
Low Carbon Economy Considerations in Designing and Operating Content Delivery Networks for VoD Services

Dragan Boscovic¹, Michael Needham¹, Faramak Vakil¹
and Jin Yang²

¹*Motorola Applied Research Center, Schaumburg, IL 60196, USA;
e-mail: dragan.boscovic@motorola.com*

²*Motorola Applied Research Center, Beijing, China*

Abstract

The advances in information, computer technology and broadband connectivity have eased the way to reach desired content. This is especially true for the Internet TV (IPTV) services such as Video-on-Demand (VoD). This has brought forth many technical challenges regarding the distribution of electronic content, such as how to send large video files, how to deal with the stream load from large numbers of users, and how to transmit the video stream to a global network over a long distance. To address this, Content Delivery Networks (CDN) have been brought into the overall VoD system design. CDNs are based on a system of data centers containing servers and massive storage that places copies from the content library closer to the end user. However, this trend brings some unfavorable conditions in terms of energy usage; worldwide, data centers account for a significant percentage of the total energy consumed by ICT industry. Better energy efficiency of data centers would significantly cut the CO₂ emission from ICT operations. As a CDN is an assembly of mutually interconnected data centers the overall efficiency has to factor for efficiency of each CDN geographical node/data center. The aggregate energy efficiency could not be simple multiplication of individual data centers/servers involved but has to introduce corrective factors to account for energy efficiency tradeoffs brought about by specific geograph-

ical CDN topology and how effectively the additional resources deployed and energy consumed by them helps improve content access performance of a given CDN design. In this article, we aim to shed additional light onto the technical challenges and recommend possible improvements as related to the energy consumption within a given CDN topology.

Keywords: Content Delivery Network (CDN), Video on Demand (VoD), Data Center Performance Efficiency (DCPE).

1 Introduction

The advances in information, computer technology and broadband connectivity have eased the way to reach desired content. TV service is not immune to these changes and watching television is today increasingly switching to an on-demand type of experience. This is especially true for the Internet TV (IPTV) services, requiring numerous precisely orchestrated technologies to deliver a quality viewing experience and services such as Video-on-Demand (VoD), live TV, private video recording (PVR) capabilities, time-shifted television (TSTV), Near VoD, Subscription VoD services, etc. This transition to the time shifted viewing has also brought forth many technical challenges regarding the distribution of electronic content, such as how to send large video files, how to deal with the stream load when all users are ordering the same or different films, how to transmit the video stream to a global network over a long distance. It is where a Content Delivery Network (CDN) comes into the overall VoD system design in order to resolve these issues.

CDN is a distributed content management and delivery system overlaid on the data network. It is based on a system of servers with massive storage that places copies from the content library closer to the end user so as to maximize the available bandwidth, and consequently reduces the data access times. A farm of servers provides computations, storage, applications and data transfer and as such are central to the CDN architecture, from which all content is sourced or passed through. For the case of a video CDN, as the popularity and penetration of video and time shifted TV viewing increases so does the need for additional clusters of video servers.

Centralized hosting model offers easier resource optimization and management by the virtue of virtualization technologies. However, the centralization trend brings some unfavorable conditions in terms of energy usage as data centers are normally over provisioned to for peak demand and require a dedicated temperature management system. Worldwide, data centers con-

sume around 40,000,000,000 kW/hr of electricity per year that amounts to 18% of the total energy consumed by ICT industry. Big portion of energy consumption, roughly 52% is wasted due to equipment inefficiencies, non-ideal design and poorly implemented algorithms and management schemes. The need for better energy efficiency of data centers is clear and would significantly cut the CO₂ emission from ICT operations projected to be around 1.5 billion of tons in 2020. The task of improving CDNs efficiency is complex and can only be achieved through better operational exploitation of individual parts and by simpler architectures and intelligent content management algorithms that will bring the energy cost into the overall performance optimization function.

In this article, we are to shed additional light onto the technical challenges and recommend possible improvements as related to the energy consumption within a given CDN topology. The issues such as datacenter cooling, power conversion or proportional computing are generic to the topic of Green Data Centers and as such not discussed here, instead this paper focuses on the topology of CDN overlay networks and related caching and replication algorithms as main domains to pay attention to when designing next generation of green CDNs.

2 State of the Art

Use of CDNs in context of the Internet is a well known phenomenon; it reduces the content acquisition time by strategically placing servers at the network edge [7]. The content selected on basis of carefully designed policies, which usually calculate content popularity in terms of temporal and spatial viewing patterns, is replicated and propagated to the edge servers. In comparison to serving all the requested content from the back end servers, this hierarchically decentralized approach brings significant advantages in terms of access time and scalability performance. These benefits will normally come at the cost of using more servers and needing more complex management algorithms to support the predictive content replication, updates and distribution.

As mentioned earlier, in the context of CDNs, the placement of video servers is an important consideration and an alternative to the geographically distributed servers is the centralized, geographically collocated server placement. Variation on the theme of centralized video servers has been thoroughly investigated [10] centralized video servers with distributed video

buffers [3], two- or multitier hierarchical distributed video servers [12], and fully replicated distributed video servers.

Equally, the energy efficient operation of datacenters has attracted a lot of attention in the recent years. Recent reports [8] show that data centers consume large amounts of energy for energy transmission/conversion, and most importantly, cooling and heat dissipation. Such overheads are captured by the Power Usage Efficiency (PUE) metric, which is defined as the ratio between the total power consumed by a data center and the power actually delivered to its IT equipment (servers, networking equipment). Efforts have been cited related to agreeing and enabling standard efficiency metric (Green Grid metrics: Describing Data Center Power efficiency, October 2007), real time measurement systems, modeling energy efficiency, suggesting optimal designs (rightsizing, efficient IT equipments, efficient cooling and NCPI equipments) [9] and developing sophisticated algorithms for designing and managing the datacenters. In [1] scalable fat tree commodity datacenter network architecture is proposed that has an efficient advantage over the traditional hierarchical enterprise structure. Furthermore, in [13], a cost effective routing algorithm is presented that moves workload to datacenters residing in different geographies to minimize energy cost and consumption.

Diametrically opposed to the centralized approach are P2P topologies that might offer scalability and cost effectiveness for certain types of CDN applications. In Nano Datacenters concept, content and service logic are stored on home gateways and the access to these services on gateways is provided by using a managed P2P infrastructure. This topology greatly reduces the need for content servers and network resources and extensive simulations has confirmed at least 20–30% in energy savings as compared to the energy spent by legacy data centers [15].

3 VoD Solution

3.1 Service Architecture

The simplified generic architecture of a VoD system is shown in Figure 1. The business ecosystem consists of the four main players, namely, the network provider, the program provider, the service provider, and the user. The user generates requests to the service provider, who will obtain the necessary material from the program providers and deliver it to the user over the infrastructure provided by the network provider. Thus, the service provider acts as an agent able to access various types of content and aggregate it according




Figure 1 Simplified generic VoD solution architecture.

to the users' preferences. It is possible for the network, program, and service providers to be the same entity, but, in general, they are more likely to be distinct organizations. Lately, this business system has been undergoing further evolution, since anyone with marketable content can offer his services directly to the user as Over the Top (OtT) services using the open nature of the Internet.

3.2 Technology and Business Trends

In order to fully understand the dynamics of techno economical forces shaping current solutions for Content Delivery Networks, Video-on-Demand Systems and allocating databases in distributed computing systems, it is very important to take notice of the following trends:

- There is a clear migration towards general purpose computing hardware by VoD vendors. Focus and innovation are in terms of software architecture and service enablers, and progressive vendors are relying

on commercial off-the-shelf (COTS) hardware and open API software to maintain their competitiveness.

- The second major shift is the transition from proprietary content management to a stand-alone content delivery network (CDN) addition to the integral VoD content propagation solution. CDN-based content management systems enable rapid ingestion and content propagation to any device from any location. At the same time it optimizes bandwidth utilization by locating and using the most adequate source and deploying transcoding, if needed, to match the capabilities of the device being served.
- There is strong evidence of growing importance of the VoD back office in supporting switched digital video (SDV) and on-demand content delivered to wireless devices, in addition to VoD session management. This expanded role of the back office is setting a trajectory by which the back office systems could become the most valuable component of any VoD system.
- The most critical parameter for the overall performance of a given VoD solution, in terms of end user experience, remains to be the download/display delays. This is to some extent relieved by the continued fierce competition among service providers offering high-speed Internet as larger bandwidth reduces the delivery delay.
- IPTV and traditional VoD services are being blended together to combine their respective benefits. Infrastructure components that exclusively serve VoD are being swapped for common ICT platforms and related services can be offered at relatively small incremental cost over the current IPTV or VoD solutions.

3.3 Content Delivery Networks

CDN nodes are normally deployed in multiple locations, linked into a virtual network by a web of different backbone connections. The control mechanisms provide logic for cooperation between servers to respond to the requests for content by end users, transparently moving content between the nodes to optimize the delivery process. Optimization can yield reduced bandwidth and/or local storage costs, shorter access/load times or improved global availability of content. The number of CDN data centers and servers varies, depending on the chosen topology; certain deployments number thousands of nodes with tens of thousands of servers in many remote Points of Presence (PoP), while others can build a global network with fewer geographical PoPs.




Figure 2 Hierarchical layered structure of a traditional communication network.

As shown in Figure 1, a video server system cloud is crucial for the overall design. It is the configuration and topology of this cloud of video servers that form a CDN network that is in turn a critical component for IPTV/VoD solution. Network topology can come in many different flavors that can be classified into three basic categories: centralized, hierarchical decentralized and hybrid. In the centralized topology all requests will be sent to be processed at the same location, usually called Headend. The distributed server system distributes requests to many sites, located closer to the users, thus alleviating the congestion in the network and the bottleneck due to the central server.

Let us now briefly present the most common of CDN topologies based on the location where the VoD servers are placed. Traditional telecommunication networks, as depicted in Figure 2, are based on a hierarchical approach, consisting of three main components core network, service routing level 1 (SR1) and service routing level 2 (SR2). Users access the services using different access technologies that are not the subject of study here.

Centralized Architecture. Video servers are organized into a single cluster which is connected through the SR1 network and are distributed to the entire network through this unique location (the Headend). Video traffic will be concentrated near the Headend and so the solution will suffer from lack of scalability.




Figure 3 Hybrid CDN architecture.

Hierarchical Architecture. This is basically a CDN play. Scalability can be improved by placing video servers in the SR1 region, or in SR1 and SR2 regions. Video servers are organized according to a tree structure having the Headend as root. When a customer wants to see a video, the request is processed by the nearest video server and, if it fails, it is then sent to the server of the upper hierarchical level.

Beside scalability performance, the CDN topology also shortens the content access times.

Additionally, the use of several servers provides redundancy. Despite these advantages, the cost of deploying and maintaining additional servers is the main tradeoff relative to the performance improvements. The additional cost is dictated by the storage capacity of each server. It is why the amount of content needed to be stored on each location is very carefully selected, usually on the basis of the expected popularity index. The more popular content is expected to be, the closer to the user location it is stored. Servers of the lowest hierarchical level store only the most popular videos and the ones that are currently requested by customers.

The popularity index of a given video increases each time a viewer requests it. The least popular videos are removed from the servers to make room for new content with a higher popularity index.

Hybrid Architecture. Figure 3 depicts a CDN with hybrid architecture. In this topology, video servers are still placed in SR1 and SR2 regions, as shown in , but the servers of the SR2 region are organized, for instance, as a Chord ring, forming P2P groups. Spatial distribution of users is not uniform; for that

reason, servers in the SR2 might have different capacities, depending on the number of customers they are responsible for. Sharing the identifier space according to these capacities enables proper resource assignment.

Each SR2 server stores only a partial set of videos and, whenever a viewer requests a video from the assigned SR2 server, the server searches for the video on its peer group, instead of propagating the request up to any of the SR1 servers. This procedure is common across all SR2 servers and consequently reduces the load for SR1 servers.

This CDN topology can be further subdivided according to the storage capacity management.

Single video copy – There is only one copy of the video in the P2P group located at the responsible server. Should a request be issued to another server, the responsible server sends the video directly to the consumer.

Multiple video copies – Multiple copies of the same video can be stored on the SR2 servers. A single server is still responsible for each video; however, a requested video is sent to another serving server rather than to the consumer directly.

4 Energy Efficiency Considerations

It is well known that the carbon emissions of ICT sector is growing at an unprecedented rate, currently data centers consume 1.5% of total US electricity and are projected to surpass the carbon emission of the air transportation industry by 2020 [14]. As the video in its all forms (TV, video on demand, Internet, and P2P) will account for over 91% of global network traffic by 2013 [11] (see Figure 4) it is obvious that overall objective to improving ICT energy efficiency is to become strongly dependent on energy efficiency of CDNs themselves.

Building and operating a large scale CDN for VoD is infrastructure demanding, and the quality of VoD service is likely to be directly proportional to the aggregate processing and storage capacity. The increased demand for more computing capabilities within a given CDNs to more efficiently handle the growth of digital video libraries makes cost a primary metric of interest in the design of the infrastructure. Cost in this case has to account for all the significant components of the solution, including hosting-facility capital and operational expenses that will include power provisioning and energy related costs. Figure 5 depicts a likely distribution of energy consumption across different system components of a modern large scale IT cluster within a given CDN solution [2].




Figure 4 Video growth in the Internet video (PB per month). Source [11].




Figure 5 Likely distribution of peak power usage across IT subsystems of a typical CDN datacenter.

Actual CDN designs, as with any system design, attempt to achieve optimal performance at acceptable overall system costs, including energy consumption. Although the optimal point may vary over time for any given workload, the cumulative behavior across wide spectra of likely workloads will tend to vary relatively slowly. Recognition of this behavior provides a useful input to system design and allows for proper dimensioning of different sub-systems. For example, a media access algorithm may best perform for a given ratio between the number of CDN nodes and overall storage capa-

city; too many nodes makes placement/retrieval of the media files relatively complex and insufficient storage capacity may not take full advantage of the intelligence built into media placement algorithms. If the desired ratio between the overall storage capacity and the CDN nodes/data centers remains constant but over time the energy efficiency of various components evolve at a different pace the overall energy efficiency can change significantly.

A good study case of different energy efficiency improvement rates is that one of CPU and DRAM. Over the past few years CPU energy efficiency has improved at much faster pace than energy efficiency of DRAM and disk storage. CPUs historically used to account for more than 60% of the system energy budget and now have dropped to 45% and keep improving. This is of very specific concern to the CDN designers since the energy usage in a well-balanced CDN design is likely to be dominated by its storage subsystem. The relatively lower energy costs of CPU versus storage might suggest, for instance, transcoding content to different formats in real time only as needed, versus pre-coding and storing all available formats, which, particularly for “long tail” content, might not even be accessed.

The exact distribution of energy consumption within a given CDN would depend on the network topology chosen, and how effectively within that topology the storage can be used by the content replication and caching algorithms.

There are energy efficiency considerations for each of the nodes (data centers or individual servers) within chosen topology. Every new generation of servers requires higher power, mainly driven by a combination of higher processor operating frequencies, larger and faster disk storage and DRAM caches, and faster networking speeds. While this trend of servers with more capable features should continue, equally important is to make sure that available computational resources are fully occupied in order to help energy efficiency. Since server efficiency seriously degrades under light computational load having too much spare capacity in terms of processing power or storage resources is not helping much manage the overall CDN efficiency. For that reason it is important to keep a server and/or storage devices either fully occupied or possibly shut them down. This type of control would require a system for work consolidation that free up and power down entire subsystems of CDN nodes and represent an avenue to create energy-proportional behavior in CDNs. The management aspect of the system for load consolidation are quite complex as it has to deal with the daily variations of load. Even though operators have perfected service load balancing, idle servers are not necessarily switched off due to the need to keep spare capacity readily available.

There are examples of the workload being dynamically moved not between the servers but between datacenters in order to maximize traffic in energy efficient data centers and reduce in less energy efficient centers [6].

Nevertheless, this will not be an easy task as it is in contradiction with distributed system architecture of CDNs which tend to disseminate chunks of content and related computation across many geographically distributed servers in order to achieve high content availability and low access latency. So any energy management mechanism to reduce the availability of sufficiently large idle periods at any one CDNs server will need to carefully balance energy saving objective versus performance degradation.

Energy efficiency of a VoD solution can be broadly defined as the amount of concurrent video streams served at requested quality divided by the total energy used in the process.

This is relatively a broad definition and we are not to make any attempt to define actual measurement metric and reference workloads. In order to get better ideas where to look for energy savings it would be more useful to factor the above efficiency definition into a number of constituting components. In that case each of the constituting components can be independently measured and optimized within related engineering disciplines. Note that this definition could reflect the number of concurrent streams at peak load, or could represent essentially “stream-hours per watt-hours” over an extended time period, reflecting the wide variability in VoD demand hour-to-hour, and day-to-day. Efficiency of a CDN node, be it in a form of dedicated data center or simple server, can be modeled by the following equation:

$$\begin{aligned}
 \text{Efficiency}_{\text{DC}} &= \frac{\# \text{ of concurrent video streams}_{\text{DC}}}{\text{Total Energy}_{\text{DC}}} \\
 &= \frac{1}{\frac{\text{Total Energy}}{\text{Energy to IT infrastructure}}} \\
 &\times \frac{1}{\frac{\text{Energy to IT infrastructure}}{\text{Energy to Computational \& Storage Resources}}} \\
 &\times \frac{\# \text{ of concurrent video streams}}{\text{Energy to Computational \& Storage Resources}} \\
 &= \frac{1}{\text{PUE}} \times \frac{1}{\text{SPUE}} \\
 &\times \frac{\# \text{ of concurrent video streams}_{\text{DC}}}{\text{Energy to Computational \& Storage Resources}}. \quad (1)
 \end{aligned}$$

The above definition is in line with the Green Grids Data Center Performance Efficiency (DCPE) definition (Green Grid metrics: Describing Data Center Power efficiency, October 2007). The first term in the above equation relates to the power usage effectiveness (PUE), and reflects the quality of the building infrastructure hosting the CDN servers. It captures the ratio between total power brought to the building and IT power, the power consumed by the actual computing equipment (servers, storage, networking equipment, etc.).

The second factor in the equation accounts for inefficiencies within the IT equipment itself. Substantial amounts of power may be lost in the servers and network storage power supply, voltage regulator modules (VRMs), and cooling fans, as such SPUE factor reflects these mainly electrical losses. It is calculated as the ratio of total server input power to its useful power, where useful power is qualified as being the power consumed by the electronic components directly involved in the computation tasks, namely motherboard components including CPUs, storage disks, DRAM, I/O cards, fans and other active components found in a server.

Because PUE and SPUE factors are agnostics to the nature of the service and/or processing load, this study will not elaborate on specific energy saving measures as related to them.

Equation (1) accounts for energy used for useful computational output, in case of CDNs it is the number of concurrent streams. No surprise this term is the hardest to measure and will need defining measure metric and specifying reference processing loads to reflect the nature of services offered. Ultimately, in context of video CDNs we would like to measure the amount of concurrent video streams delivered by the network from perspective of the energy invested in the related computational tasks across all the servers/data centers constituting a given CDN. Such a measure should come handy when comparing the relative energy efficiency of two specific CDN designs.

Equation (1) examines whether or not we need to involve additional factors to account for interconnectivity between CDN nodes as well for the efficiency of related activities such as transcoding, content chunking, predictive caching and replication, etc. In order to transition from energy efficiency of a self contained data center to a distributed CDN constellation we invoke efficiency analysis of parallel computing [4].

$$\begin{aligned} \text{Efficiency}_{\text{CDN}} &= \frac{\# \text{ of concurrent streams}_{\text{CDN}}}{\text{Total Energy}_{\text{CDN}}} \\ &= \left[\frac{\alpha}{\beta} + (1 - \alpha) \right] \times \text{Efficiency}_{\text{DC}}. \end{aligned} \quad (2)$$

As a CDN is an assembly of mutually interconnected data centers the overall efficiency has to factor for efficiency of each CDN geographical node/data center. But it has also to take into the account energy efficiency of the CDN topology and efficiency of the content management algorithms such as predictive caching and content replication mentioned earlier. These considerations are taken into account by parameters α and β in Equation (2). Factor β defines the increase ratio in terms of processing and storage resources deployed in a CDN overlay network relative to the reference VoD data center. Equally parameter α represents extent to which these additional resources are used to servicing additional video streams. For example, the border case when $\alpha = 0$ denotes a design in which all additional resources are used to increase the number of serviced video streams; in contrast $\alpha = 1$ corresponds to the design where all additional resources in a given CDN constellation are used to improve content access time and do not contribute to increasing the number of video streams being serviced by the CDN. If we make an assumption that a CDN solution is made of K identical data centers ($\beta = K$) and that all additional resources are used for strategically caching and replicating content closer to the users, then the efficiency of entire CDN constellation will be equivalent to the efficiency of a single constituting Data Center only if there is a commensurate increase in stream count – otherwise, the efficiency will be less. This basically leads to the conclusion that the CDN's access time performance improvements come from the additional processing and storage equipment deployed and the extra energy needing to power them.

$$\max_{\alpha \rightarrow 0, \beta = K} \text{Efficiency}_{\text{CDN}} = \frac{K \times \# \text{ concurrent streams}_{\text{DC}}}{K \times \text{Total Energy}_{\text{DC}}} = \text{Efficiency}_{\text{DC}}. \quad (3)$$

Energy efficiency of a CDN is always lower than that one of a single DC streaming same # of concurrent videos.

Equation (2) captures the tradeoff between CDN performance in terms of the content access time and its overall energy performance for a given overall number of video streams served. The two factors α and β are important parameters of the CDN's genetic code; as such they capture the related computation performance, the CDN's overlay network topology and account for the efficiency of caching and replication algorithms.

Computation servicing video streaming: Clearly, the same application binary can consume different amounts of power depending on the servers and CDN architecture and, similarly, an application can consume more or less of a servers capacity depending on software performance tuning. Disk drives,

for example, spend a large fraction of their energy budget simply keeping the platters spinning, possibly as much as 70% of their total power for high RPM drives. Creating additional energy efficiency and energy proportionality may require smaller rotational speeds, smaller platters, or designs that use multiple independent head assemblies. Pioneering work in this domain is provided by Carrera et al. [5] who considered the energy impact of multispeed drives and proposed combining server-class and laptop drives in order to achieve proportional energy behavior. Significant energy savings might be available by more effectively controlling and utilizing inactive low-power modes such as spun-down disk drives. A spun-down disk might use almost no energy; nevertheless its transition to active mode incurs a latency penalty 100 times higher than a regular access. And it is not all; spinning up the disk platters adds a large energy penalty that practically limits usage of spin-down modes to situations in which the device is likely to stay idle for several minutes, which should further reinforces the need for effective caching and replication algorithms for CDNs.

Network topology: CDN data centers are built in multiple locations to reduce access time and communication costs associated with the service offered. Building a Data center is rather expensive business proposition, for that reason number of the data centers in a given CDN constellation needs to be minimized for desired performance objectives in order to keep the investments under control. Thus the number and geographical locations of CDN nodes have to be selected carefully and usually placed at major network interconnection points and relatively close to the end users. Apart from reducing access delays, content proximity benefits energy efficiency as well since it keeps at minimum involvement of the networking gear and saves on energy needed for powering and cooling of such equipment.

The hybrid CDN topology described earlier is well placed to exploit locality and reduce the redundancy requirements, effectively putting less burden on the resources needed to power their operation.

CDN chunking and replication algorithms: For the purpose of effective delivery over CDNs video content is usually split into smaller chunks called data windows to make it more manageable. Length of data windows can vary from solution to solution and it determines the number of windows constituting a movie. It is then this number of chunks that needs, at minimum, to be distributed across available CDN nodes/servers under the constraint imposed by memory resources at each gateway. Content chunking when combined with data encoding and replica distributions across an entire CDN network instead of concentrating them within only a small number of machines provides

additional resilience of the system. The entire video can be reconstructed from an incomplete set of chunks retrieved. The consequence of such otherwise sound designs is that low traffic levels translate into lower activity for all machines instead of full idleness of a significant subset of them.

5 Conclusion

Interactive Television over the Internet (IPTV) is happening at the intersection of television, IT and telephony convergence. Most interactive television systems nowadays are still video-on-demand oriented and the interaction is merely between the audiences and the content servers. Such services are used to deliver programs to the users on request and find many applications in education, entertainment and business. It is in this context that CDNs have emerged as a new technology to overcome the problems arising on the Internet due to the fast growth of the web related traffic, such as slow response times and heavy server loads. CDN is a distributed video management and delivery system integrated within the data network. It is based on a system of servers with massive storage that places copies from the content library closer to the end user so as to maximize the available bandwidth, and consequently reduces the data access times.

Energy optimization of a given CDN is a complex end-to-end problem, requiring an intricate coordination across hardware, operating systems, virtual machines, middleware and applications. Energy efficiency of a VoD solution can be broadly defined as the amount of concurrent video streams served at requested quality divided by the total energy used in the process. This is relatively a broad definition and we are not to make any attempt to define actual measurement metric and reference workloads. Solving this end-to-end problem at scale of CDN will be much more difficult and here we made an attempt to divide this complex problem into smaller constituting problems and look closer into related operational mechanisms that have an implication on the energy efficiency.

As a CDN is an assembly of mutually interconnected data centers the overall efficiency has to factor for efficiency of each CDN geographical node/data center. The aggregate energy efficiency could not be simple multiplication of individual data centers/servers involved but has to introduce corrective factors to account for energy efficiency tradeoffs brought about by specific geographical CDN topology and how effectively the additional resources deployed and energy used by them is using content chunking and

its interplay with predictive caching and replication algorithms to improve content access performance of a given CDN design.

References

- [1] M. Al-Fares, A. Loukissas, and A. Vahdat, A., A scalable commodity data center network architecture, in *SIGCOMM*, 2008.
- [2] L.A. Barroso and U. Holzle, An introduction to the design of ware-house scale machines, in *The DataCenter As Computer*, University of Wisconsin, Madison, 2009.
- [3] T. Bektas, T., *Discrete Location Models for Content Distribution*, Centre de Recherche sur les Transports, Université de Montréal, 2009.
- [4] B.L. Buzbee, The efficiency of parallel processing. *Frontiers of Supercomputing, Los Alamos Science*, 71, 1983.
- [5] E. Carrera, E. Pinheiro, and R. Bianchini, Conserving disk energy in network servers, in *Proceedings of the 17th International Conference on SuperComputing*, June 2003.
- [6] J. Dean, Designs, lessons and advice from building large distributed systems, in *LADIS. Big Sky: The 3rd ACM SIGOPS International Workshop on Large Scale Distributed Systems and Middleware*, 2009.
- [7] A. Duta, K. Dutta, H. Thomas, and D. VanderMeer, D., WorldWideWaita study of Internet scalability and cache-based approaches to alleviate it, *Management Science*, 49(10):1425–1444, October 2003.
- [8] J. Hamilton, Internet scale service infrastructure efficiency, in *Proceedings ISCA2009*, Austin, Texas, June 2009.
- [9] J. Hass and T. Pierce, T., Data center design guide, *2nd Annual Green Grid Technical Forum*, 2009.
- [10] V.K. Li, W. Lao, X. Qiu, and E.W. Wong, Performance model of interactive video-on-demand systems, *IEEE Journal on Selected Areas in Communication*, 14(6), August 1996.
- [11] W. Paper, *Cisco Visual Networking Index: Forecast and Methodology, 2008–2013*, Cisco, 2008.
- [12] G. Peng, *CDN: Content Distribution Network*, State University of New York at Stony Brook, New York, 2006.
- [13] A. Qureshi, R. Weber, H. Balakrishnan, J. Gutttag, and B. Maggs, Cutting the electric bill for Internet-Scale systems, in *ACM SIGCOMM*, Barcelona, 2009.
- [14] M. Report, *Revolutionizing Data Center Efficiency*, Orlando, April, 2008.
- [15] V. Valancius, N. Laoutaris, L. Massoulie, and C. Diot, Greening the Internet with nano data centers, in *Proceedings of ACM CoNEXT'09*, New York, 2009.

Biographies



Dragan Boscovic envisions, develops and drives the convergence of wireless and wireline edge network solutions across Motorola's businesses. As senior director of Wireless Systems and Networks for Motorola's Applied Research Center, Boscovic leads a global team of researchers that identify opportunities for new products, systems and services. His team leverages wireless technologies, including WiMAX and LTE, as well as IP connectivity to bring service intelligence onto home gateways and other on premise devices in order to blend voice, video and internet user experiences. He also creates the network architecture necessary to deliver services such as mobile and Social TV, media sharing, and adaptive video streaming over fixed and wireless networks.

Currently, Boscovic is focused on leading user experience research and video technology enablers and content delivery solutions that facilitate advanced and converged video, internet and voice services. The migration of service enablers toward home platforms such as home gateways, set-tops, WiFi access points and cellular CPE devices in context of the Wireless Healthcare is another area of particular interest.

Prior to this role, Boscovic served as director of Engineering and Technology for Motorola's Strategic Growth Engine, where he developed technologies that enhanced consumer's abilities to connect with multimedia and one another. He also worked within Motorola's R&D sites in Europe for 12 years. Boscovic is known for his work on Radio Ecology, which describes better use of the limited resources of the radio frequency (RF) spectrum. His vision in this domain led to a number of large collaborative research programs in Europe. He is also credited with the vision of Cognitive Networks which he published in 2005 and has been enhancing ever since.

Boscovic has amassed 18 patents, with eight more pending, and has published numerous papers on topics including wireless technologies and

cognitive networks. He is often called upon to share insights on wireless topics at industry events such as the International Consumer Electronic Show, Global Semiconductor Forum, AlwaysOn, Wireless World Initiative and Wireless World Research Forum.

Boscovic received his bachelor's and master's degrees from the University of Belgrade in Serbia and his doctorate degree from the University of Bath in the UK.



Michael Needham is a Distinguished Member of the Technical Staff in the Applied Research Center, Motorola, Inc., located in Schaumburg, Illinois. He holds Bachelor's and Master's degrees in Electrical Engineering from the University of Illinois at Urbana-Champaign. He has been with Motorola research since 1984, and has worked in several areas, including network and traffic modeling, air interface protocols, and network quality of service technologies. He holds 25 issued US patents, and is an Associate of the Science Advisory Board at Motorola.



Faramak Vakil is a Distinguished Member of Technical Staff in the Motorola Applied Research Center working on network and service management. Prior to joining Motorola in 2006, he was a Senior Scientist in the Applied Research Area of Telcordia. He received the Ph.D. degree in Electrical En-

gineering from Columbia University in the City of New York in 1984, and immediately joined Bellcore (now Telcordia).

Since 1998, Faramak has almost exclusively focused on the architecture, design, control and management of a mobile Internet. He has devised novel mobility, network self-organization and management, and service creation schemes for a mobile Internet. From 1993–1997, he worked on performance management of ATM networks; where he led the design and implementation of an ATM monitoring system, devised a self-assembling network simulator, and conducted experimental studies on the performance of TCP/IP over ATM in hybrid terrestrial-satellite networks. From 1987–1993, he exclusively worked on the specifications of ATM protocols where he developed the error, flow and congestion control; as well as circuit emulation and service clock recovery mechanisms for ATM. In his first assignment, from 1984–1986, he devised a performance monitoring system for the Bellcore MAN trials in Boston and New York. Faramak is a member of ACM and IEEE. He has published numerous papers and standard contributions, and has several patents.



Jin Yang is the Director of Broadband Communication and Application Research Centre, Motorola, based in Beijing, China. He leads a research team working on innovative video and mobility solutions. Prior to this role, Jin was working for Motorola wireless network product research and development in UK.

Dr. Yang began his career as a lecturer at Tsinghua University, where he obtained his MSC and BSC degrees. Later he got his PhD degree from Imperial College, University of London in the field of parallel and distributed computing. He then became a Research Fellow in UCL (University College

London). He then joined Lucent Technologies, UK, in the advanced wireless technology group, before joining the Motorola in UK.

His published works include more than 30 papers, book chapters and standard contributions. Jin Yang holds several patents. Jin is now a member of Motorola SABA (Science Advisory Board Associate).