Journal of Network and Computer Applications 32 (2009) 941–942

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

Editorial

Next generation content networks

Content Delivery Networks (CDNs) have recently been proposed to improve the performance, in terms of response time, bandwidth, accessibility, of Internet-based content delivery through coordinated content replication (Pallis and Vakali, 2006). CDNs maintain geographically distributed clusters of surrogate servers which are placed at the network edge and store copies of identical content, so that user requests can be satisfied by the optimal surrogates.

CDNs offer infrastructure and mechanisms to deliver content and services in a scalable manner, and enhance user experience. However, modern applications do not just perform retrieval or access operations on content, but also create, modify and manage it, and actively place content at appropriate locations. In order to deal with such new requirements, Content Networks (CNs) have emerged.

Content Networks are experiencing an enormous growth in many of their aspects, such as the amount of provided content, the number of involved servers and hosts, the number and variety of users, and the plethora of traditional and innovative services, often involving large bandwidth and real-time connections (Buyya et al., 2008). This evolution demands new technologies and paradigms, which can keep pace with growing complexity through the use of decentralized algorithms and protocols. This trend was explored in a recent special issue on CDNs (Fortino and Mastroianni, 2008), which focuses on the use of peer-to-peer (P2P), multi-agent, and Grid technologies. Since then, however, a new acceleration has been experienced in this field, for several reasons: the increased need for large size content, partly due to the explosion of social networks, the wider exploitation of P2P technologies, and the emerging of the new Cloud paradigm (Hayes, 2008), which let the content be managed by hundreds of server hosts transparently located in the "storage clouds".

Several approaches to content management and delivery are currently being adopted (Leighton, 2008). The "centralized hosting" approach is only acceptable for small- or medium-sized networks, due to its poor scalability properties, while the "big data canters" CDN paradigm has been successful so far, but begins to show its limits in large and dynamic systems. The management and delivery of content will mostly be guaranteed by next-generation technologies, specifically "highly distributed CDNs", which leverage hundreds of servers placed within end-user Internet Service Providers, and P2P networks, in which the downloaders, peers requesting content, can also act as uploaders, which offer content to other peers. These new paradigms are particularly efficient in scenarios where the services requested by users are video-centered and real-time, as is often the case today.

The six papers of this special issue successfully capture the state of the art in the content networking domain, and offer very interesting perspectives both on current trends and on the innovation forces that will drive research and development efforts in the coming years. In particular, significant key issues are tackled in the papers, among which: the design and implementation of collaborative CN networks, which exploit grouping and peer-to-peer techniques to provide improved performance in terms of efficiency, fault tolerance, and load balancing; the provisioning of efficient and flexible data management and data transfer services; the beneficial adoption of CN technologies for novel and emerging environments, such as Cloud Computing, streaming multimedia applications, and mobile transient networks.

The first paper of the special issue, authored by Ravi et al. is, conveniently, a comprehensive survey paper that presents a wide overview on content networks technologies and currently available systems. The paper has two key merits: firstly, it illustrates and compares a large number of devised solutions for such important issues as the generation of dynamic, personalized, and real-time content, the caching and delivery of content, and the provisioning of important security requirements; the second merit is the accurate and very useful categorization of the most important available systems with respect to the technical solutions adopted and the different scenarios for which they have been devised. Finally, the paper discusses the possible future directions for web content generation and delivery, with particular emphasis given to four emerging trends that will improve the perceived quality of service: the use of P2P for caching and delivery, the exploitation of statistics about client access patterns, the adoption of semantic-based techniques, and the quest for improved guarantees concerning the correctness of caching.

The paper authored by Fedak et al. deals with the important issue of data management in large-scale, dynamic, and highly distributed networks. Currently, data management relies on ad-hoc solutions, which may not easily be generalized. A new class of data management service is desirable to efficiently integrate the large variety of existing protocols, including client/server and P2P, and the new and emerging Cloud storage services. The authors present BitDew, an environment designed to automate large-scale data management and distribution and provide a consistent interface to data transfer protocols (Http, BitTorrent, GridFTP) and storage services (Amazon S3, Internet Back-Plane Protocol). The framework adopts a specific set of metadata to drive key data management operations such as life cycle handling, data distribution, placement, replication, and fault tolerance, with a high level of abstraction. BitDew has been evaluated as a support for the execution of Basic Local Alignment Search Tool (BLAST) bio-informatics applications; performance analysis demonstrates that the high level of abstraction and transparency is obtained with a reasonable overhead, while offering the benefits of scalability, improved performance, and fault tolerance with low programming cost.

1084-8045/$ - see front matter
doi:10.1016/j.jnca.2009.05.001

In their paper, Pathan and Buyya present an architecture for a peer arrangement of multiple CDNs that can cooperate to serve the user requests and guarantee adequate delivery performance in situations when the incoming request load is overwhelming for a single provider. A discovery protocol and a redirection algorithm are devised to route the requests to appropriate peering CDNs, so as to combine good performance and load balancing. In this approach, when any Web server reaches an overload condition, the load distribution strategy reacts to redirect loads by selecting available optimally underloaded server(s), while not compromising network proximity. This can alleviate the problems experienced with the commonly used DNS-dispatching policies for load balancing, which do not provide sufficient control on user requests. The proposal is validated with the aid of simulations, considering practical constraints and significant system parameters; performance are analyzed through metrics such as number of on-going connections, number of completions, service disruptions, and server utilization.

The paper authored by Lloret et al. shows a new architecture of Content Delivery Networks, based on the grouping of surrogates. Their work is inspired by the scalability and performance improvements that are achieved by group-based systems in many application environments. The paper illustrates the development of a two-layer CDN architecture that supports the grouping of surrogates and establishes connections among the surrogates of neighboring groups, by taking into account their "remoteness" parameter and a parameter based on the capacity of the surrogates. The authors describe the protocol developed to connect surrogates within the same group and among different groups, and the flow of the designed messages. Notably, they provide real measurements in a controlled environment, which helps to analyze the network control traffic and the performance of the surrogates in different scenarios. Measurements demonstrate that the presented CDN architecture requires low bandwidth connections to run properly, and scales very well, in this way justifying the benefits of the proposal.

In their paper, Molina et al. investigate the efficiency and benefits of a mobile transient network in which the nodes collaborate, by means of intelligent agents, to offer a multi-homed network with community services. The paper presents an architecture that supports a peer-to-peer transient community with two "modes": file download and multimedia streaming. The first mode, implemented with intelligent and cooperative agents, provides two negotiation models, fair and dynamic, in order to guarantee wider flexibility for different possible scenarios and community interactions. The multimedia streaming mode implements a content delivery system into a transient community, taking advantage of the multi-homed features of mobile devices, i.e., the capacity of mobile devices to provide multiple and heterogeneous connection paths. The authors describe an on-going work included in two EC-funded projects (Multinet and Expeshare), and analyze through simulation the suitability and performance of robust cooperative download in mobile ad-hoc networks (MANETs).

The last paper, by Broberg et al., gives an interesting insight into the new stimulus that content network technologies will receive from a field that has rapidly emerged both in the academic and in the industrial world, that is, Cloud Computing. The presented "MetaCDN" architecture aims to exploit the infrastructures provided by companies that offer on demand data storage and delivery facilities by the exploitation of a large number of hosts available in the "Storage Clouds". This approach may allow clients to obtain remarkable cost reductions with respect to the existing Content Delivery Networks, which are today out of reach for the majority of small and medium enterprizes. Moreover, MetaCDN offers a single unified namespace that easily integrates with the origin systems, and is completely transparent for end users. MetaCDN is currently under active testing. The tests reported in the paper show that the exploitation of Storage Clouds allows considerable performance improvements, in terms of throughput and response time, to be achieved, while the MetaCDN system itself introduces minimal overhead compared to the direct usage of the Storage Clouds.

This special issue is the outcome of a collaborative work. We would like to thank the editor in chief of JNCA, Mohammed Atiquzzaman, for his precious and continuous support during the review and publication process. The special issue selected 6 high-quality papers among 18 initially submitted. This was possible thanks to the work of the renowned researchers that provided their anonymous reviews. They are, in alphabetical order: Jemal Abbawajy, Ronald Geoffrey Addie, Oscar Martinez Bonastre, Cyril Briquet, Eugenio Cesario, Jiann-Liang Chen, Carmela Comito, Floriano De Rango, Giuseppe Di Fatta, Zomgmin Fei, Francesco Folino, Agostino Forestiero, Alfredo Garro, Yunhong Gu, Andrew Harrison, Xiaojun Hei, Stratos Idreos, Jussi Kangasharju, Thilo Kielmann, Panayiotis Koutsabasis, Emilio Leonardi, Claudio Lucchese, Raffaele Montella, George Pallis, Guillaume Pierre, Andrea Pugliese, Nadia Ranaldo, Farzad Safaei, Corrado Santoro, Marc Schiely, Giandomenico Spezzano, Anna Cinzia Squicciarini, Álvaro Suárez, Hiroyuki Takizawa, Domenico Talia, Michela Taufer, Emiliano Tramontana, Paolo Trunfio, Athena Vakali, Giuseppe Vizzari, Lizhe Wang, Norihiko Yoshida, Yifeng Zhu, and Eugenio Zimeo.

Finally, we are most grateful to the authors for their valuable contributions and for their willingness and efforts to improve their papers in accordance with the reviewers' suggestions and comments.

References

Buyya Rajkumar, Pathan Mukaddim, Vakali Athena, editors. Content delivery networks, Lecture Notes in Electrical Engineering, vol. 9. Berlin: Springer; 2008.

Fortino Giancarlo, Mastroianni Carlo. Enhancing content networks with P2P, grid and agent technologies. Future Generation Computer Systems 2008;24(3):177–9.

Hayes Brian. Cloud computing. Communications of the ACM 2008(7):9–11.

Leighton Tom. Improving performance on the internet. Queue 2008;6(6):20–9.

Pallis George, Vakali Athena. Insight and perspectives for content delivery networks. Communications of the ACM 2006;49(1):101–6.

Giancarlo Fortino *
DEIS University of Calabria, 87036 Rende (CS), Italy
E-mail address: fortino@deis.unical.it

Carlo Mastroianni [1]
ICAR-CNR, 87036 Rende (CS), Italy
E-mail address: mastroianni@icar.cnr.it

* Corresponding author. Tel.: +39 0984 494063; fax: +39 0984 494713.
[1] Tel.: +39 0984 831725; fax: +39 0984 839054.