIZE DISTRIBUTION MODELS

In this section, we briefly describe the network and content size distribution models in preparation for the modeling and analyses in Sections III and IV.

### *A. Network Model*

We focus on core networks, which in this paper are assumed to consist of both long-haul and metropolitan networks. We follow the convention of representing the core network as a graph $G(N, L)$, with $N$ and $L$ denoting the number of nodes (representing core routers, ROADMs, or optical cross-connects OXCs) and the number of bidirectional edges (representing WDM links), respectively.

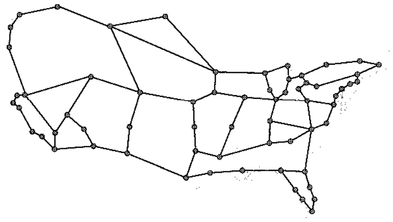

Fig. 3. Deployed networks: a hypothetical backbone network US64, with $N = 64$ and $\bar{\Delta} = 2.64$. The mean node degree of the network, $\bar{\Delta}$, is given by $\bar{\Delta} = 2L/N$.

In the context of CDN, a total of $n$ replicas[2] of the same content are cached in $n$ distributed data centers. That is, $N - n$ nodes in the core network have to traverse through one or more hops to reach a copy. In this context, the average hop distance to a replica, $H_r$, is an important measure of the placement efficiency. $H_r$ drives the tradeoff between the caching and the transport energy, as will be discussed in Sections III and IV. The exact form of $H_r$ depends on both the network topology as well as the replica placement algorithm. In this paper, we derive and estimate $H_r$ (as a function of $N$ and $n$) for both symmetric and deployed networks.

We first consider topologies with symmetric structures, in particular ring and grid topologies. For a ring topology of $N$ nodes, the optimal placement is achieved by maximizing the hop distance among the replicas. Based on this, an analytical expression of $H_r$ can be derived and approximated as $H_r(n) \approx \frac{1}{4}\left(\frac{N}{n}\right)$.

For an $M \times M$ grid topology, $H_r$ can be estimated as follows. When $n = 1$, the optimal placement of the copy is at the center of the grid, as shown in Fig. 2 (a). An expression for $H_r(1)$ can be derived and approximated as $H_r(1) \approx \frac{\sqrt{N}}{2}$. For $n > 1$, the $M \times M$ grid is divided into $n$ sub-grids with (almost) equal size. It can be shown[3] that the optimal placement is to put each of the $n$ copies at the center of each sub-grid, as shown in Fig. 2 (b). For $n > 1$ and $N \gg 1$, $H_r$ can thus be estimated as $H_r(n) \approx \frac{1}{2}\left(\frac{N}{n}\right)^{\frac{1}{2}}$.

As a representation of the topological characteristics of commercial networks, we use a hypothetical backbone network US64 as shown in Fig. 3. This network is modified from the prototypical Continental United States (CONUS) topology used in the DARPA CORONET program [10]. For this irregular and asymmetric network, we take a semi-analytical approach in deriving $H_r$ as a power law function of $n$. We first numerically compute the average hop distance to a replica[4]. We next fit the computed data to the function

---

[2] We use the terms replica, copy, and cached content interchangeably herein.

[3] The proof is omitted in the paper due to space limitation.

[4] A genetic algorithm approach is used to find optimal or near-optimal placement of replicas. Due to space limitation, the details are omitted here, but are to be covered in our future work.

TABLE I
NOTATIONS AND VALUES OF THE KEY PARAMETERS

| Notations | Parameter | Values [4] |
|---|---|---|
| $N$ | No. of nodes | Multiple values |
| $t$ | Time duration | Multiple values |
| $R_k$ | No. of downloads of the files in the $k$th category | Multiple values |
| $B_k$ | Average size of the files in the $k$th category | Multiple values |
| $p_d^r$ | Power density of a core router | $1.2 \times 10^{-8}$ J/bit |
| $p_d^{wdm}$ | Power density of a WDM per hop (based on an average distance of 1000 km) | $1.48 \times 10^{-9}$ J/bit |
| $p_d^{roadm}$ | Power density of a ROADM | $1.95 \times 10^{-11}$ J/bit |
| $w_{st}$ | **Power** of storage equipment | $7.84 \times 10^{-11}$ **W/bit** |
| $n$ | No. of caching replicas | Multiple values |
| $H_r$ | Average number of hops to a replica | Multiple values |

$H_r(n) \approx A\left(\frac{N}{n}\right)^{\alpha}$, with $\alpha > 0$. For US64, $H_r$ can be estimated as $0.36(N/n)^{0.6}$.

### B. Content Size Distribution Model

In this paper, we consider content in the form of data files. These files vary by several orders of magnitude in size. In addition, the sizes of these files follow a heavy-tailed distribution. Specifically, we use a bounded Pareto distribution [8]. Let $B_L$ and $B_U$ denote the lower and upper bound of the content size, respectively. The probability density function of the file size $x$ is given by:

$$f(x) = \frac{\beta B_L^{\beta} x^{-(\beta+1)}}{1 - (\frac{B_L}{B_U})^{\beta}},\ B_L \leq x \leq B_U, \tag{1}$$

where $\beta > 0$ is the shape parameter. A large $\beta$ indicates a relatively small percentage of very large files. For our analysis, we divide the files into different categories according to their sizes. Let the size of category $k$ be between a lower bound $B_L^k$ and an upper bound $B_U^k$. If there are a total of $R$ downloads in a time duration $t$, the number of downloads in category $k$, $R_k$, is given by

$$R_k = R \int_{B_L^k}^{B_U^k} f(x)dx. \tag{2}$$

The average file size of category $k$, denoted as $B_k$, is

$$B_k = \frac{\int_{B_L^k}^{B_U^k} x f(x)dx}{\int_{B_L^k}^{B_U^k} f(x)dx}. \tag{3}$$

## III. ENERGY CONSUMPTION MODELS

In preparation for the analysis in Section IV, we set up energy consumption models for both decentralized CDNs with and without dynamic optical bypass. For clarity, we summarize the notation of the key parameters in Table I.

### A. Energy Model for a Decentralized CDN without Optical Bypass

The total energy for a decentralized CDN consists of three parts: the storage energy $E_{st}$, the server energy $E_{sr}$, and the transport energy $E_{tr}$. The transport energy, in turn, includes the energy consumption in the core, edge, and access networks. In this work, the main difference between the two architectures lies in the transport mechanism of core networks, i.e., whether dynamic optical bypass is used or not. In edge and access networks, the transport mechanism is the same for both architectures. Thus, the energy per bit is the same and only adds an identical constant to each case. As such, the energy consumed in edge and access networks is not included in the analysis. As expressed in equations (1)-(3), in a duration $t$ the files of category $k$ (with an average size $B_k$) are downloaded $R_k$ times. If on average it takes $H_r$ hops to access a file, the energy consumed to transport $R_k B_k$ bits in core network, $E_{tr}^k$, is

$$E_{tr}^k(H_r) = 4B_k R_k[(p_d^r + p_d^{roadm})(H_r + 1) + p_d^{wdm} H_r]. \tag{4}$$

Note that a factor of four is used to account for redundancy and overhead [4].

If $n$ copies of a particular content are cached in $n$ different storage sites, the energy consumed by storing $nB_k$ bits is

$$E_{st}^k(n) = B_k t w_{st} n = (B t w_{st} A^{1/\alpha} N)/H_r^{\frac{1}{\alpha}}, \tag{5}$$

where we assume the network topologies used all follow the power law scaling of $H_r(n) = A(N/n)^{\alpha}$.

Since the server energy consumption is independent of transport mechanism and the number of storage sites [4], it is the same for both architectures. Therefore, it is not included in the following analysis.

Combining the transport and storage energy, the total energy for the delivery of $R_k B_k$ bits, $E_{tot,k}^{CDN}$, is

$$E_{tot,k}^{CDN} = E_{tr}^k(H_r) + E_{st}^k(H_r). \tag{6}$$

### B. Energy Model for a Decentralized CDN with Dynamic Optical Bypass

With dynamic optical bypass, a file is delivered across the core network via transparent optical connections. We assume that the file is transmitted at the full capacity of a wavelength $C_\lambda$. Resources dedicated to the transmission (e.g., transceivers and wavelengths) are occupied and maintained during the lifetime of the file transmission. After the transmission is over, these resources are released. Implementing dynamic optical bypass requires an overhead for both signaling and physical reconfiguration of the network elements. This overhead increases the wavelength switching time and thus reduces the utilization of each wavelength. This, in turn, increases the effective energy consumption per download. We use an under-utilization factor to account for the penalty on transport energy. Let $T^{\circ}$ denote the overhead for signaling and reconfiguration, the under-utilization factor for downloading a file of size $B_k$, denoted as $\gamma_k$, is defined as

$$\gamma_k = \frac{T^{\circ} + B_k/C_\lambda}{B_k/C_\lambda}. \tag{7}$$

Clearly, for a fixed overhead $T^{\circ}$ and wavelength capacity $C_\lambda$, delivering a file with a large size $B_k$ has a better utilization of the capacity, thus a smaller under-utilization factor $\gamma_k$. We make a simplification by assuming that during the signaling and reconfiguration, the network equipment consume approximately the same amount of energy that it would consume for sending traffic. By taking into consideration the under-utilization factor $\gamma_k$, we express the transport energy as

$$E^k_{tr}(H_r) = 4\gamma_k B_k R_k [p_d^{roadm}(H_r + 1) + p_d^{wdm} H_r]. \tag{8}$$

For dynamic optical bypass, extra buffering is required at the edge/access network. The energy for buffering a file with a size of $B_k$, $E^k_{bf}$, is

$$E^k_{bf} = t_{bf} B_k w_{st}, \tag{9}$$

where $t_{bf}$ denotes the time duration of buffering.

The energy consumption for content storage used in a network using optical bypass has the same form as that in a network without optical bypass (c.f. equation (5) ). The total energy for a decentralized CDN with dynamic optical bypass, $E^{BP}_{tot,k}$, is:

$$E^{BP}_{tot,k} = E^k_{tr}(H_r) + E^k_{bf} + E^k_{st}(H_r). \tag{10}$$

## IV. PROBLEM FORMULATIONS AND SOLUTIONS

In this section, we first find the minimal per-bit energy consumption of content delivery via the optimization of content placement for decentralized CDNs with and without dynamic optical bypass. Next we analyze and compare the relative energy benefit of these two architectures.

### *A. Energy Optimization for Decentralized CDNs without Optical Bypass*

For a decentralized CDN, an examination of equations (4) and (5) shows that $H_r$ generally drives the transport energy in an increasing direction and the storage energy in a decreasing direction. To find the optimal trade-off between transport and storage energy, we formulate the following optimization problem:

$$\begin{aligned} \underset{H_r}{\text{minimize}}: \quad & E^{CDN}_{tot,k}(H_r) \\ \text{subject to} \quad & 0 \le H_r \le H_{max}. \end{aligned} \tag{11}$$

Here $H_{max}$ denotes the maximal hop distance. We observe that $E^{CDN}_{tot}$ is a convex function of $H_r$. As such, optimal solutions exist. The optimal value $H_r^*$ can be solved as

$$H_r^* = \min[H_r^\circ, H_{max}], \tag{12}$$

where $H_r^\circ$ has the form of

$$H_r^\circ = \left[\frac{A^{\frac{1}{\alpha}} N t w_{st}}{4R_k(p_d^r + p_d^{roadm} + p_d^{wdm})}\right]^{\frac{\alpha}{\alpha+1}}. \tag{13}$$

The optimal number of copies $n^*$ can be solved as

$$n^* = \max[1, \min[n^\circ, N]], \tag{14}$$

where $n^\circ$ is

$$n^\circ = \left[\frac{4AR_k(p_d^r + p_d^{roadm} + p_d^{wdm})}{t w_{st}}\right]^{\frac{1}{\alpha+1}} N^{\frac{\alpha}{\alpha+1}}. \tag{15}$$

If the file size follows a bounded Pareto distribution, equations (13) and (15) show that small sized files (with a large $R_k$) are replicated more often (larger $n^*$) and are placed closer to the end user (smaller $H_r^*$).

The functional form of the optimal total energy depends on the value of $n^*$. We consider the following cases:

*1)* $1 < n^* < N$*:* In this case, the optimal total energy consumption is found to be

$$\begin{aligned} E^{CDN*}_{tot,k} &= 4B_k R_k(p_d^r + p_d^{roadm}) \\ &+ c_0 B_k R_k^{\frac{1}{1+\alpha}} (t w_{st} N)^{\frac{\alpha}{\alpha+1}} [(p_d^r + p_d^{roadm} + p_d^{wdm})]^{\frac{1}{\alpha+1}} \end{aligned} \tag{16}$$

where $c_0 = (2^{\alpha+3} A)^{\frac{1}{\alpha+1}}$. The energy per bit $E^{CDN}_{b,k}$ can be obtained via dividing $E^{CDN}_{tot,k}$ by the total number of downloaded bits $R_k B_k$. Thus, the per-bit energy consumption for downloading files in category $k$ is given by

$$\begin{aligned} E^{CDN*}_{b,k} &= \frac{E^{CDN*}_{tot}}{R_k B_K} = 4(p_d^r + p_d^{roadm}) \\ &+ c_0 R_k^{-\frac{\alpha}{1+\alpha}} (t w_{st})^{\frac{\alpha}{\alpha+1}} [(p_d^r + p_d^{roadm} + p_d^{wdm})N]^{\frac{1}{\alpha+1}} \end{aligned} \tag{17}$$

*2)* $n^* = 1$ *or* $n^* = N$*:* For the case that $n^* = 1$, the optimal per bit energy is:

$$E^{CDN*}_{b,k} = \frac{w_{st} t}{R_k} + 4[(p_d^r + p_d^{roadm})(AN^\alpha + 1) + p_d^{wdm} AN^\alpha]. \tag{18}$$

For the case that $n^* = N$, the optimal per bit energy is:

$$E^{CDN*}_{b,k} = \frac{N w_{st} t}{R_k} + 4(p_d^r + p_d^{roadm}). \tag{19}$$

### *B. Energy Optimization of Decentralized CDNs with Dynamic Optical Bypass*

For the decentralized CDNs with dynamic optical bypass, the tradeoff between transport and storage energy can be analyzed using an optimization formulation similar to equation (11). The optimal value $H_r^*$ is obtained as

$$H_r^* = \min[H_r^\circ, H_{max}], \tag{20}$$

with

$$H_r^\circ = \left[\frac{A^{\frac{1}{\alpha}} N t w_{st}}{4\gamma_k R_k(p_d^{roadm} + p_d^{wdm})}\right]^{\frac{\alpha}{\alpha+1}}. \tag{21}$$

The optimal number of copies $n^*$ can be solved as

$$n^* = \max[1, \min[n^\circ, N]], \tag{22}$$

where $n^\circ$ is

$$n^\circ = \left[\frac{4AR_k\gamma_k(p_d^{roadm} + p_d^{wdm})}{t w_{st}}\right]^{\frac{1}{\alpha+1}} N^{\frac{\alpha}{\alpha+1}}. \tag{23}$$

Similarly, the optimal per-bit energy is found to be

$$\begin{aligned} E^{BP*}_{b,k} &= \frac{E^{BP*}_{tot}}{R_k B_K} = 4p_d^{roadm} + \frac{t_{bf} w_{st}}{R_k} \\ &+ c_0 \gamma_k^{\frac{1}{1+\alpha}} R_k^{-\frac{\alpha}{1+\alpha}} (t w_{st})^{\frac{\alpha}{\alpha+1}} [(p_d^{roadm} + p_d^{wdm})N]^{\frac{1}{\alpha+1}} \end{aligned} \tag{24}$$

when $1 < n^* < N$. For the case that $n^* = 1$, the optimal per-bit energy is

$$E^{BP*}_{b,k} = \frac{(t_{bf} + t) w_{st}}{R_k} + 4\gamma_k[p_d^{roadm}(AN^\alpha + 1) + p_d^{wdm} AN^\alpha]. \tag{25}$$

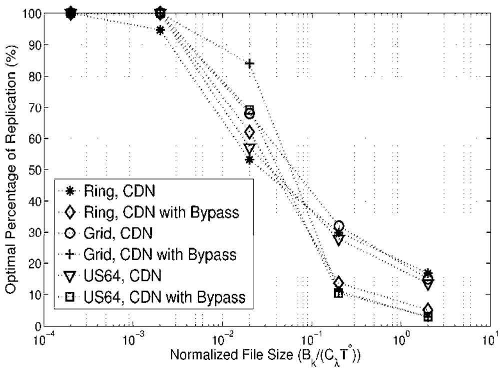

Fig. 4. The normalized optimal number of replications ($n^*/N$) for files of different size category.

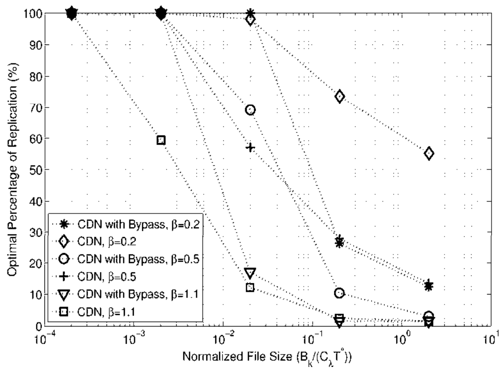

Fig. 5. The normalized optimal number of replications ($n^*/N$) for files of different size category. The US64 network is used.

For the case that $n^* = N$, the optimal per-bit energy is

$$E_{b,k}^{BP*} = \frac{t_{bf} w_{st}}{R_k} + \frac{N t w_{st}}{R_k} + 4\gamma_k p_d^{roadm}. \tag{26}$$

### *C. Energy Benefit of Decentralized CDNs with and without Dynamic Optical Bypass*

Equations (17) and (24) show that optical bypass could potentially save transport energy per bit by an order of magnitude $(p_d^{roadm} + p_d^{wdm}/p_d^r + p_d^{roadm} + p_d^{wdm}) \approx 10^{-1}$. However, under-utilization caused by signaling and reconfiguration overhead increases the effective transport energy. Moreover, the buffering at edge/access networks incurs an additional energy cost per download. As such, except for the trivial case where $E_b^{BP}$ is strictly less than $E_b^{CDN}$ under the conditions $t_{bf} \ll t$ and $\gamma_k \to 1$, it is difficult to carry out an analytical comparison of the averaged per-bit energy. Nonetheless, we can compare the optimal combined per-bit energy of transport and storage of both architectures (denoted as $E_{tr,st}^{CDN*}$ and $E_{tr,st}^{BP\,*}$, respectively) as follows:

$$\frac{E_{tr,st}^{BP\,*}}{E_{tr,st}^{CDN*}} = \gamma_k^{\frac{1}{1+\alpha}} \left[ \frac{(p_d^{roadm} + p_d^{wdm})}{\eta (p_d^r + p_d^{wdm} + p_d^{roadm})} \right]^{\frac{1}{\alpha+1}}, \tag{27}$$

where $\eta$ is a scaling factor. In order to have $E_{tr,st}^{BP\,*} < E_{tr,st}^{CDN*}$, the under-utilization factor $\gamma_k$ needs to satisfy the following condition:

$$\gamma_k < \frac{\eta (p_d^r + p_d^{wdm} + p_d^{roadm})}{(p_d^{wdm} + p_d^{roadm})}. \tag{28}$$

It in turn requires that

$$B_k > \frac{1}{\eta} C_\lambda T^\circ \frac{p_d^{wdm}}{p_d^r}. \tag{29}$$

That is, dynamic optical bypass saves combined storage and transport energy only when delivering files larger than a certain size. This size threshold is proportional to the bandwidth-overhead product scaled by the ratio between the per-bit energy of WDM equipment and routers.

## V. RESULTS AND DISCUSSIONS

In this section, we first evaluate the tradeoff between transport and storage energy. We next compare the energy benefit of decentralized CDNs with and without dynamic optical bypass. For the purpose of comparison with previous work, we use the same values of the key parameters as those in [4], as shown in Table I. We note that the values of $p_d^{wdm}$ and $p_d^r$ indicate a power density ratio $p_d^{wdm}/p_d^r \approx 10^{-1}$.

To understand the tradeoff between the transport and the storage energy, we consider the case of delivering files with sizes ranging from 10 megabits ($B_L = 10$ Mb) to 1 terabit ($B_U = 1$ Tb). We assume that there are a total of $10^5$ downloads during 1 day ($t = 8.64 \times 10^4$ seconds). The file size follows a bounded Pareto distribution (c.f. equation (1)) with $\beta = 0.5$. For our case studies, we divide these files into 5 categories, with the file size of the $k$th ($1 \le k \le 5$) category being in the range of $10^k$ Mb to $10^{k+1}$ Mb. In Fig. 5 we plot the normalized optimal number of replications ($n^*/N$) for different size category. We consider decentralized CDNs with and without dynamic optical bypass for three differently topologies: a 64-node ring, a $8 \times 8$ grid, and US64. For dynamic optical bypass, the wavelength capacity and signaling overhead are $C_\lambda = 10$ Gb/s and $T^\circ = 5$ seconds, respectively. In the plot, we normalize the file size by the capacity-overhead product ($C_\lambda T^\circ = 50$ Gb). The plot shows that, since the downloads of small files (10 Mb to 100 Mb) account for a significant portion of the total downloads, it is optimal to replicate these files at every node of the core network ($n^* = N$). As such, the transport energy is minimal and the storage energy is amortized by the large number of downloads. In comparison, the downloads of large files (100 Gb to 1 Tb) account for a very small portion of the total downloads. Therefore, they are replicated only at a small portion (less than 16%) of the nodes in the core network. We also observe the influence of the topology on the trade-off between transport and storage energy. For instance, a ring topology requires more replications in core networks to deliver large files (100 Gb to 1 Tb), compared to a grid topology and US64 network. This indicates that a ring topology is less efficient for content

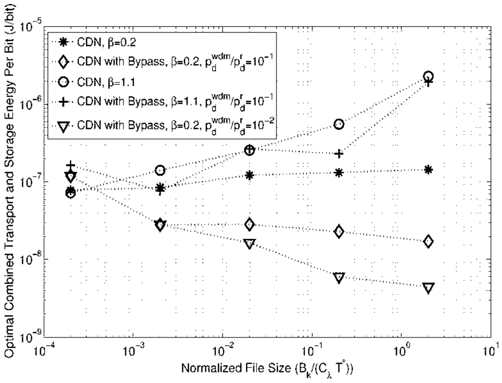

Fig. 6. Comparison of the optimal combined transport and storage energy per bit (*J*/bit) for delivering files of different sizes, between a decentralized CDN with and without dynamic optical bypass.

delivery, i.e., it requires more replications to reach a given hop distance $H_r$.

To evaluate the shape (tail weight) of Pareto distribution on the tradeoff between transport and storage energy, in Fig. 5 we plot the normalized optimal number of replications ($n^*/N$) for different values of the shape parameter $\beta$. The topology is the US64 network. The plot indicates that as the value of $\beta$ increases (e.g., $\beta = 1.1$), the replications of the large files (100 Gb to 1 Tb) decrease drastically to a few or none ($n^* = 1$). This can be explained by the fact that, at a low download rate the energy used for storing large files becomes dominant. As such, it is optimal to keep only a single copy of the file in the network.

We next compare the energy efficiency of decentralized CDNs with and without dynamic optical bypass. In Fig. 6, we plot the optimal combined storage and transport energy per bit for delivering files of different sizes over the US64 network. The plot shows that whether dynamic optical bypass can achieve per-bit energy savings depends on both the file size and the download rate. Significant energy savings are achieved in delivering large files (100 Gb above) with high download rate (a few hundreds per hour, e.g., $\beta = 0.2$). For smaller files (10-100 Mb), the potential energy savings of dynamic optical bypass are offset by the under-utilization. In this case, it is actually more energy efficient not to use dynamic optical bypass. For large files (100 Gb above) and very low download rates (a few downloads per **day**, e.g., $\beta = 1.1$), the file storage becomes under-utilized. As a result, the per bit storage energy dominates the total energy consumption. The savings of per bit transport energy through optical bypass are counter-balanced by the significant increase of the per bit storage energy. We also examine the effect of the energy efficiency improvement of WDM equipment relative to that of routers. In Fig. 6 we also plot the optimal combined energy per bit for the case where $p_d^{wdm}/p_d^r \approx 10^{-2}$. For the purpose of better illustration, we omit plotting the per bit energy for decentralized CDNs (without optical bypass) with $p_d^{wdm}/p_d^r \approx 10^{-2}$, as the corresponding values are very close to those for decentralized CDNs with $p_d^{wdm}/p_d^r \approx 10^{-1}$. The plot shows that in this case dynamic optical bypass achieves savings less than two orders of magnitude when delivering large files with high download rates. This indicates that some of the savings in per-bit transport energy are offset by the per-bit storage energy.

## VI. CONCLUSIONS

In this paper, we evaluate the energy efficiency of dynamic optical bypass for decentralized CDNs. Our results show that, due to the under-utilization associated with signaling and reconfiguration overhead, CDNs with dynamic optical bypass achieve the largest savings in delivering very large files with high download rate. Current Internet traffic sees a significant and continuing increase in the content size as well as the tail weight of the size distribution. As the technology for dynamic optical bypass matures and the energy efficiency of WDM equipment improves, decentralized CDNs using dynamic optical bypass are likely to gain more savings in combined transport and storage energy.

We note that we have employed simplifying assumptions in our modeling and analyses. Some practical issues are not explicitly addressed. For example, additional signal conditioning equipment, if needed for optical bypass, will increase the energy consumption and thus limit the benefit of optical bypass. In addition, we note that the overall network energy efficiency is still limited by edge/acces equipment and servers. We plan to address these issues in our future research.

## REFERENCES

[1] V. Valancius, N. Laoutaris, L. Massoulie, C. Diot, and P. Rodriguez, "Greening the Internet with Nano Data Centers," *CoNext 2009*, Rome, Italy, Dec. 2009.

[2] D. Kilper, et al., "Power Trends in Communication Networks," accepted for publication in *IEEE Journal on Selected Topics in Quantum Electronics 2011 Green Communication Issue*.

[3] K. Guan, G. Atkinson, D. Kilper, and E. Gulsen, "On the Energy Efficiency of Content Delivery Architectures," *The 4th IEEE International Conference on Communications (ICC) Workshop on Green Communications*, Kyoto, Japan, June 2011. (accpted)

[4] J. Baliga, R. Ayre, K. Hinton, and R. Tucker, "Architectures for Energy-Efficient IPTV Networks," *Optical Fiber Communication Conference (OFC)*, San Diego, CA, Mar. 2009.

[5] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy Consumption in Optical IP Networks," *Journal of Lightwave Technology*, vol. 27, pp. 2391-2403, 2009.

[6] A. Bianco, E. Bonetto, D. Cuda, G. G. Castillo, M.Mellia, and F. Neri, "Energy-Efficient Design of Wavelength-Routing Networks," *European Conference on Optical Communication (ECOC) 2010*, Torino, Italy, Sept. 2010.

[7] W. Vereecken, W. V. Heddeghem, B. Puype, D. Colle, M. Picakvet, and P. Demeester, "Optical Networks: How Much Power Do They Consume and How Can We Optimize This?" *European Conference on Optical Communication (ECOC) 2010*, Torino, Italy, Sept. 2010.

[8] M. Crovella, "Performance Evaluation with Heavy Tailed Distributions," *Revised Papers from 7th International Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP '01)*, Springer-Verlag, London, UK, 1-10.

[9] H. Kim, et al., "Performance Impact of Large File Transfer on Web Proxy Caching: A Case Study in a High Bandwidth Campus Network Environment," *Journal of Communications and Networks*, vol. 12, no. 1, pp. 52-66, Feb. 2010.

[10] Dynamic Multi-Terabit Core Optical Networks: Architectures, Protocols, Control and Management (CORONET), http://www.darpa.mil/sto/programs/coronet/index.html, last accessed Nov. 2010.