Multiagent and Grid Systems – An International Journal 5 (2009) 133–135
DOI 10.3233/MGS-2009-0123
IOS Press

Guest-editorial

Special Section: Content management and delivery through P2P-based content networks

Giancarlo Fortino[a,*] and Carlo Mastroianni[b]
[a]*DEIS, University of Calabria, Via P. Bucci, 41C, 87036 Rende* (*CS*), *Italy*
[b]*ICAR-CNR, Address: Via P. Bucci 41C, 87036 Rende* (*CS*), *Italy*

Content Networks (CNs) are distributed networks/systems tailored to the management, distribution, discovery and delivery of content to Internet users. They can be considered as an extension and improvement of Content Distribution Networks that emerged previously as an innovative technology that can improve the efficiency of static, time-dependent and rich media content delivery atop large-scale IP-based networks.

This special section is dedicated to "Content Management and Delivery through P2P-based Content Networks" and contains an invited paper and three papers selected from those presented at the second edition of the UPGRADE-CN workshop, which took place at the ACM/IEEE International Symposium on High Performance Distributed Computing (HPDC 2007), in Monterey Bay, CA, USA, in June 2007. The intent of the UPGRADE-CN workshop is to provide a forum to discuss recent enhancements of Content Networks. Design and implementation of robust and efficient CNs in large-scale computing environments demand for emerging paradigms and software technologies that can improve the performance and extend the functionalities of CNs. For example, multi-agent systems, peer-to-peer (P2P) and Grid computing paradigms and technologies have demonstrated their suitability in supporting the development of content-oriented high-performance systems over large-scale, dynamic and heterogeneous environments. In particular, the best three papers presented at UPGRADE-CN'07 focus on the use of decentralized techniques and protocols, based on the P2P paradigm, and on their efficient adoption for the management and delivery of content in distributed systems. Moreover, the invited paper written by two very active researchers in this field Schiely and Felber focuses on the same topics and, particularly, on P2P-based multimedia content delivery.

These four articles were fully peer-reviewed according to the practice of this journal. They offer a very interesting and broad range of views on the actual use of peer-to-peer technologies for content distribution infrastructures. Though this research area is very recent, the selected papers already present an evaluation of real systems or prototypes, discuss architectural aspects of content-oriented systems based on decentralized algorithms, and provide simulation or analytical performance analysis. In the following the main content of these papers is briefly described.

*Corresponding author. Tel.: +39 0984 494063; E-mail: g.fortino@unical.it.

ISSN 1574-1702/09/$17.00

In their paper, Briquet, Dalem, Jodogne and de Marneffe focus on the relevant issue of optimizing the transfer of large data files in a distributed network. This paper offers a thorough comparison of widespread data transfer protocols, for example GridFTP, which is an extended FTP protocol specifically designed for high performance computational Grids, and BitTorrent, a protocol that exploits techniques such as swarming and simultaneous download of file pieces from multiple sources. The main contribution of this paper is the definition and implementation of a data transfer algorithm that combines the use of the BitTorrent P2P file sharing protocol with a novel task selection scheduling algorithm, namely "Temporal Tasks Grouping", or TTG. The TTG algorithm extends the classical data caching paradigm to a temporal basis. While data caching schemes exploit the use of cache nodes to favour the download requests issued from consumer peers that are located in the same spatial area, these schemes do not take into account the time at which download requests are issued. On the other hand, the attempt of TTG is to group those tasks that require the same input data, and schedule and serve them within a restricted lapse of time, so as to maximize the efficiency of transfers and, counter-intuitively, have a maximum of client nodes in parallel. The use of TTG can reduce the use of caching because it does not necessarily require the displacement of data to caches before the download requests; caching is used only when tasks cannot be temporally grouped to multiple client nodes.

The article written by Pathan and Buyya copes with an important practical issue that may hinder the efficient use of Content Delivery Networks, i.e. their proprietary nature. This means that CDNs are closed and do not naturally cooperate, and that an enormous amount of capital and labour is required by a single CDN to offer high performance Internet content delivery through global coverage. The article presents a novel architecture that allows distinct CDNs to coordinate and cooperate in a P2P fashion, thus improving the performance of each CDN provider and, in the end, increasing the quality of service perceived by end-users. In a cooperative multi-provider environment, users are redirected across distributed set of Web servers deployed by partnering CDNs as opposed to individual servers belonging to a single CDN. Peering is based on the Virtual Organization (VO) paradigm, and can be short-term, with the aim to prevent or resolve the creation of hotspots, or long-term, to establish a durable collaboration among CDNs. Beyond describing the collaborative architecture, the article presents a policy-based framework for SLA negotiation among peering CDNs to ensure that requests are effectively served, meeting user QoS requirements. Moreover, a performance model, based on queue networks, is used to predict user perceived performance and demonstrate the advantageous effects of peering. This model is also exploited to analyze the impact of the redirection of requests among peering CDNs and compare four different request-redirection policies.

In the paper by Schiely and Felber, the CrossFlux P2P system for media streaming is presented. P2P media streaming is an interesting approach for end-system-based media delivery over the Internet which may reduce and, sometimes, substitute the use of costly CDN infrastructures. In particular, CrossFlux is designed around four important features overcoming the shortcoming of client/server architectures: scalability, efficiency, reliability and fairness. The basic idea is that a stream to be delivered from one source to a large group of clients is split into multiple chunks that are transmitted over different delivery trees. Beyond a common tit-for-tat strategy, peers that contribute more to content distribution are rewarded by additional robustness, as each connection can be used as a backup link in the reverse path. This allows for a media streaming delivery at a more uniform rate. Moreover, CrossFlux employs adaptive algorithms to provide dynamic reorganization of delivery trees to take into account bandwidth fluctuations and content distribution optimization. In particular, stable nodes having high bandwidth are moved upward the delivery trees toward the root where they can act in a more productive way. The evaluation phase of CrossFlux based on Modelnet shows that the proposed system allows for balanced

and efficient delivery trees and fast recovery from node failures due to the backup links. The CrossFlux prototype is currently implemented in Java and deployed on Planetlab testbed.

The article written by Eger, Hoßfeld, Binzenhöfer and Kunzmann focuses on simulation techniques for BitTorrent-like P2P networks. BitTorrent is a popular protocol for P2P file delivery over the Internet. The basic idea of BitTorrent is to fragment a file into chucks that, when downloaded by peers, can be quickly uploaded to other peers. Moreover each peer controls to which other peers it uploads data so to upload to those peers providing highest download rates. BitTorrent causes a significant amount of traffic in the current Internet so it is important to define simulation techniques to analyze its performances. In this paper, two techniques for the evaluation of the protocol are compared: packet-level and flow-level simulation. While flow-level simulation shows that BitTorrent is near to optimal, packet-level simulation provides results that deviate more from the flow-level results and also from the analytical results. Nevertheless packet-level simulation is very important for studying cross-layer interactions between application and transport layer. These results can be generalized and applied to bandwidth trading among peers and peer rate control.

The quality and variety of the topics covered by these papers reflect the vitality of this research area, which was also confirmed by the success of the third edition of the UPGRADE-CN workshop, which took place at HPDC 2008, in Boston, Massachusetts, in June 2008. We hope that these papers will be a valuable source of information for researchers who want to become familiar with the field of content-oriented P2P systems.