Available online at www.sciencedirect.com
ScienceDirect
Future Generation Computer Systems 24 (2008) 202–212
FGCS FUTURE GENERATION COMPUTER SYSTEMS
www.elsevier.com/locate/fgcs

# An economic model for resource management in a Grid-based content distribution network

Antonella Di Stefano[a], Corrado Santoro[b,*]

[a] University of Catania, Engineering Faculty, Department of Computer Science and Telecommunication Engineering, Viale A. Doria, 6 - 95125 - Catania, Italy
[b] University of Catania, Computer Science Faculty, Department of Mathematics and Computer Science, Viale A. Doria, 6 - 95125 - Catania, Italy

Received 7 December 2006; received in revised form 11 July 2007; accepted 14 July 2007
Available online 6 August 2007

**Abstract**

This paper presents an architecture to federate Content Distribution Networks (CDNs) in order to share computational resources, thus building an infrastructure that we call a *Content Distribution Grid (CDG)*. The purpose of a CDG is to use a community policy allowing each CDN to put together a portion of its own resources to meet each other's requirements, and therefore to guarantee a stronger quality of service to users. Since CDNs of a CDG belong to different organisations, the interaction scenario can be considered *competitive*, that is, organizations are mainly self-interested and are aiming at maximizing the performances of its own system. For this reason, the resource sharing policy proposed here is based on a *offer/demand* competitive model, in which resources are *purchased* by paying for them a certain amount of (virtual or real) money. An economic model is thus derived to guide such a sale; here 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. A multi-agent system is then proposed to realise the software architecture supporting this model of CDG.


*Keywords:* Content distribution networks; Software agents; GRID; Economics; Resource sharing

## 1. Introduction

Content Distribution Networks (CDNs) [1–9] and Grid systems [10–14] are two well-known wide-area distributed systems, aiming at supporting different application scenarios. A CDN (essentially an extension to the WWW) duplicates web resources from an *origin server* to different *replica servers*, in order to provide a better quality of service to the users. A Grid system aims instead at federating the resources of several interconnected machines belonging to different administrative domains, thus building a distributed supercomputer whose resources can be leased on demand.

A main difference exists between CDNs and Grids: in a CDN, all the servers (origin and replicas) are owned by the *same organisation*, since they serve the same web site; conversely, Grid nodes typically belong to *different organisations* that have decided to federate their resources. This last aspect gives a user or an organisation the opportunity to exploit a large computational power just by offering and lending only a limited part of it. Obviously, such a resource usage has to be suitably placed under a security and safety control, in order to avoid things such as unauthorised access to a specific resource, resource exhausting, etc. This is the reason why Grid middlewares place a strong focus on the security infrastructure of grids [15].

This view of "resource federation", typical of a Grid, can be very useful also in a CDN, and this is the key aspect of the work proposed in this paper. Actually, we can assume that in a CDN, the larger the number of replicas, the higher the computational power that can be exploited, and thus the higher the quality of service provided to users; therefore, the possibility to rent or buy some resources from other "federated network nodes", as it happens in a Grid environment, allows content providers to improve such aspects. We can thus consider that the current scenario, in which there are several independent CDNs, could

* Corresponding address: Engineering Faculty, University of Catania, Department of Computer Science and Telecommunication Engineering, Viale Andrea Doria, 6 - 95125 - Catania, Italy.
*E-mail address:* csanto@diit.unict.it (C. Santoro).

0167-739X/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2007.07.014

lead to a near future environment in which such CDNs are interconnected together, exploiting one another's resources. This enables a new kind of distributed system that can be called a *Content Distribution Grid (CDG)*, which can be considered a specialisation of a traditional Grid for the specific purpose of improving *computational power* and *quality-of-service* in providing web contents.

Such a kind of system poses some new problems that have to be properly faced. The main issue is that in a traditional Grid, administrative organisations collaborate to provide a common pool of resources, but the situation changes in a CDG: each CDN of the CDG belongs to a different organisation which is mainly interested in improving its *own* performances. As a result, while the organisations of a traditional Grid are *cooperative*, organisations of a CDG are *selfish*: this means that, since the scenario moves from cooperation to *competition*, the policies and algorithms employed in traditional Grids for resource management cannot be used, and new strategies are required. Indeed, this new scenario stresses not only the policies to perform resource finding, selection and leasing, but also the security techniques to guarantee a safe and trusted access to such shared resources.

With these considerations in mind, this paper proposes a strategy for resource selection, negotiation, and allocation in a CDG. Resources, which in particular are *disk space*, *CPU time* and *network bandwidth*, are *sold* by servers belonging to a certain CDN to another CDN requesting them. An appropriate selling/purchasing strategy and a form of electronic money called *credit* are used, so a certain amount of resource can be sold at a given price (in credits), according to the agreement reached between the requester and the seller. In order to regulate negotiation and purchase, an *economic model* is derived [16], which establishes the trend of the price of a resource with respect to the quantity being offered (by sellers) and requested (by buyers). As a result, as it happens in classical economic theories [17], an *offer curve* and a *demand curve* are derived, which are used to guide the interaction in order to find a point of equilibrium in which an agreement can be reached. If such an equilibrium cannot be found, a suitable negotiation strategy, based on the concept of *utility*, is derived, in order to guide both the seller and the buyer in trying to find a new possible point of intersection (between the curves) and thus to reach agreement.

The paper is structured as follows. Section 2 describes the system model of a CDN and its modification to support a CDG; it also introduces the basics of the interaction model and the software architecture of the CDG. Section 3 deals with the economic model, focusing on the principles and the way in which negotiation is conducted, proposals are evaluated, and agreements can be reached. Section 4 compares the proposed approach with related techniques. Section 5 concludes the paper.

## 2. System model

### 2.1. Overview of CDNs

Before presenting the economic approach, let us introduce the system model that will be considered throughout the paper.

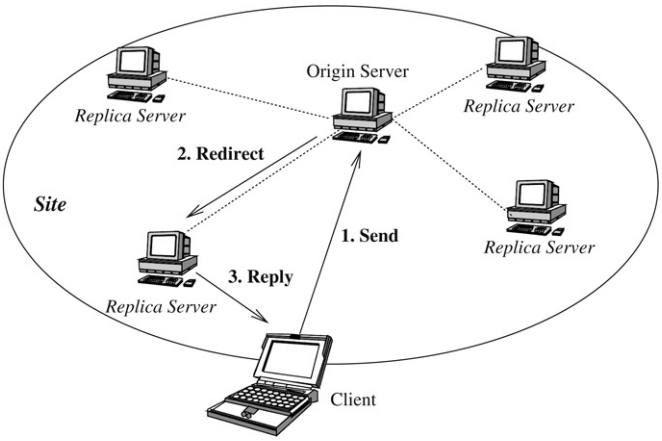


Fig. 1. Components and interactions in a site of a CDN.

We will first deal with CDNs and then we will extend the CDN model to support the concept of CDG.

By borrowing and extending the terminology used in [6], we briefly define, for a CDN, the concepts of *site*, *server*, *node* and *organisation*, together with their relationships. In a CDN, a *site* collects a set of web resources, stored on different *nodes* (i.e. single machines or clusters), connected through the Internet and managed by several *servers*; for each site, there exists an *origin server* and one or more *replica servers* (one per node) holding some content replicated from the origin server. Any *organisation* owns one or more nodes and the relevant web sites; it can be, for example, a user, a company, an enterprise, etc. For the sake of generality, we assume that each node can host several origin or replica servers belonging to different sites of the same organisation.

Referring to the outlined system model, a typical algorithm that is used in a CDN can be summarised in the following steps (sketched in Fig. 1):

1. A client that wants to interact with a site of organisation $A$ sends an HTTP request to the *origin server*.
2. The origin server determines a node, say $n^{(A)}$, which hosts a replica server for the site and is near the client.
3. The request is redirected to the selected replica server and finally handled.

Step (ii) requires us to determine the node that is considered "quite near" the client, basing this *neighbourhood property* on the capabilities (i.e. speed, throughput, propagation delay) of the network link; therefore we can assume that, in a computer network, a host $a$ is "nearer" $b$ than $c$ if e.g. the overall throughput of the link from $a$ to $b$ is greater than that from $a$ to $c$. In other words, we can assume a function $d(\cdot, \cdot)$, taking two nodes as arguments and returning a scalar numerical value that gives the distance between two such nodes; with reference to the example above, we say that $d(a, b) < d(a, c)$.

### 2.2. From CDNs to CDG

The above schema works taking into consideration the fact that all involved nodes belong to the same organisation. The *CDG* proposed here consists of many CDNs "federated