Computer Networks 55 (2011) 3987–3990

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Editorial

Editorial for special issue Internet-based Content Delivery

We are in the midst of an Internet computing revolution. One vision of 21st century computing is that users will access Internet services and "resource-hungry" applications (e.g. gaming, streaming media, video on demand, and voice-over-IP) over lightweight portable devices rather than through some descendant of the traditional desktop PCs. In this context, distributing and processing Internet-based data in an efficient and cost-effective manner is a challenging issue in Internet technology.

Content Delivery Networks (CDNs) have emerged to overcome the inherent limitations of the Internet in terms of user perceived Quality of Service (QoS) when accessing Web data. They offer infrastructure and mechanisms to deliver content and services in a scalable manner, and enhance users' Web experience. Specifically, a CDN is an overlay network across the Internet, which consists of a set of servers (distributed around the world), routers and network elements. Edge servers are the key elements in a CDN, acting as proxy caches that serve directly cached content to users. With CDNs, content is distributed to edge cache servers located close to users, resulting in fast, reliable applications and Web services for the users. Once a user requests content on a Web provider (managed by a CDN), the user's request is directed to the appropriate CDN server. The perceived high end-user performance and cost savings of using CDNs have already urged many Web entrepreneurs to make contracts with CDNs. For instance, Akamai – one of the largest CDN providers in the world – claims to be delivering 20% of the world's Web traffic.[1] While the real numbers are debatable, it is clear that CDNs play a crucial role in the modern Internet infrastructure.

The main value proposition for traditional CDN services has shifted over time. Initially, the focus was on improving end-user perceived experience by decreasing response time, especially when the customer Web site experiences unexpected traffic surges. Nowadays, CDN services are treated by content providers as a way to use a shared infrastructure to handle their peak capacity requirements, thus allowing reduced investment cost in their own Web site infrastructure. Moreover, recent trends in CDNs indicate a large paradigm shift towards a utility computing model, which allows customers to exploit advanced content delivery services without having to build a dedicated infrastructure. These trends foster the necessity and success of a well-designed content-utility system to provide highly scalable Web content delivery over the Internet. One approach to address these issues is to exploit the recent emergence of "Cloud Computing". Cloud Computing is a recent trend in Information Technology (IT) that moves computing and data away from desktop and portable PCs into computational resources such as large Data Centers ("Computing") and make them accessible as scalable, on-demand services over a network (the "Cloud"). The main technical underpinnings of Cloud Computing infrastructures and services include virtualization, service-orientation, elasticity, multi-tenancy, power efficiency, and economics of scale. The perceived advantages for Cloud-service clients include the ability to add more capacity at peak demand, reduce cost, experiment with new services, and to remove unneeded capacity. The use of Clouds for content delivery is highly appealing as Cloud providers (e.g. Amazon S3, Amazon CloudFront, Mosso Cloud Files, and Nirvanix Storage Delivery Networks) charge customers for their utilization of storage and transfer of content (*pay-as-you-go*), typically in the order of cents per gigabyte. They also offer SLA-backed performance and uptime guarantees for their services.

The integration of Cloud Computing in content delivery opens new perspectives in Internet technologies, raising new issues in the architecture, design and implementation of existing CDNs. Moreover, the evolution of next-generation Internet-based Content Networks (CNs) in a large-scale heterogeneous environment demands adaptation within the research community in terms of the technologies used. Therefore, the integrated uses of existing content delivery technologies and emerging technologies (i.e. Cloud Computing, Peer-to-Peer (P2P), utility computing, and agent technology) are anticipated to augment the effectiveness and boost the efficiency of future CN infrastructures.

Given the continued, intense activity in this area, we invited researchers and practitioners to submit papers to this special issue of Computer Networks describing research efforts and experiences in the content networking domain.

[1] http://www.akamai.com/html/perspectives/index.html.

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2011.09.009

This special issue has a three-fold integrated contribution: to deliver the state-of-the art in the current research about Internet-based content management and delivery; to promote the content networking discipline by integrating different perspectives and by innovating through sound propositions; and to challenge the future research by indicating open research issues and addressing key problems whose solution will enhance the evolution of next-generation Cloud-based CNs. From among the 33 submissions, and after rigorous review, we selected the following 10 papers as representative of ongoing research and development activities.

CDNs are striking a balance between the bandwidth costs for Web content providers and the quality of service for Web customers. In the paper "Content Delivery and Caching from a Network Provider's Perspective" Haßlinger and Hartleb [1] present a survey of the current Internet-based content delivery schemes, providing a comparison between CDN and P2P overlays. The impact of caching and traffic engineering on the performance of CDN and P2P infrastructures, based on delay measurement, is also discussed. On a similar topic, the paper "Algorithms for Optimizing the Bandwidth Cost of Content Delivery" by Adler et al. [2] deals with the bandwidth cost optimization problem for CDNs and multi-homed enterprises. Authors study from a theoretical point of view both offline and online algorithms using realistic contract models.

As the volume of data produced at the Internet surges, the pub/sub paradigm becomes increasingly popular for content access and dissemination. In the paper "Storage Planning and Replica Assignment in Content-Centric Publish/Subscribe Networks" Sourlas et al. [3] enhance the pub/sub communication paradigm with an advertisement and a request/response mechanism. The authors focus on content-centric pub/sub networks and describe a new placement and replica assignment algorithm.

P2P technology has a major impact on Internet streaming (video-on-demand, live streaming, etc.) and other content distribution services. In the paper "On Tradeoffs Between Cross-ISP P2P Traffic and P2P Streaming Performance" Yang and Xu [4] propose an upload bandwidth allocation mechanism to achieve tradeoffs between the cross-ISP P2P traffic and the P2P streaming performance for P2P assisted video-on-demand streaming. Another challenging issue for the performance of P2P content distribution systems is to investigate how peers use distributed algorithms to simultaneously use multiple neighbors to help them achieve their downloading needs and minimize the content server's support. This problem is known as the "load balancing problem" and it is addressed in the paper "Design and Evaluation of Load Balancing Algorithms in P2P Streaming Protocols" by Wang et al. [5]. In this work, the authors analyze this problem from several aspects and study a number of practical algorithms.

As more facets of work and life move online and the Internet expands beyond a communications network to become a platform for business and society, we observe an increase in popularity of live streaming media. In the paper "Live Streaming of User Generated Videos: Workload Characterization and Content Delivery Architectures", Silva et al. [6] characterize the workload of a real Web 2.0 live streaming media service, providing distribution models for different aspects related to live transmissions and user viewing sessions. On a similar topic, the paper "Improving Multi-View Peer-to-Peer Live Streaming Systems with the Divide-And-Conquer Strategy" by Wang et al. [7] proposes a flexible, efficient and scalable protocol for multi-view P2P streaming systems using a divide-and-conquer strategy. Extensive simulations are used to evaluate the proposed protocol.

Video on Demand, IPTV and other large-scale content distribution services have emerged in our lives. In this context, efficient content delivery over the Internet has become an important element for improving Web performance. In the paper "On Predictable Large-Scale Data Delivery in Prefix-based Virtualized Content Networks" Wählisch et al. [8] present a novel overlay content delivery approach that distributes data on prefix-based overlay trees.

Data traffic in mobile cellular networks is growing rapidly. In the paper "Optimised Local Caching in Cellular Mobile Networks", Olof Arvidsson et al. [9] deal with the problem of increasing backhaul transmission costs in cellular mobile networks. The authors propose a novel distributed caching scheme where caches are organized in domains so as to fetch content from other caches and share content between caches in the domain. The characteristics of such a scheme are investigated by applying an analytical model to a set of realistic examples.

The advent of Cloud computing has resulted in a transition from "traditional desktops" that have dedicated hardware and software installations into "virtual desktop clouds" that are accessible via thin-clients. In the last paper presented in this special issue, "Utility-Directed Resource Allocation in Virtual Desktop Clouds" Calyam et al. [10] study the impact of increasingly constrained memory and network health conditions on the performance of various application tasks in a virtual desktop cloud environment. The authors propose a conceptual framework as well as results from an actual implementation to incorporate offline benchmarking and guide online resource allocation within virtual desktop clouds.

This special issue came to be due to the direct and indirect involvement of many researchers, academicians, developers, designers, and industry practitioners. We would like to express our gratitude to the authors of accepted papers and the reviewers for their contributions to this special issue. Throughout the working period of the special issue, its organization, review, and selection of papers, we were constantly involved in discussion to fine tune our work. We have directed the authors of accepted papers to provide high quality manuscripts with minimal errors within the texts.

With the contributions from authors, we have tried to make this issue as rich as possible within our capacity. Because of the rigorous review process and selection policy enforced by the expert review panel, we could accept only a limited number of high quality papers. However, we also thank those authors whose contributions could not be selected for this special issue. Besides these, we are very thankful to Dr. Harry Rudin, Editor-in-Chief of Computer Networks journal for accepting our journal special issue proposal, for giving us the opportunity to work on it, and for his continuous support throughout this project from

establishing a successful call for papers to the time-consuming editorial fine-touch he put on for each accepted paper. Last but not least, we also express our gratitude to the production staff of the Computer Networks journal for their professional work to keep the outcome in line with the journal's policies and key dates.

We hope that through this special issue we are delivering a state-of-the-art glimpse of current Internet-based content delivery networking topics, bringing to the attention of the community novel problems that must be investigated. We hope that this special issue will serve as a valuable reference for researchers and practitioners working in the CDN domain and its emerging consumer applications. We envision this special issue as establishing a pathway for the integrated use of existing technologies for the development of future-generation Content Delivery Networks.

References

[1] G. Haßlinger, F. Hartleb, Content delivery and caching from a network provider's perspective, Computer Networks 55 (18) (2011) 3991–4006.

[2] M. Adler, R. Sitaraman, H. Venkataramani, Algorithms for optimizing the bandwidth cost of content delivery, Computer Networks 55 (18) (2011) 4007–4020.

[3] V. Sourlas, P. Flegkas, G.S. Paschos, D. Katsaros, L. Tassiulas, Storage planning and replica assignment in content-centric publish/subscribe networks, Computer Networks 55 (18) (2011) 4021–4032.

[4] P. Yang, L. Xu, On tradeoffs between cross-ISP P2P traffic and P2P streaming performance, Computer Networks 55 (18) (2011) 4033–4042.

[5] Y. Wang, T.Z.J. Fu, D. Ming Chiu, Design and evaluation of load balancing algorithms in P2P streaming protocols, Computer Networks 55 (18) (2011) 4043–4054.

[6] T. Silva, J. Almeida, D. Guedes, Live streaming of user generated videos: workload characterization and content delivery architectures, Computer Networks 55 (18) (2011) 4055–4068.

[7] M. Wang, L. Xu, B. Ramamurthy, Improving multi-view peer-to-peer live streaming systems with the divide-and-conquer strategy, Computer Networks 55 (18) (2011) 4069–4085.

[8] M. Wählisch, T.C. Schmidt, G. Wittenburg, On predictable large-scale data delivery in prefix-based virtualized content networks, Computer Networks 55 (18) (2011) 4086–4100.

[9] Å. Olof Arvidsson, A. Mihály, L. Westberg, Optimised local caching in cellular mobile networks, Computer Networks 55 (18) (2011) 4101–4111.

[10] P. Calyam, R. Patali, A. Berryman, A.M. Lai, R. Ramnath, Utility-directed resource allocation in virtual desktop clouds, Computer Networks 55 (18) (2011) 4112–4130.

Giancarlo Fortino is an Associate Professor of Computer Science at the Department of Electronics, Informatics and Systems (DEIS) of the University of Calabria, Rende (CS), Italy. He received his Laurea degree (bachelor + master) in Computer Science Engineering from University of Calabria in 1995 and a PhD in Computer Science and Systems Engineering from the same institution in 2001. He also received an MSc in advanced technologies of computer science and communications from the IIASS (International Institute for Advanced Scientific Studies), Salerno (Italy), in 1997. In 1997 and 1999 he was visiting researcher at the International Computer Science Institute (ICSI), Berkeley (CA), USA. From 2001 to 2006, he was Assistant Professor of Computer Science at the Department of Electronics, Informatics and Systems (DEIS) of the University of Calabria, Rende (CS), Italy. His current research interests agent-based systems, agent oriented software engineering, content distribution networks, wireless sensor networks, distributed multimedia systems, GRID systems. He is author of more than 150 papers in international journals, books and conferences. He is an editorial board member of JNCA (Elsevier) and is member of IEEE, IEEE Computer and Communications Society, and ACM. For more information, please visit http://si.deis.unical.it/fortino.

Carlo Mastroianni is a researcher at the Institute of High Performance Computing and Networks of the Italian National Research Council (ICAR-CNR) in Cosenza, Italy, since 2002. He received his PhD in Computer Engineering from the University of Calabria in 1999. Previously, He was a Computer Engineer at the Computer Department of the Prime Minister Office, in Rome. His research interests focus on distributed systems and networks, and in particular on Grid and Cloud Computing, peer-to-peer networks, content distribution networks, multi agent systems, bio-inspired algorithms and protocols. He published more than 25 scientific papers in international journals, among which, IEEE/ACM TON, IEEE TEVC, ACM TAAS, IEEE Multimedia, Parallel Computing (Elsevier). He published more than 60 papers in proceedings of international conferences and books. He served as guest editor for special issues of Future Generation Computer Systems, Elsevier (vol. 24, No. 3, 2008), Multiagent and Grid Systems, Ios Press, (vol. 5/2, 2009), and Journal of Network and Computing Applications, Elsevier (vol. 32/5, 2009). He is a member of the Editorial Board of the Journal of Network and Computing Applications, Elsevier and he serves in the program committee of several international conferences. He is member of ACM and IEEE.

George Pallis is a lecturer at the Computer Science Department, University of Cyprus. Previously, he was a visiting lecturer and Marie Curie Fellow in the same department. He received his BSc and PhD degrees from the Department of Informatics at Aristotle University of Thessaloniki, Greece. His research interests focus on computer system technologies in the area of the Internet with emphasis on web data management, web data caching, content-based distribution networks, cloud and grid computing, information retrieval, and web data clustering. He has published several papers in international journals (e.g., CACM, IEEE Internet Computing, WWWJ) and conferences (e.g., IDEAS, ADBIS, CCGrid, ISMIS). He is an editorial board member of IEEE Internet Computing. For more information, please visit http://www.cs.ucy.ac.cy/~gpallis/.

Mukaddim Pathan is a research scientist at the Commonwealth Scientific and Industrial Research Organization (CSIRO), the national government body of scientific research in Australia. He also holds the position of an adjunct lecturer at the Australian National University. Previously, he was a researcher at the Cloud Computing and Distributed Systems (CLOUDS) Lab of the University of Melbourne, Australia. He holds a PhD in Computer Science and Software Engineering from the University of Melbourne. His research interests include data management, resource allocation, load balancing, and coordination policies in wide-area distributed systems such as Content Delivery Networks, Cloud computing, and Sensor networks. He is one of the developers of MetaCDN that leverages the capabilities of existing storage Clouds for high performance content delivery. He is the editor of the book Content Delivery Networks, Lecture Notes in Electrical Engineering, Vol. 9, Springer-Verlag, Germany. He has authored and co-authored a number of research papers in internationally recognized journals and conferences. He is involved in the organization of the UPGRADE-CN and IDCS workshops and is a PC member of several international conferences. He has edited a few research issues in reputed international journals and also serves as the reviewer of a few renowned journals such as IEEE Transactions on Circuits and Systems for Video Technology (TCSVT), International Journal of Management Science (OMEGA), Journal of Network and

Computer Applications (JNCA), Computer Communications, Computer Networks, Journal of Systems and Software, and IEEE Software. He is a member of IEEE, IEEE computer society, and ACM. For more information, please visit: http://www.ict.csiro.au/staff/mukaddim.pathan.

Athena Vakali has been a faculty member (now an associate professor) in the Department of Informatics at Aristotle University of Thessaloniki since 1997. She is head of the Operating Systems Web/Internet Data Storage and management research group OSWINDS (http://oswinds.csd.auth.gr). Her research activities are on various aspects and topics of web information systems, including web data management (clustering techniques), content delivery on the web, web data clustering, web caching, XML-based authorization models, text mining, and multimedia data management. Her publication record is now at more than 100 research publications, which have appeared in several journals (e.g., CACM, IEEE Internet Computing, WWWJ), book chapters, and scientific conferences (e.g., IDEAS, ADBIS, ISCIS, ISMIS). She is a member of the editorial board of the Computers and Electrical Engineering Journal (Elsevier), and since March 2007 she has been the coordinator of the IEEE TCSC technical area of Content Management and Delivery Networks. For more information, please visit: http://www.csd.auth.gr/~avakali.

Giancarlo Fortino [a]
Carlo Mastroianni [b]
George Pallis [c,*]
Mukaddim Pathan [d]
Athena Vakali [e]
[a] *Dept. of Electronics, Informatics and Systems, University of Calabria, Rende (CS), Italy*
[b] *ICAR-CNR, Italian National Research Council, Rende (CS), Italy*
[c] *Dept. of Computer Science, University of Cyprus, Nicosia, Cyprus*
[d] *Information Engineering Laboratory, CSIRO, Canberra, Australia*
[e] *Dept. of Informatics, Aristotle University, Thessaloniki, Greece*
* Corresponding author.
E-mail addresses: g.fortino@unical.it (G. Fortino), mastroianni@icar.cnr.it (C. Mastroianni), gpallis@cs.ucy.ac.cy (G. Pallis), mukaddim.pathan@csiro.au (M. Pathan), avakali@csd.auth.gr (A. Vakali)

Available online 17 September 2011