Content Delivery Networks: A Bridge between Emerging Applications and Future IP Networks

Hao Yin and Xuening Liu, Tsinghua University
Geyong Min, University of Bradford
Chuang Lin, Tsinghua University

Abstract

With the rapid development of network applications, the Internet has evolved from a content-based communication infrastructure to a social-based community network. The emerging applications require the Internet to preserve not only the existing advantages of simplicity and scalability, but also demand varying amounts of capability, availability, reliability, flexibility, and differentiated quality of service. Therefore, it is of paramount importance to bridge the gap between these applications and the IP networks which were originally designed and developed for supporting one-size-fits-all functionality. An efficient solution is to build a virtual network on top of a generic IP transport layer in order to provide additional functionality and flexibility. The content delivery networks technique is one of the successful virtual networks rapidly developed over the last decade with the specific advantage of optimizing the Internet. Nowadays, the CDN has become one of the most important parts of the Internet architecture for content distribution. In this article we highlight the innovative technologies in CDNs and present their evolution triggered by ever newer emerging applications. By presenting in-depth discussion about the architecture, challenges, and applications of CDNs, we demonstrate their importance for the future Internet.

The Internet has become an indispensable part of our daily life as a major platform for information communications. Internet applications have experienced great changes over the past decade. The trend is shifting toward applications that facilitate interactive information sharing, interoperability, user-centered design, and collaboration, such as web-based communities, online hosting services, video sharing, and social networking sites (SNSs). The mainstream vision of 21st century computing and communications is that users will access interactive Internet services and *resource-hungry* applications (network gaming, media streaming, etc.) over lightweight handheld devices, intelligent appliances, and any other devices that can connect to the Internet through heterogeneous access networks, rather than over traditional desktop PCs only.

The status and trends of emerging applications in the future Internet bring significant challenges to the network infrastructure. The importance of these applications involves:

Higher scalability: The number of the Internet users has reached 1.73 billion (as of September 30, 2009), and the penetration continues to increase rapidly. Meanwhile, various electronic devices have been simultaneously applied to access the Internet. Thus, the Internet needs to process and distribute content and applications in a highly cost-effective and efficient manner.

Higher capability: Users prefer putting most of their content on the Internet, and the content is heading toward rich media: large files, audio and video streams, high definition TV (HDTV), and so on. Thus, the Internet needs to offer higher capability in processing, storage, and communication.

Higher quality of service (QoS): The multimedia streaming and emerging SNS applications require higher QoS, such as lower startup delay, lower source-to-end delay, and higher continuity. Any degradation in any of these factors may cause users to turn away to other competing platforms to pursue the benefits.

Stronger interactivity: The Internet has become a community network, where users are not limited to viewing information produced and provided by a few professional content creators, but also want to interact with other users or update web contents. Thus, the content delivery path is not only from a central server through backbone networks to end users, but also from end users to end users. Thus, the Internet needs to offer stronger interactivity.

Heterogeneity: In coming years users will prefer to connect to the Internet through various terminals and via heterogeneous access networks. Thus, the Internet needs to offer transparent services across heterogeneous networks and devices.

Security: Different applications have set various security demands to the Internet service, such as content confidentiality, integrity, and service availability against attacks.

The underlying packet-switched IP networks were designed on the end-to-end principle [1]. This design and service model provides incredibly robust and scalable functions that sacrifice speed and efficiency for ruggedness. In the face of growing traffic volumes and constantly emerging applications, it is difficult for the current Internet to satisfy their demands. To

0890-8044/10/$25.00 © 2010 IEEE

bridge the gap between the applications and best effort IP networks, the technique of content delivery networks (CDNs) has emerged and been developed over the last decade. In Internet Engineering Task Force (IETF) RFC 3466 CDN, which builds a virtual network on top of a generic IP networks, was defined as a type of content network. Such overlay networks provide the flexibility and control necessary to customize and evolve with the carried content. They also have the advantage of easy extension across any network platform, allowing them to evolve with the underlying technology rather than being embedded, and thus offering the desirable flexibility.

CDNs typically deploy servers in multiple geographically diverse locations in order to redirect user requests to the nearest available server. Thus, end users observe higher QoS, and content providers offer more reliable and larger volumes of service. At the same time, Internet service providers (ISPs) can also benefit from deploying CDN servers in their networks as the total amount of traffic transmitted in the backbone is reduced. Initial CDNs addressed static content, later adding audio and video streaming, video sharing, dynamic contents, and applications. To match the different requirements of rapidly and constantly emerging applications, CDNs experienced many active innovations of architecture and techniques during the last decade. These innovations led CDNs to experience successful development, and they now play an important role in the Internet. For example, Akamai Technologies, the leading commercial CDN, handles 20 percent of total Internet traffic today.

The rest of this article is organized as follows. The next section describes the existing problems in the current IP networks and introduces the idea of building *content networks* in layers 4 through 7 in the open systems interconnection (OSI) stack to bridge the gap between new applications and IP networks. We then introduce CDNs and the related technical taxonomy. The following section brings an in-depth discussion of the evolution of significant technologies of CDNs triggered by emerging applications, through which we can declare that CDNs are suitable to leverage the shortages of best effort IP networks. The emerging challenges of CDNs are highlighted in the following section. This article is concluded in the final section.

Potential Solutions to Problems in IP Networks

In this section we first analyze the existing problems in the current IP networks and then discuss the potential solutions.

Existing Problems in IP Networks

The Internet is one of the most successful technology achievements ever. In its 40-year history, the Internet has grown from a small experimental platform to a global infrastructure that connects billions of people. The design principles of the current Internet can be characterized as follows [2]: layering, packet switching, a network of collaborating networks, and intelligent end systems as well as the end-to-end argument.

The critical issue facing the Internet is to support differentiated QoS for heterogeneous applications in a flexible and inexpensive way. However, the current IP networks were originally designed for providing one-size-fits-all functionality, that is, for *best effort* applications not covered by QoS guarantees and other control mechanisms. However, the emerging applications (e.g., online gaming, SNSs, videoconferencing, live streaming, and video sharing) require varying amounts of reliability, functionality, speed, efficiency, cost effectiveness, and scalability. Therefore, it is necessary to bridge the gap between the emerging heterogeneous applications with various demands of service capacity and IP networks.

Incremental and Clean Slate

Despite its simplicity and scalability, the IP network lacks QoS differentiation, traffic control, and management mechanisms. This fact brings significant challenges to Internet service. In recent years many academic and industrial communities have put in hard work rethinking how to improve the current Internet architecture to accommodate the emerging applications.

There are two alternative design principles to develop the future Internet:

- **Incremental design**: A system is moved from one state to another with incremental patches. To date, network researchers have focused on solutions that incrementally improve the Internet with the implicit assumption that radical new solutions are not needed or have no chance of being deployed.
- **Clean-slate design**: To cater to future needs, the Internet has to be extended. It has to be redesigned for the present requirements, at the same time ensuring enough flexibility to adequately incorporate future requirements.

Due to the industrial and economic reasons, it is impossible to enable more than 13,000 competing ISPs to abandon the current Internet infrastructure and establish a brand new one. Thus, the incremental design principle will still inhabit the mainstream in coming decades.

Performance Enhancement of IP Networks

The OSI stack of the Internet can be characterized in two logical layers:

- Infrastructure layer, which traditionally processes data at layers 1 through 3, centered on the routing, forwarding, and switching of frames and packets
- Service layer, which includes layers 4 through 7 and deals with the routing and forwarding of requests and responses for content

The latter is also defined as content networks in RFC 3466.

To preserve stateless best effort IP networks requires putting more state-relative functions in the upper (service) layer to solve the flexible QoS control problem over the global Internet. The way the IP layer provides only simple and general service can reduce cost, facilitate future network upgrades, and enable new applications to be added without the need for changes to the existing network. OpenFlow, a famous research activity from the Clean-Slate Program, also agreed that the IP layer should be as thin as possible, handling only simple forward functionality.

An efficient way is to build a virtual network on top of a generic IP transport layer in order to add on additional functionality, security, and flexibility to IP networks. Such overlay networks provide flexibility, traffic control, and resource management with the advantages of easy extension across a heterogeneous network platform without significant changes to the underlying technology.

CDNs and Technical Taxonomy

Overview of CDNs

CDNs provide enormous utility in a cost-effective way as an overlay network is built on top of generic IP networks to improve overall system performance and offer fast and reliable applications and services. Flexibility is one of the most important advantages of CDNs.

A CDN is a system of computers containing copies of data placed at various locations in order to maximize bandwidth for access to the data from clients throughout the network. A client accesses a copy of the data near the client. A CDN usually has the following common components:

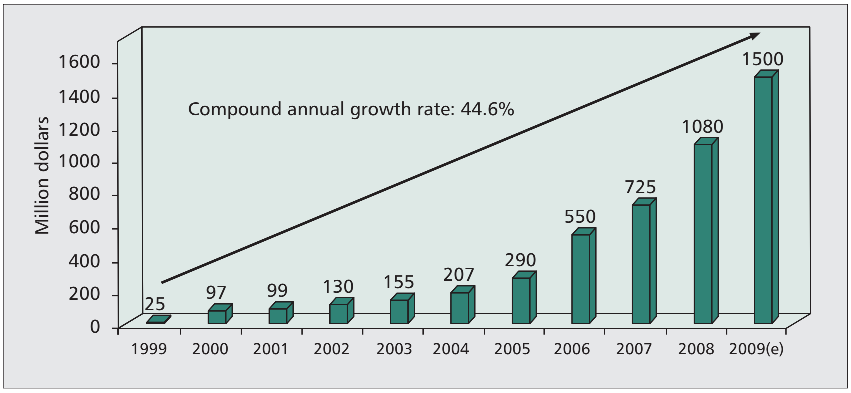

Figure 1. *Global CDN market [4].*

- A content outsourcing unit consists of mechanisms that move contents from the source server to the surrogates.
- A content delivery unit consists of a set of surrogate servers that deliver copies of content to users.
- A request routing unit consists of mechanisms that move a client request toward a rendezvous with a surrogate.
- A management unit tracks and collects data on request routing, outsourcing, and delivery functions within the CDN.

In CDNs the traditional communication flow between clients and the original server is devided into two communication flows: one between the client and the surrogate server, and another between the surrogate server and the original server. CDNs can bring benefits to end users, content providers, and ISPs. End users observe higher QoS, while content providers can offer more reliable service and handle larger numbers of client requests. Meanwhile, ISPs can benefit from deploying CDN servers in their networks as the total amount of upstream and transit traffic is reduced.

The first CDN appeared in 1998. Akamai Technologies, one of the earliest successful companies in CDNs, aimed to solve the flash crowd problem and developed a set of breakthrough algorithms for intelligently routing and replicating content over a large network of distributed servers around the world in a reliable and scalable manner [3]. Over the last decade, there have been many commercial CDNs established to provide the distributed computing platform for delivering Internet contents and applications, including Akamai Technologies, Limelight Networks, Digital Island, CDNetworks, Verisign, BitTorrent, ChinaCache, and so on. Moreover, a number of non-commercial (academic) CDNs, such as CoDeeN, Coral, and Globule, which mainly use decentralized peer-to-peer (P2P) technologies [3], have also appeared.

Currently, CDNs have become a huge market generating large revenues. The global CDN market was as high as $1.5 billion in 2009, a 60-fold growth over the past 10 years. As shown in Fig. 1, the market for CDNs has increased at a higher speed since 2005 because video streaming applications became the most popular and penetrated worldwide.

A Technical Taxonomy of CDNs

This subsection highlights important technologies with sufficient coverage of CDNs. The technologies can be classified into two categories: infrastructure and service (Fig. 2).

The infrastructure part involves three separate but relevant technologies, placement, data center, and networking, which work together to find out a good trade-off between QoS and cost. The questions to be addressed include:

- How many data centers should the system have?
- Where should the servers be allocated?
- How many servers should be deployed and organized in a given data center?
- How should the relationship between the data centers and ISPs be established?

The service part is composed of three main technologies. *Request routing* is responsible for routing client requests to an appropriate server for delivery of contents. The request routing system uses a set of metrics such as network proximity, client perceived latency, distance, and replica server load in an attempt to direct users to the most suitable server to best serve the request.

Content distribution and management includes content outsourcing, content delivery, and content management technologies; the latter is largely dependent on the techniques for cache organization (i.e., caching techniques, cache maintenance, and cache update).

System management includes operational support systems (OSSs) and business support systems (BSSs). An OSS mostly deals with supporting processes such as maintaining inventory, providing services, configuring components, and managing faults. A BSS typically deals with customers, supporting processes such as taking orders, processing bills, and collecting payments.

Evolution of CDN Technologies

Internet applications have been undergoing tremendous changes with higher requirements for quality and performance, and this trend is continuing. We discuss the evolution of CDN technologies according to significant changes in applications and types of content/service delivered by CDNs. These technologies can be classified into four categories: highly distributed CDNs, big data center CDNs, P2P-assisted CDNs, and cloud CDNs. The first two have been widely deployed by many commercial CDNs. The third attempts to solve the cost problem of servers for delivering rich media, while the last one is an emerging solution for leveraging economies of scale.

Highly Distributed CDNs

When the CDN technique appeared 10 years ago, the major content types over the Internet were small-sized text files and images. To achieve efficient delivery of these contents, CDN providers (e.g., Akamai Technologies) usually adopt the highly distributed CDN approach to build their platforms. This type of CDN can be described as *entering deep into ISPs* [5], deploying content distribution servers inside ISP points of presence (POPs). The idea is to allocate the expected content closer to end users to improve user-perceived performance in terms of delay and throughput. This design results in a large number of server clusters scattered around the globe. Such an approach is an ideal solution to handle the last mile problem.

The common request routing mechanism used in most commercial CDNs is Domain Name System (DNS)-based request routing. Its advantage lies in maintaining lower overhead in the original server than URL rewriting and having higher security as it conceals the source to end users. Because of this highly distributed design, the task of network maintenance and management becomes very challenging. Sophisticated algorithms are required to shuffle data among the servers across the public Internet. Moreover, edge side includes (ESI) technology has been adopted to achieve dynamic web page acceleration where the dynamic web page is separated into two parts, static and dynamic. The static part is cached in an

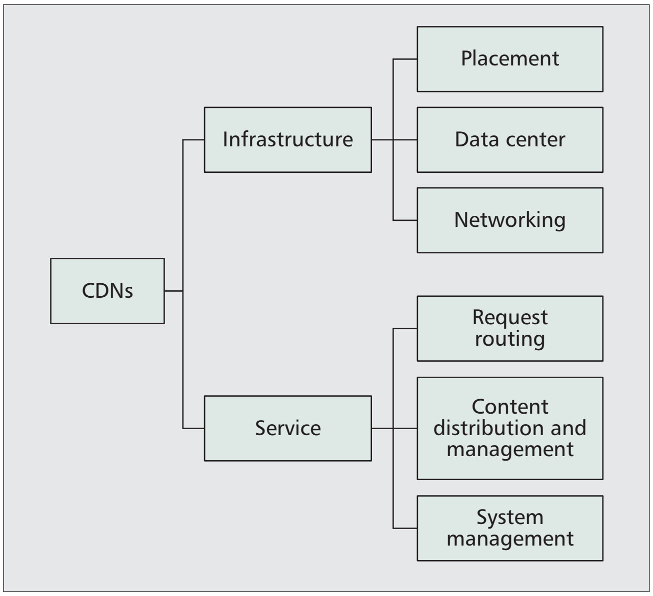

Figure 2. *Technical taxonomy of CDNs.*

edge server, and the CDN edge server retrieves the dynamic part from the original server through an optimal path in order to reduce total response time.

Big Data Center CDNs

With the growth of new applications and higher access speeds, Internet users do not feel satisfied with only small-size simple contents. As a result, multimedia content (e.g., audio and video) delivery through the Internet has received much attention. To achieve efficient delivery of media contents, the big data center approach has been adopted to build the platform of some CDNs (e.g., Limelight Networks). This approach can be described as *bring ISPs to the home* [5] by building large centers of content distribution at only a few key locations and connecting these centers using private high-speed connections. Instead of getting inside ISPs' POPs, these CDNs typically place each distribution center at a location that is simultaneously near the POPs of many large ISPs. Such an approach is considered a suitable solution to the *middle-mile* problem.

Compared to the aforementioned highly distributed CDNs, this design approach typically results in lower overhead for maintenance and management, possibly at the expense of higher delay to end users. Such CDNs usually adopt IP Anycast technology as their request routing mechanism.

P2P-Assisted CDNs

P2P is an alternative (or competing) technology for delivering Internet contents. It has recently become very popular to cope with the growing demand of end users and has been used in applications such as file sharing, video streaming, and distributed computing. P2P systems solve the scalability issue by leveraging the resources of the participating peers. P2P systems achieve high scalability while reducing the server requirements. However, the decentralized and uncoordinated operation implies that this appealing scalability comes with undesirable side-effects including network unfriendliness.

A natural question in P2P-assisted CDNs is whether P2P and CDNs can be integrated to obtain the advantages of scalability inherent in P2P, and the reliability and manageability of CDNs. Indeed, with the recent rapid growth of P2P applications and CDNs, many initiatives and efforts have been made to combine these advantages to get the best of both worlds. Akamai, with its acquisition of Red Swoosh P2P technology [5], is expected to combine P2P file distribution software with its back-end control system and global network of edge servers. VeriSign, CacheLogic, Grid Networks, Internap, Joost, and ChinaCache have all announced their own P2P-assisted CDN services as well.

Such an approach is the appropriate complement to the highly distributed and big data center CDNs. It can be established on any of the previous two approaches in order to efficiently reduce the server cost even when delivering large files or live video streaming. A P2P-assisted CDN live streaming system, LiveSky, was introduced in [6]. This system can efficiently resolve the cost and scalability issues by involving P2P technologies in existing CDN architecture without the need for noticeable changes to the existing infrastructure. Moreover, such a system can also efficiently overcome the problems of traditional P2P streaming systems, including low quality, network unfriendliness, and upload unfairness, by adequately adopting the existing CDN infrastructure when establishing hybrid approaches. On the other hand, the P2P-assisted approach can offer additional QoS guarantee for real-time HDTV applications. Such applications require very high bandwidth when downloading streaming content. When the path condition between client and the surrogate server does not satisfy the performance requirements of the applications, the P2P-assisted approach can provide complementing capacities to increase the available downloading bandwidth by leveraging the resources of the participant clients, thus enhancing the system QoS.

Cloud CDNs

Cloud computing is a new term for a long-held dream of computing as a utility and has recently emerged as a commercial reality. The construction and operation of extremely large-scale data centers are the key part of cloud computing. Cloud computing will enable new types of applications; meanwhile, important existing applications, such as software as a service (SaaS) and Facebook apps, will become even more compelling [7]. Cloud computing, as a new IT trend, opens innovative perspectives in the architecture, design, and implementation of CDNs. One direction is that CDN providers may build their own clouds, and another is to use existing clouds to assist content distribution.

Content delivery in cloud CDNs could be achieved by leveraging existing cloud providers. Cloud CDNs can intelligently match and place contents on one or many storage clouds based on their QoS, coverage, and budget preferences. The important value of cloud CDNs is economies of scale, and they enable customers to distribute web contents in a pay-as-you-go manner. The initial cases of cloud CDNs are Amazon CloudFront and MetaCDN. Cloud CDNs benefit not only customers but also CDN providers as complements since their service capacity can be enhanced in a pay-for-use manner in the presence of a flash crowd, thus avoiding large investments in infrastructure.

Emerging Challenges in CDNs

In this section we highlight the important challenges in CDNs.

Server Placement

How many data centers are enough? This problem has made it even more important to decide whether we should adopt highly distributed CDNs or big data center CDNs. The general view is that the larger the number of data centers, the better the user experience, but the higher the cost. However, a recent study [8] indicated that the number of nodes in Akamai can be reduced to 60 without noticeable degradation of system performance. On the other hand, Leighton [9] has noted that highly distributed CDNs can also reduce cost by deploying servers in some regions that can host their servers for free.

Data Center Structure and Server Organization

It is inefficient for each server in a CDN to serve only one customer or application type due to load imbalance and cost ineffectiveness. Thus, it is important to achieve efficient task migration among the servers in a data center. Several related techniques including virtualization will be involved in the design of a data center structure. Moreover, it is a great challenge to organize the servers in a data center in order to reduce the burden of configuration and accelerate deployment speed.

Content Distribution

The recent challenge of content distribution in CDNs focuses on the acceleration of dynamic contents and computing applications, which have been considered uncacheable. The general solution is to optimize the path between the source and edge servers [9]. One feasible way is to allocate them in the same data center, but this approach cannot offer desirable performance for end users worldwide. Thus, the significant challenge is to find potential solutions to cache or partly cache dynamic contents or applications so that edge servers can be widely distributed to enhance the user experience.

Request Routing

The current mechanism of request routing is based on DNS infrastructure, which is not suitable in some environments. A critical challenge is how to combine or find other mechanisms to achieve efficient request routing. For example, in IP multimedia subsystem (IMS), the elements are identified by TEL/SIP URI. Thus, CDN request routing needs an integrated solution to accommodate heterogeneous systems. Moreover, several network applications (e.g., online gaming) are latency-sensitive, and need more efficient and accurate routing strategies. A recent study [10] introduced a latency prediction system, Htrae, which combines both geolocation and network coordinate systems to achieve the benefits of both, thus efficiently reducing the prediction error.

System Management

It is a significant challenge for the service provider to manage a large-scale distributed system. The traditional mechanisms typically focus on the management of servers. However, CDN servers are deployed in multiple ISPs, so it is essential to manage multiple resources including servers and networks. Moreover, efficient system management can reduce investment in CDNs and becomes a significant challenge. Krishnan *et al*. [11] introduced a tool, WhyHigh, to identify, diagnose, and prioritize the network problems of CDNs. By virtue of WhyHigh, the network administrators of CDNs (e.g., Google) can efficiently use and configure existing nodes to improve system performance without adding new nodes. An efficient tool, Entact, was introduced in [12] to dynamically redirect client requests and adaptively achieve a route injection mechanism to efficiently change the content distribution path, thus reducing total traffic cost without compromising performance.

Conclusions

The fundamental design principle of the current Internet architecture is best-effort-based packet switching. Due to its advantages of simplicity and efficiency, the Internet has grown quickly. Simple best effort traffic has been shown to work well in the Internet over the past decades. The future Internet architecture should preserve the features of best effort while offering differentiated QoS and various amounts of accessibility, reliability, feasibility, and security.

The CDN technique was developed to offer the service of large-scale content delivery based on IP networks and improve system performance in order to satisfy the QoS requirements of heterogeneous Internet applications. The rapid development over the past decade indicates that CDNs can efficiently satisfy the demands of ever emerging applications by adopting innovative architecture and technologies. Thus, CDNs have become a significant bridge between various emerging applications and best effort IP infrastructure thanks to their flexibility. CDNs will grow to enable future IP networks to offer the requirements of emerging Internet applications and users.

Acknowledgment

This work was funded by Project 60873254 supported by NSFC, Project 20090460317 supported by the China Postdoctoral Science Foundation, and the Tsinghua-ChinaCache CDN Research Institute Project.

References

[1] M. S. Blumenthal and D. D. Clark, "Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New *World," ACM Trans. Internet Tech.*, vol. 1, no. 1, Aug. 2001, pp. 70–109.

[2] A. Feldmann, "Internet Clean-Slate Design: What and Why?" *ACM SIGCOMM Comp. Commun. Rev.*, vol. 37, no. 3, July 2007, pp. 59–64.

[3] R. Buyya, M. Pathan, and A. Vakali, Content Delivery Networks, Springer, 2008.

[4] China Mobile, "The Evolution Status and Trend of CDNs," (in Chinese); http://labs. chinamobile.com/report/view_15707

[5] C. Huang *et al.*, "Understanding Hybrid CDNP2P: Why Limelight Needs its Own Red Swoosh," *Proc. NOSSDAV*, May 2008, pp. 75–80.

[6] H. Yin *et al.*, "Design and Deployment of a Hybrid CDN-P2P System for Live Video Streaming: Experiences with LiveSky," *Proc. ACM Multimedia*, Oct. 2009, pp. 25–34.

[7] M. Armbrust *et al.*, "Above the Clouds: A Berkeley View of Cloud Computing," Tech. Rep. UCB/EECS-2009-28, Feb. 2009, p. 25.

[8] S. Triukose, Z. Wen, and M. Rabinovich, "Content Delivery Networks: How Big is Big Enough?" *Proc. ACM SIGMETRICS*, June 2008, pp. 59–60.

[9] T. Leighton, "Improving Performance on the Internet," *Commun. ACM*, vol. 52, no. 2, Feb. 2009, pp. 44–51.

[10] S. Agarwal and J. R. Lorch, "Matchmaking for Online Games and Other Latency-Sensitive P2P Systems," *Proc. ACM SIGCOMM*, Aug. 2009, pp. 315–26.

[11] R. Krishnan *et al.*, "Moving Beyond End-to-End Path Information to Optimize CDN Performance," *Proc. IMC*, Nov. 2009, pp. 190–201.

[12] Z. Zhang *et al.*, "Optimizing Cost and Performance in Online Service Provider Networks," *Proc. NSDI*, Apr. 2010, p. 15.

Biographies

HAO YIN (h-yin@tsinghua.edu.cn) is an associate professor in the Department of Computer Science and Technology, Tsinghua University, Beijing, China. He is serving as chief scientist of Beijing Blue I.T. Technologies Co., Ltd. (ChinaCache), and director of the Tsinghua-ChinaCache Content Delivery Network Research Institute (http://www.cdnlab.org/). His research interests are content delivery networks, performance evaluation for Internet and wireless networks, multimedia communications, and network security. He has published more than 70 papers in refereed journals and conferences, and holds several patents in the previous mentioned research areas.

XUENING LIU (xnliu@tsinghua.edu.cn) is a research associate in the Department of Computer Science and Technology, Tsinghua University. He received his Ph.D. degree in computer science and technology from Tsinghua University in 2009. His research interests include multimedia communications, content distribution networks, Internet measurement, and performance evaluation.

GEYONG MIN (g.min@brad.ac.uk) is a senior lecturer in the Department of Computing at the University of Bradford, United Kingdom. He received his Ph.D. degree in computing science from the University of Glasgow, United Kingdom, in 2003, and his B.S. degree in computer science from Huazhong University of Science and Technology, China, in 1995. His research interests include performance modeling/evaluation, mobile computing and wireless networks, multimedia systems, grid computing, and parallel and distributed systems.

CHUANG LIN (chlin@tsinghua.edu.cn) is a professor with the Department of Computer Science and Technology, Tsinghua University. His current research interests include computer networks, performance evaluation, network security analysis, and Petri net theory and its applications. He has published more than 300 papers in research journals and IEEE conference proceedings in these areas and has published three books. He is a member of the ACM Council and the Chinese Delegate in TC6 of IFIP.