

Energy Trade-offs among Content Delivery Architectures

Anja Feldmann Andreas Gladisch Mario Kind
Christoph Lange Georgios Smaragdakis Fritz-Joachim Westphal

Abstract—It is envogue to consider how to incorporate various home devices such as set-top boxes into content delivery architectures using the Peer-to-Peer (P2P) paradigm. The hope is to enhance the efficiency of content delivery, e.g., in terms of reliability, availability, throughput, or to reduce the cost of the content delivery platform or to improve the end user experience. While it is easy to point out the benefits of such proposals they usually do not consider the implications with regards to the energy costs.

In this paper we explore the energy trade-offs of such P2P architectures, data center architectures, and content distribution networks (CDNs) by building upon an energy consumption model of the transport network and datacenters developed in the context of Internet TV (IPTV). Our results show that a CDN within an ISP is able to minimize the overall power consumption. While a P2P architecture may reduce the power consumption of the service provider it increases the overall energy consumption.

I. INTRODUCTION

According to the Climate Group report SMART 2020 [1] the Information and Communication Technology (ICT) sector is currently responsible for 2% of the global emissions. Moreover, it is expected to grow at a rate of 6% per year (long-term predictions claim a 10% growth [2]). The report also points out that going forward a rethinking of how we optimize for energy efficiency is needed. Indeed, energy usage has not only become a top level political item [3] but also a hot topic in networking and systems research as underlined by a number of workshops including Hotpower'08, GreenMetrics'09, Green Networking'10 among others.

The main approaches for saving energy in the context of the Internet include turning off unused devices [4], [5], aggregating traffic streams [5], [6], adapting rates [6], network planing and configuration [7] and consolidating usage, e.g., via virtualization [8] or migration [9]. In addition, whenever mobile devices are involved, e.g., in wireless or sensor networks,

Financial support was granted by Deutsche Telekom AG. Anja Feldmann is with Deutsche Telekom Laboratories and Technical University of Berlin, Ernst-Reuter-Platz 7, 10587 Berlin, Germany; e-mail: anja.feldmann (at) telekom.de. Andreas Gladisch is with Deutsche Telekom Laboratories, Goslarer Ufer 35, 10589 Berlin, Germany; e-mail: andreas.gladisch (at) telekom.de. Mario Kind is with Deutsche Telekom Laboratories, Goslarer Ufer 35, 10589 Berlin, Germany; e-mail: mario.kind (at) telekom.de. Georgios Smaragdakis is with Deutsche Telekom Laboratories and Technical University of Berlin, Ernst-Reuter-Platz 7, 10587 Berlin, Germany; e-mail: georgios.smaragdakis (at) telekom.de. Christoph Lange is with Deutsche Telekom Laboratories, Goslarer Ufer 35, 10589 Berlin, Germany; e-mail: christoph.lange (at) telekom.de. Fritz-Joachim Westphal is with Deutsche Telekom Laboratories, Goslarer Ufer 35, 10589 Berlin, Germany; e-mail: fritz-joachim.westphal (at) telekom.de

energy efficiency is one of the main device and protocol design criteria [10].

In order to determine how to best reduce energy consumption in the Internet it is crucial to consider how it is actually used today, i.e., the application mix. Recent studies of the Internet application mix claim that P2P is the most dominant application [11], [12], [13]. However, even more recent studies show that P2P is on the decline and that the Web is regaining its ground [14] with contributing more than 50% to the overall traffic. Among the reasons are the popularity of streaming content, e.g., videos from YouTube, as well as of direct download provider, e.g., RapidShare. This highlights that content distribution is an important contributor to todays Internet traffic. Moreover, the popularity of user generated content as well as advent of High Definition TV (HDTV) means that likely to increase [15]. However, there are multiple ways to distribute content and their relative popularity can change, e.g., from the Web to P2P back to the Web. This motivates us to study the energy trade-offs of content delivery architectures.

Possible architectures within the design space are based on the following concepts

Data center: A centrally managed pool of compute and storage servers with generally good Internet connectivity at a single location.

CDN: A centrally managed pool of compute and storage servers with generally good Internet connectivity distributed at strategically chosen locations throughout the Internet or within an ISP.

P2P architecture: A pool of clients that contribute their decentralized managed compute and storage resources to a distributed content delivery system, e.g., BitTorrent, eDonkey.

There are two recent research trends involving such architectures: (a) to explore how to reduce the energy consumption within data centers, e.g., [16], [17], [18], (b) to incorporate various home devices such as set-top-boxes into content delivery architectures using the P2P paradigm [19], [20], [21], [22], [23]. While the latter appears to be in general feasible it is still an open question on what the energy tradeoffs are [24], [25].

In this paper we build upon the energy consumption models developed by Baliga et al. [26], [27], [28] in the context of IPTV. We extend the model to account for replicas needed in P2P system to guarantee data availability, different content popularity, and take into consideration the network access cost. Our results show that a CDN within an ISP is able to minimize




Fig. 1. Typical IPTV network

the overall energy consumption. However, a P2P architecture may reduce the energy consumption of the service provider even further. Yet, it increases the overall energy consumption.

This opens the question on who pays the bill? Given that energy should be viewed as a social utility and that the leaders of the top eight industrial countries have agreed to try to limit global warming to just two degree Celsius above pre-industrial levels by 2050 [3] it seems inappropriate to migrate the problem from the ISPs to the home users.

II. NETWORK SCENARIO

In order to develop an energy model we have to understand the network components and their energy consumptions. In this paper we focus on the DSL access network and use the IPTV architecture as an example network architecture as it is currently the only one involving the set-top boxes.

A. IPTV network components

An IPTV network typically consists of three parts (a) the storage and server components inside the network to store the content, (b) the transport network which includes the access, aggregation, and backbone networks, and (c) the set-top-boxes for delivering the content to the end-user. Figure 1 gives an example of a typical IPTV network of a major European ISP utilizing VDSL as access technology.

We assume that the set-top-boxes are connected to the access network via VDSL. VDSL is an asymmetric technology and typically offers downstream bandwidth of 25 Mbps, upstream bandwidth of 5 Mbps, and is managed by the DSLAM. The access network utilizes a redundant two stage aggregation network based on carrier grade Ethernet-aggregation switches. Basic management service including termination of the users PPP sessions is provided via a Broadband Network Gateway (BNG) / Broadband Remote Access (BRAS) devices. The BRAS is connected to the IP backbone via an edge router (LER). The IP backbone itself utilizes core routers. The backbone is built utilizing optical transport technology (OTN) on the basis of wavelength multiplexing (WDM).

The content of the IPTV network is hosted in a well provisioned data center. Such a data center is connected to the IP backbone via another LER. In addition, some IPTV platforms utilize caching and/or redistribution servers located in the aggregation network.

Given the different tasks of the different components it is natural that they offer different network capacities, ranging from the throughput of WDM/OTN equipment of $C_{WDM} = 3.2$ Tbps to the upload capacity of set-top boxes with $C_{SRB}^{up} = 5$ Mbps (the set-top-box capacity is limited by the DSLAM upload capacity). Table I contains a list of the typical capacities of the above mentioned components.

TABLE I
IPTV NETWORK COMPONENTS X AND THEIR CAPACITY IN $Gbps$, THEIR POWER CONSUMPTION P_X IN KW AS WELL AS THEIR ENERGY CONSUMPTION PER BIT ECb IN μWs .

	Equipment	Capacity	Power	ECb
Access	DSLAM (VDSL)	$C_{DSLAM}^{up}=0.005$	$P_{DSLAM}=0.0035$ (consumption per subscriber)	0.70
		$C_{DSLAM}^{down}=0.025$		0.14
Metro	Switch	$C_{ES}=320$	$P_{ES}=3.55$	0.0111
	BRAS	$C_G=8$	$P_G=1.1$	0.1375
	Edge Router	$C_{PE}=140$	$P_{CE}=2.7$	0.0193
Core	Core Router	$C_C=640$	$P_C=9.15$	0.0143
	WDM/OTN	$C_{WDM}=3200$	$P_{WDM}=22.6$	0.0071
Data Center	LAN	$C_{ES}=320$	$P_{ES}=3.55$	0.0111
	Edge router	$C_{PE}=140$	$P_{PE}=2.7$	0.0193
	Server	$C_{SR}=0.8$	$P_{SR}=0.35$	0.4375
Set-top box	Storage	$S_{SD}=604800$	$P_{SD}=4.9$	10^{-5}
	Server	$C_{SRB}^{up}=0.005$	$P_{STB}=0.03$	6.0
	Server	$C_{SRB}^{down}=0.025$		1.2
Storage	$S_{SDB}=2560$	10^{-5}		




Fig. 2. Power consumption per user in DSLAM.

B. Power consumption of IPTV components

To quantify the power consumption of the components of the above network we measured the power consumption of each of the components. Since the power consumption depends on the equipment vendor we studied multiple devices from multiple vendors. While there is a huge variation in the power consumption P of different components the differences between vendors are small. Therefore, Table I only contains typical values of the base power consumption for each class of devices. In addition, we also compute the energy consumption for communication by calculating the Energy Consumption per bit (ECb), see last column of Table I.

Next, we discuss how the power consumption profile scales. **DSLAMs:** While one might expect that the power consumption per user is constant this is not the case for DSLAM line cards. Figure 2 shows how the power consumption per subscriber changes with the number of subscribers. With less than one hundred users the base power consumption dominates. On the other hand the power consumption per user is almost constant once the utilization of the line card exceeds 30%. Given that we assume a well designed network we focus on the case of well utilized DSLAMs [29].

Routers: Consistent with the results from Chabarek et al. [7] we found that power consumption is basically scales with the number of line-cards, number of ports, and number of acti-

vated subsystems. While power consumption is proportional to the traffic load this dependency is minimal. Therefore, it is justified to ignore this part in our initial energy study.

Data Centers: Typically, data centers are usually very well designed for redundancy, both in terms of physical protection as well as server components and cooling, but not necessarily for energy efficiency. Nevertheless, it is currently possible to achieve Power Usage Efficiencies (PUE), i.e., (total facility equipment)/(IT equipment power), of less than two [17]. In order to not give an unfair advantage to data centers we use the PUE of 2 for the rest of the paper.

Set-top-boxes: Currently, set-top-boxes are low-end user device with integrated hard-disk. They are not yet optimized for low power transmission, storage maintenance and/or retrieval. Therefore, while the power consumption itself is small the ECb is rather high, especially for the upstream. Moreover, significant power savings are possible by enabling the sleep modus of the device or turning them off [24].

III. ENERGY CONSUMPTION OF CONTENT DISTRIBUTION ARCHITECTURES

Given our understanding of the energy consumption of individual network components we now develop energy models for various content distribution architectures.

A. Content distribution architectures

In order to develop the energy model we first have to revisit the content distribution architectures to see how (a) the content is actually delivered to the users and (b) how reliability and availability of the content is assured.

Data centers: The design of the data center architecture itself should guarantee reliability and availability. Nevertheless, we assume that there are two copies of each object (original and backup) for load balancing and/or redundancy. Typically, the data center is located in the core of the network.

CDNs: Within the scope of this study we focus on a single ISP. Accordingly, we assume that the CDN has access to two (for redundancy) identical CDN servers/caches at each location. Given that each server has limited storage capacity we presume that the CDN will push popular objects close to the users while unpopular objects are pulled from the centralized origin servers using similar techniques than those used by existing CDNs, such as Akamai, Limelight, or Coral [30].

P2P: While the architectures for data center and CDN are fairly simple there are many possible candidates for setting up a P2P based architectures on the set-top-boxes. They differ regarding the number of maintained replicas and the mechanisms used for handling churn. The first fundamental choices is to use a structured or unstructured approach.

In structured P2P systems based on distributed hashing (DHT), e.g., Chord [31], each object has an ID and each set-top-box is responsible for a range of IDs. Each box maintains $O(\log N)$ pointers to other boxes where N is the population of the set-top-boxes in the network. In order to handle churn it has been proposed that each box also stores replicas of the predecessor and the successor ones. To enable fast access,

load balancing, and fast replication under node failures or departures the proposal is to maintain more than the required minimum of three replicas [32]. To reduce the per box load for popular files and the overhead imposed by churn multiple boxes might be responsible for one ID and boxes might be responsible for multiple ID ranges [33]. In addition, each object can be split in many pieces that are distributed across multiple boxes [34]. Accordingly, our evaluation studies the cases of 2, 10, and 100 replicas for a population of 10,000 users (set-top-boxes). Note that we refer to the total number of replicas in the P2P network. Portions of the replicas can be stored in different set-top-boxes.

In unstructured networks, e.g., Bubblestorm [35], one typically maintains \sqrt{N} replicas of each object [36] to enhance searching, improve availability of objects, and balance node load.

B. Energy Consumption Models

To derive the energy consumption models for the different content distribution architectures we combine the model proposed by Baliga et al. [27] with a detailed model of the access network (ADSL/VDSL technology). For this we assume that B is the size of object (in bits) and D is the access frequency (in number of downloads per hour).

We start with the data center architecture. The two key parameter are R the number of servers in the network and H the average core network distance to the data center (in number of hops). We only need the core network distance as the power consumption of the access network is captured separately. Typical values are two replicas ($R = 2$) and a hop distance less than 12 ($H = 12$ due to the high redundancy). We refer to this model as DC. The energy E consumed to download a single file is:

$$E = \frac{B}{3600} \frac{P_{DSLAM}}{C_{DSLAM}^{down}} + \frac{B \cdot R}{D} \frac{P_{SD}}{S_{SD}} + \frac{4B}{3600} \left(\frac{P_{SR}}{C_{SR}} \right) + \frac{4B}{3600} \left(\frac{3P_{ES}}{C_{ES}} + \frac{P_G}{C_G} + \frac{2 \cdot P_{PE}}{C_{PE}} + \frac{(H+1)P_C}{C_C} + \frac{H \cdot P_{WDM}}{C_{WDM}} \right) \quad (1)$$

Interestingly, the same model also captures a distributed caching architecture (DCache) just with other parameters. The number of servers is larger, e.g., $R = 10$, while the replicas are closer to the access network and thus the hop count is smaller, e.g., $H = 3$. Next, we note that a CDN is a hybrid between a caching and a data center architecture (depending on the popularity of the object). As such it is the minimum of the two parameter settings of the data center and the caching architecture.

The peer-to-peer energy model differs as the data has to traverse the access network twice. On the other hand, the costs for the data center are eliminated. Assuming a reasonable locality of peers we get an average hop distance of $H = 3$. Regarding the number of replicas this depends on the specifics of the P2P network as discussed above. We use values of $R = 2, 10, 100$ to reflect a basic P2P network (P2P-base), a DHT based P2P

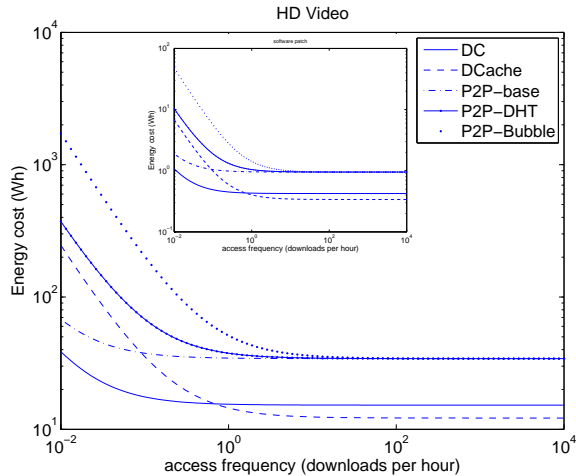


Fig. 3. Energy costs for downloading a single movie ($B = 1.8$ GB) or patch ($B = 50$ MB, small plot) across access frequency.

network (P2P-DHT) with redundancy to handle churn, and an unstructured P2P network with higher redundancy (P2P-Bubble). Note, that our energy model includes all energy costs, especially those of the set-top-boxes, and not only the costs of the ISP, namely the network and server costs. The energy consumed to download a single file is:

$$E = \frac{B}{3600} \frac{P_{DSLAM}}{C_{DSLAM}^{up}} + \frac{B}{3600} \frac{P_{STB}}{C_{SRB}^{up}} + \frac{B \cdot R}{D} \frac{P_{STB}}{S_{SDB}} + \frac{4B}{3600} \left(\frac{3P_{ES}}{C_{ES}} + \frac{P_G}{C_G} + \frac{2 \cdot P_{PE}}{C_{PE}} + \frac{(H+1)P_C}{C_C} + \frac{H \cdot P_{WDM}}{C_{WDM}} \right) \quad (2)$$

IV. ENERGY COST TRADEOFFS

To explore the energy trade-offs of the various different content distribution architectures we start our evaluation with two use cases: download of a high definition movies of size $B = 1.8$ Gigabytes and a software update or patch of size $B = 50$ Megabytes.

A. Per download energy cost

The initial step is to understand how the energy consumption scales with the popularity of the objects under different content distribution architectures. Accordingly, Figure 3 shows the energy consumption per download (y-axis) as the popularity of the file (x-axis) increases for both use cases and 10,000 users. The first observation is that the energy cost decreases with increased object popularity. The main reason for this is that caches become effective and the cost of storage is amortized.

For unpopular objects the data center architecture is the most energy efficient one. For popular objects the DCache architecture is the best in terms of energy consumption. Indeed, there is a threshold in terms of number of downloads when it is better to cache an object closer to the user rather than pulling it from the data center. This is the threshold that any CDN tries to realize. Accordingly, the CDN architecture which

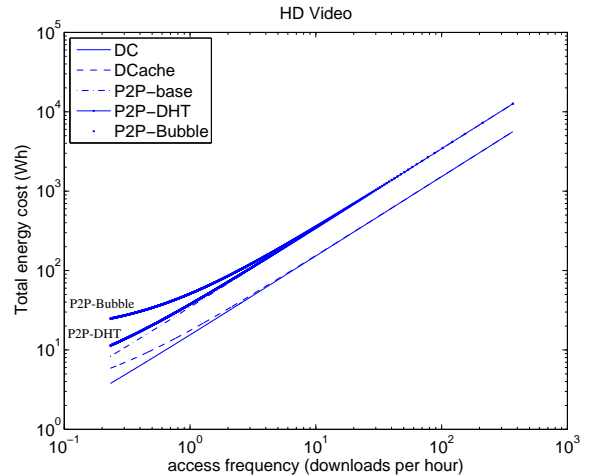


Fig. 4. Total energy costs for downloading movies (popularity Zipf distribution $\alpha = 0.8$) across access frequency.

combines the advantages of the DC and DCache architecture is the overall winner. Even if we change the specific parameters of the architectures, R or H , the overall observations hold.

We note that all P2P architectures converge to roughly the same energy costs per download. However, for unpopular objects the costs for keeping the replicas are substantially. We also observe that even the ideal P2P architecture does not seem to be able to beat the CDN even though our energy model for the P2P architecture does not even include the cost of maintaining the P2P network. The main reason is that each object has to be transferred over two DSLAMs. Due to the limited upload capacity of each peer this increases the usage of the DSLAM which increases the energy costs. Given that our energy model assumes well utilized DSLAMs this impact is kept reasonable at the cost of reducing the available benefits of the CDN architecture. Our observations are consistent for different file sizes ranging from 1.8 Gigabytes to 50 Megabytes.

B. Total energy costs

So far we focused on the cost of a single object. But to justify a content distribution architecture we have to look at multiple objects. Regarding their popularity measurement studies have shown that it is consistent with a Zipf distribution, e.g., for channels in IPTV [37] with $\alpha \sim 0.8$.

Accordingly, Figure 4 shows the overall energy consumption for a population of 10,000 users and 10,000 objects with an object popularity of $\alpha = 0.8$. As expected the energy consumption is dominated by the popular objects. Accordingly, it is not surprising that both data center and DCache architectures are more energy efficient than the P2P architectures. As such we again expect the CDN architecture to be the most energy efficient one.

To explore the impact of the differences in popularity we vary the α of the Zipf distribution and compute the total energy consumption. Figure 5 shows the normalized total energy consumption. Normalization is done with respect to the energy




Fig. 5. Total energy costs for downloading movies across α normalized by the energy consumption of the CDN architecture.

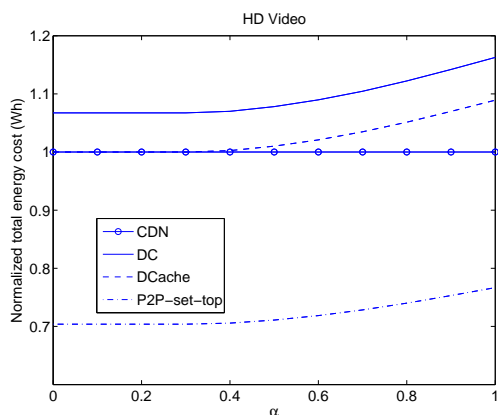


Fig. 6. Total energy costs for downloading movies across α normalized by the energy consumption of the CDN architecture (without the energy costs for the set-top-boxes).

consumption of the CDN architecture. In general we note that the skewness of the popularity distribution has some albeit limited impact. Note, how close the energy consumption of DCache is to the optimal one. Even the data center architecture is only 17% off. The main reason is that the energy costs of the network is only one contributor to the overall energy costs. However, the energy consumptions of all P2P architectures are significantly higher. Thus we conclude that a CDN is the best architecture with regards to the overall energy consumption.

C. Total ISP energy costs

So far our energy models for the P2P architecture included the energy costs for the set-top-box. However, this costs do not have to be paid by the ISP. Therefore, the actual energy cost of an ISP that uses a set-top-box based P2P system are significantly lower. Figure 6 again shows the normalized total energy consumption. But this time for CDN, DC, DCache, and P2P-set-top (without the energy costs for the set-top-boxes). For this scenario the P2P architecture shows significant potential. The possible energy savings range from 25% to 30%. However, this specific P2P architecture is only based on two replicas. As the number of replicas is increased to handle churn

the benefits decrease. Moreover, recall that our current P2P energy model does not account for the cost of updating and maintaining replicas which is left to future work.

V. DISCUSSION AND CONCLUSION

In this paper we study the energy tradeoffs of three different content delivery architectures: Data centers, CDNs, and P2P systems. We find that in terms of total energy costs CDNs are the clear winner. However, an ISP can potentially minimize its energy costs by incorporating the users set-top-boxes into a P2P architecture. However, we believe that it is important to minimize the overall energy consumption rather than migrating the problem to the end users. Moreover, in the future it is quite likely that users will turn off their set-top-boxes so that the P2P architecture will have to cope with high churn which imposes extra energy costs for maintaining additional replicas.

So far our energy model assumes that the energy costs are the same throughout. However, we plan to investigate to what degree it is possible to further reduce the energy costs of an ISP that uses a CDN architecture by distributing the load across his data centers based on the dynamic prices for energy in a similar manner as proposed by Qureshi et al. [38]. Our future research agenda also includes the study of the corresponding energy costs when new access technologies such as fiber-to-the-home are launched.

REFERENCES

- [1] "Smart 2020: enabling the low carbon economy in the information age," <http://www.theclimategroup.org/assets/resources/publications/Smart2020R%eportSummary.pdf>.
- [2] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Devellder, D. Colle, B. Dhoedt, and P. Demeester, "Worldwide Energy Needs for ICT: The Rise of Power-aware Networking," in *Proc. of Advanced Networks and Telecommunication Systems*, 2008.
- [3] "BBC article on G8 summit, July 2009," <http://news.bbc.co.uk/2/hi/europe/8141352.stm>.
- [4] Y. Agarwal, S. Hodges, R. Chandra, J. Scott, P. Bahl, and R. Gupta, "Somniloquy: Augmenting Network Interfaces to Reduce PC Energy Usage," in *Proc. NSDI*, 2009.
- [5] M. Gupta and S. Singh, "Greening of the Internet," in *Proc. ACM SIGCOMM*, 2003.
- [6] S. Nedeveschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," in *Proc. NSDI*, 2008.
- [7] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power Awareness in Network Design and Routing," in *Proc. IEEE INFOCOM*, 2008.
- [8] L. Liu, H. Wang, X. Liu, X. Jin, W. B. He, Q. B. Wang, and Y. Chen, "GreenCloud: a new architecture for green data center," in *Proc. of ICAC-INDST*, 2009.
- [9] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford, "Virtual routers on the move: live router migration as a network-management primitive," in *Proc. ACM SIGCOMM*, 2008.
- [10] C. E. Jones, K. M. Sivalingam, P. Agrawal, and J. C. Chen, "A Survey of Energy Efficient Network Protocols for Wireless Networks," *Wireless Networks*, vol. 7, no. 4, pp. 343–358, 2001.
- [11] A. Parker, "CacheLogic: P2P Media Summit."
- [12] H. Schulze and K. Mochalski, "Ipoque: Internet Study 2007," <http://www.ipoque.com/resources/internet-studies/internet-study-2007>.
- [13] —, "Ipoque: Internet study 2008/2009," http://www.ipoque.com/resources/internet-studies/internet-study-2008_20%09.
- [14] G. Maier, A. Feldmann, V. Paxson, and M. Allman, "On dominant characteristics of residential broadband internet traffic," in *Proc. ACM Internet Measurement Conference*, 2009.
- [15] "Cisco Visual Networking Index," 2009.

- [16] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *Proc. ACM SIGCOMM*, 2008.
- [17] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks," *ACM Comp. Comm. Review*, vol. 39, no. 1, pp. 68–73, 2009.
- [18] J. Liu, F. Zhao, X. Liu, and W. He, "Challenges Towards Elastic Power Management in Internet Data Centers," in *IEEE ICDCS WCPs workshop*, 2009.
- [19] V. Valancius, N. Laoutaris, L. Massoulie, D. Christophe, and P. Rodriguez, "Greening the internet with nano data centers," in *ACM CoNEXT*, 2009.
- [20] J. He, A. Chaintreau, and C. Diot, "A performance evaluation of scalable live video streaming with nano data centers," *Comput. Netw.*, 2009.
- [21] N. Laoutaris, P. Rodriguez, and L. Massoulie, "Echos: edge capacity hosting overlays of nano data centers," *ACM Comp. Comm. Review*, 2008.
- [22] "Vudu," <http://www.vudu.com>.
- [23] X. Yang, M. Gjoka, P. Chhabra, A. Markopoulou, and P. Rodriguez, "Kangaroo: Video Seeking in P2P Systems," in *USENIX IPTPS*, 2009.
- [24] G. Lefebvre and M. J. Feeley, "Energy Efficient Peer-to-Peer Storage."
- [25] S. Nedeveschif, S. Ratnasamy, and J. Padhye, "Hot data centers vs. cool peers," in *Proc. USENIX NSDI Workshop on Power Aware Computing and Systems*, 2009.
- [26] J. Baliga, R. Ayre, W. Sorin, K. Hinton, and R. Tucker, "Energy Consumption in Access Networks," in *Conference on Optical Fiber communication/National Fiber Optic Engineers Conference*, 2008.
- [27] J. Baliga, R. Ayre, K. Hinton, and R. Tucker, "Architectures for Energy-efficient IPTV Networks," in *Conference on Optical Fiber communication/National Fiber Optic Engineers Conference*, 2009.
- [28] J. Baliga, K. Hinton, and R. Tucker, "Energy Consumption of the Internet," in *Optical Internet, 2007 and the 2007 32nd Australian Conference on Optical Fibre Technology. COIN-ACOFT*, 2007.
- [29] C. Lange, D. Kosiankowski, C. Gerlach, F.-J. Westphal, and A. Gladisch, in *European Conference on Optical Communication*, 2009.
- [30] M. J. Freedman, E. Freudenthal, and D. Mazières, "Democratizing content publication with Coral," in *Proc. NSDI*, San Francisco, California, 2004.
- [31] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *Proc. ACM SIGCOMM*, 2001.
- [32] B.-G. Chun, F. Dabek, A. Haeberlen, E. Sit, H. Weatherspoon, M. F. Kaashoek, J. Kubiawicz, and R. Morris, "Efficient Replica Maintenance for Distributed Storage Systems," in *Proc. NSDI*, 2006.
- [33] B. Godfrey, K. Lakshminarayanan, S. Surana, R. Karp, and I. Stoica, "Load Balancing in Dynamic Structured P2P Systems," in *Proc. IEEE INFOCOM*, 2004.
- [34] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica, "Wide-area cooperative storage with CFS," in *Proc. ACM SOSP*, 2001.
- [35] W. W. Terpstra, J. Kangasharju, C. Leng, and A. P. Buchmann, "Bubblestorm: resilient, probabilistic, and exhaustive peer-to-peer search," in *Proc. ACM SIGCOMM*, 2007.
- [36] C. Leng, W. W. Terpstra, B. Kemme, W. Stannat, and A. P. Buchmann, "Maintaining Replicas in Unstructured P2P Systems," in *Proc. ACM CoNEXT*, 2008.
- [37] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, "Watching television over an IP network," in *Proc. ACM Internet Measurement Conference*, 2008.
- [38] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electrical bill for internet-scale systems," in *Proc. ACM SIGCOMM*, 2009.