Available online at www.sciencedirect.com
ScienceDirect
Future Generation Computer Systems 24 (2008) 177–179
FGCS
FUTURE GENERATION COMPUTER SYSTEMS
www.elsevier.com/locate/fgcs

Editorial

Special section: Enhancing content networks with P2P, Grid and Agent technologies

Giancarlo Fortino[a,*], Carlo Mastroianni[b]

[a] DEIS – University of Calabria, Via P. Bucci, 87036 Rende (CS), Italy
[b] ICAR – Italian National Research Council, Via P. Bucci, 87036 Rende (CS), Italy

Received 6 August 2007; accepted 10 August 2007
Available online 17 August 2007

Content Distribution Networks (CDNs) have recently emerged as an innovative technology for improving the efficiency of static, time-dependent and rich media content delivery atop large-scale IP based networks. This special section contains six papers selected from those presented at the first edition of the UPGRADE-CDN (the Use of P2P, Grid and Agents for the Development of CDN) workshop, which took place at the IEEE HPDC 2006 Conference, in Paris, France, in June 2006.

The intent of the UPGRADE-CDN workshop was to provide a forum for discussing recent enhancements of Content Distribution Networks, and more generally, of Content Networks, which can be defined as distributed networks/systems tailored to the management, distribution, discovery and delivery of content to Internet users. Design and implementation of robust and efficient CDNs in large-scale computing environments demand emerging paradigms and software technology able to support the core functionalities of CDNs. Agent, P2P and Grid computing paradigm and technology have already demonstrated their suitability in supporting the development of high performance systems over large-scale, dynamic and heterogeneous environments. In content distribution infrastructures, the P2P paradigm can be adopted to increase dynamism and fault-tolerance, Grid technologies can favour robustness and support for multi-organization applications, and agent based techniques and algorithms can foster intelligent behaviour and self-organization features.

The six papers selected for this special section were extended and revised by the authors and were fully peer-reviewed according to the practice of this journal. They offer a very interesting and broad range of views on the potential use of the above-mentioned new technologies for content distribution infrastructures. Though this research area emerged very recently, some of the selected papers already present an evaluation of real systems or prototypes, discuss architectural aspects of content-oriented systems based on innovative technologies, and provide simulation or analytical performance analysis.

This special section is opened by the paper written by Fortino and Russo which can be considered as a "manifesto" motivating and promoting the use of the P2P, Grid and Agent technologies and paradigms for the development of novel CDNs. The authors claim that the combined use of all three paradigms and technologies, which are characterized by dynamism, robustness, and intelligence, is expected to increase the effectiveness and boost the efficiency of future CDNs. In particular, after giving an interesting introduction to major, currently available research efforts which exploit the aforementioned paradigms and technologies, they present UPGRADE-CDN, an experimental component based platform which jointly uses P2P, Grid, and Agent based mechanisms for client request redirection, CDN monitoring, and content delivery. UPGRADE-CDN is composed of (i) GRedirector, a redirection system based on Grid, (ii) AMonitor, a system for monitoring the CDN servers and network, (iii) CoDelivery, a system for content delivery based on P2P and clusterization. The benefits of designing and implementing such components through Grid, Agents and P2P are highlighted and discussed. Moreover, UPGRADE-CDN is classified according to a recently proposed CDN taxonomy which allows for its direct comparison with currently available commercial and academic CDN systems. Finally, a performance evaluation through simulation of the CoDelivery component is provided which

* Corresponding author. Tel.: +39 0984 494063; fax: +39 0984 494713.
E-mail addresses: g.fortino@unical.it (G. Fortino), mastroianni@icar.cnr.it (C. Mastroianni).

0167-739X/$ - see front matter
doi:10.1016/j.future.2007.08.002

shows that UPGRADE-CDN has the potential to be more efficient with respect to conventional CDN architectures.

The paper by Lloret et al. focuses on the grouping of content delivery servers into logical networks according to some criteria such as similar characteristics, functionalities, type of data delivered, proximity. Current research efforts in the context of this research topic are mainly based on centralized solutions which use a single point for data delivery between groups of content delivery servers. This paper proposes a novel hierarchical architecture for the interconnection of content delivery servers selected from different groups and describes how they establish logical connections. The topology of logical connections established between servers clustered in different groups changes depending on their capacity or on their failures. The authors also present an analytical model for this interconnection architecture and results obtained for some defined parameters. Moreover, performance evaluation of the proposed architecture, deployed on a real testbed, has been carried out to obtain accurate measurements. These measurements have shown that the bandwidth consumption for managing the system is low so that the system can be effectively used in real environments.

The paper by Di Stefano and Santoro proposes a Grid based architecture for federating Content Distribution Networks (CDNs) in order to share computational resources. Such architecture, called the Content Distribution Grid (CDG), aims at using a community policy which allows each CDN to share a portion of its own resources to meet the requirements of each federated CDN so as to guarantee a better quality of service to users. As CDNs federated by a CDG belong to different organizations, organizations are mainly self-interested, each aiming at maximizing the performances of its own CDN, so the interaction among CDNs is performed on a competitive basis. To tackle this issue, the authors propose a resource sharing policy based on an offer/demand competitive model, in which resources are purchased, a certain amount of (virtual or real) money being paid for them. In the defined economical model, the CDN requesting resources (buyer) and the CDN offering resources (seller) agree on the quantity and the price by means of a utility based negotiation approach. Finally, a multi-agent system is proposed for implementing the software architecture supporting the proposed model of CDG.

In their paper, Meo and Milan focus on the use of P2P algorithms for the management of shared content in a distributed network. An active peer participates in the community by downloading contents from other nodes and by allowing others to download contents which are stored in its own memory. Efficient exchange of content is partly dependent on the behaviour of single nodes: whenever a node retrieves some new content, it may store it in the shared memory and make it available to other users, or it may "consume" it on the fly. The storing policy impacts the effectiveness with which content can be provided to users. Indeed, if only few copies of a given content are available, it may happen that, while the nodes storing those copies are off, the content is not available. On the basis of the probability that this content unavailability occurs, the authors define a notion of Quality of Service (QoS) and propose criteria for the QoS design of content management policies. Moreover, they consider and evaluate a simple policy based on these criteria. An interesting aspect of this paper is that performance evaluation is performed by an elegant analytical model based on a Markovian approach. This allows for the accurate estimation of performance indices characterized by very low values, for example blocking probabilities, which would be a difficult task if only the simulation approach were exploited.

The paper by Cannataro et al. also deals with the use of P2P algorithms for the management of content, but they focus on a practical application domain, and cope with the interoperability and information sharing between health care providers. This paper proposes a P2P framework that enables health operators of different hospitals to share and aggregate clinical information about patients and therapy effects. Patient records are mapped into a simple XML based meta-Electronic Patient Record (meta-EPR). The meta-EPR is a lightweight data structure defined to contain relevant and aggregate information extracted from the different EPRs adopted by each hospital. This approach allows hospital operators to formulate queries against a meta-EPR schema; queries are then distributed to the connected hospitals through a P2P infrastructure. The proposed framework has been fully implemented in a system called SIGMCC, which implements a view mechanism to allow for the protection of personal (patient) information against unauthorized users. As a case study, an application of the proposed meta-EPR to the cancer medical domain has been developed.

Finally, Forestiero et al. propose an approach that combines multi-agent and P2P techniques to improve the management of content in a Grid. The main novelty of their paper is the adoption of bio-inspired techniques, borrowed from the behaviour of ants. Content, specifically metadata descriptors that specify the characteristics of Grid resources, is disseminated and reorganized through the collective activity of a large number of ant-like agents, which travel the Grid through P2P interconnections established among Grid hosts. These agents perform simple operations at the local level, but together engender an advanced form of "swarm intelligence" at the global level. The authors performed a simulation analysis and showed that the proposed ant based protocol is capable of reducing the entropy of the system, efficiently propagating content, and specifically enforcing the dissemination of descriptors related to high QoS resources. The Grid information system built with the proposed protocol enables the use of a semi-informed discovery algorithm which efficiently drives query messages towards "representative peers" that maintain information about a large number of high QoS resources.

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 second edition of the UPGRADE-CN workshop, which took place at HPDC 2007, in Monterey Bay, California, in June 2007. 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 distributed systems.

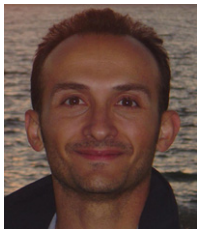

Giancarlo Fortino is an Associate Professor of Computer Science at the Department of Electronics, Computer Science and Systems of the University of Calabria, Cosenza, Italy. He received a Ph.D. in Computer Engineering from the University of Calabria, Italy, in 2000. He was visiting researcher at the International Computer Science Institute (ICSI), Berkeley (CA), USA, in the 1997 and 1999. His research interests include distributed computing and networks, agent systems, agent-oriented software engineering, streaming content distribution networks, distributed multimedia systems, Grid computing. He is currently involved in research and development projects for the application of intelligent and mobile agents (in the framework of the international AgentLink III project), Grid and content distribution networks to industry and education. He is author of more than 80 papers in international journals and conferences. He is member of IEEE, IEEE Computer and Communications Society, and ACM, and is involved in several IPCs of international conferences. Web Site: http://wwwinfo.deis.unical.it/~fortino.

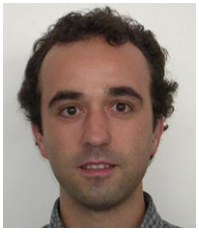

Carlo Mastroianni is a Researcher at the Institute for High Performance Computing and Networks of the Italian National Research Council (ICAR-CNR) in Cosenza, Italy. He received a Ph.D. in Computer Engineering from the University of Calabria, Italy, in 1999. His research interests include distributed systems and networks, Grid computing, peer-to-peer networks, multimedia distribution networks. He is currently involved in research and development projects for the application of P2P and Grid (in the framework of the CoreGrid and Grid.it projects) to industry and education. He is author of more than 60 papers in journals and conferences. He also is involved in several IPCs of international conferences. Web Site: http://www.icar.cnr.it/mastroianni.