Economy-based Content Replication for Peering Content Delivery Networks

Al-Mukaddim Khan Pathan[1] and Rajkumar Buyya
Grid Computing and Distributed Systems (GRIDS) Laboratory
Department of Computer Science and Software Engineering
The University of Melbourne, Parkville, VIC 3010, Australia
{apathan,raj}@csse.unimelb.edu.au

Abstract

Existing Content Delivery Networks (CDNs) by nature are closed delivery networks which do not cooperate with other CDNs and in practice, islands of CDNs are formed. The current logical separation between contents and services in this context results in two content networking domains. In addition to that, meeting the Quality of Service requirements of clients according to negotiated Service Level Agreements is crucial for a CDN. Present trends in content networks and content networking capabilities give rise to the interest in interconnecting these networks. Hence, in this paper, we present an open, scalable, and Service-Oriented Architecture (SOA)-based system that assist the creation of open Content and Service Delivery Networks (CSDNs), which scale and support sharing of resources through peering with other CSDNs. To encourage resource sharing and peering arrangements between different CDN providers at global level, we propose using market-based models by introducing an economy-based strategy for content replication.

1. Introduction

Content Delivery Networks (CDNs) [1], which evolved first in 1998, replicate content over several mirrored Web servers, strategically placed at various locations around the globe to deal with *flash crowds* and to enhance response time. In a typical content delivery environment Web server clusters are located at the edge of the network to which the clients are connected. A content provider can sign up with a CDN provider for service and have its content placed on the content servers. The content is replicated either on-demand when clients request for it, or it can be replicated beforehand, by pushing the content to the surrogate servers. A client is served with the content from the nearby replicated Web server. Thus the clients end up unknowingly communicating with a replicated CDN server close to them and retrieves files from that server. Figure 1 depicts the different content/services served by the CDN to the clients.

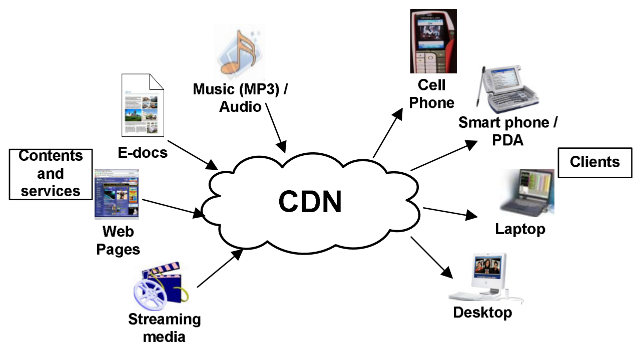

Figure 1: Content/services provided by a CDN

1.1. Motivation and scope

Existing CDNs are proprietary in nature–individual companies own and operate them. Each of them comprises of own closed delivery network, which is expensive to set up and maintain. They also have limited scalability. Running a global CDN is even more costly, requiring an enormous amount of capital and labor. In addition, content providers typically subscribe to one CDN provider and thus can not use the resources of multiple CDNs at the same time. Such a closed and non-cooperative model results in "islands" of CDNs. Moreover, the logical separation between content and services (e.g. application for video processing) under the "content distribution" and "content services" domains is undesirable considering the ongoing trend in content networking. A unified network that supports coordinated composition and delivery of content and services would be much better.

Furthermore, commercial CDNs make specific commitments to their customers by signing a Service Level Agreement (SLA). It describes provider's commitment and specifies penalties if those

[1] The author is in the first year of his candidature as postgraduate research student and Associate Professor Buyya is the supervisor.

commitments are not met. So, if a particular provider is unable to provide quality of service to the client requests, it may result in SLA violation and end up costing the provider. To cut expenses and to avoid adverse business impact, peering with other providers could be a solution to consider the issues stated above.

In this paper, we present a model for an open, scalable and SOA-based system. This system helps to create open Content and Service Delivery Networks (CSDNs) [3] that scale well and share resources with other CSDNs through peering, thus evolving past the current landscape where "islands" of CDNs exist. To encourage peering among CDNs at global level, we propose using market-based models in resource allocation and management inspired from their successful utilization in the management of autonomous resources, especially in global grids [4]. Hence, we introduce an economy-based replication strategy that involves on-demand placement of outsourced content to the surrogates of peering CDNs. The use of economic mechanisms in this context has the following benefits:

- We propose a Virtual Organization (VO) model to peer the CDNs. A VO consists of real CDN providers, which are self-interested and autonomous stake-holders. Hence, an economic model is suitable to represent this scenario and to regulate interactions among the participants. The problem can be manageable in this way through analyzing emergent marketplace behavior.
- A peering environment of content networks is highly dynamic in nature, where the availability of resources changes over time. Thus, an economic model is appropriate to exploit the dynamism of the market to make more informed decisions on the fly.
- An economic model could be the basis of a self-regulating replication strategy that dynamically adapts to the changes in the client request patterns.

The rest of the paper is structured as follows: Section 2 establishes the significance and relevance of our proposal; Section 3 addresses the shortcomings of related work; Section 4 presents the proposed model; Section 5 enlightens the economy-based model for content replication; and Section 6 provides a summary with expected contributions and future directions.

2. Significance of peering among CDNs

In our approach, a VO is formed through the coordination of Web server clusters operated by different CDNs who have come together to share resources and to collaborate on common goal(s). A VO in the peering CDN environment may vary in terms of purpose, scope, size, and duration. VOs in such an environment are of two types: *short-term on-demand VOs* and *long-term VOs* with established SLAs.

A short-term VO is formed for short time duration, based on current client request pattern to prevent the generation of *hotspots*. Consider the following scenario as a motivation for short-term VO formation. Suppose that the content of www.cnn.com is hosted by the CDN provider Akamai [5]. Akamai's Web servers receive significant client requests to serve the latest content on behalf of www.cnn.com. A sudden news outburst, demanding to the clients (e.g. 9/11 incident in USA), may cause heavy workload on Akamai's Point-of-Presences (POPs) in a particular region. As a result a *hotspot* can be generated. It could cause Akamai's POPs in that region to be unable to cope with the strain. Eventually the Web servers will be totally overwhelmed with the sudden increase in traffic, and CNN's Web site will be temporarily unavailable. Such sudden spikes in Web content requests is termed as *flash crowd* [2] or *SlashDot effect* [6].

In the peering CDN environment, the generation of hotspots due to flash crowd can be resolved through the formation of short-term VOs. A short-term VO intervenes with sudden spike in requests for particular Web content(s), which results in heavy workload on certain Web server(s) of a particular CDN. Hence, Web servers of peering CDNs form a goal-oriented constellation of distributed semi-autonomous entities and excess load is distributed to the less loaded Web servers of other CDNs. Such peering arrangement should be automated within a short time frame to address the evolving situation. A short-term VO is formed on-demand and the policy for such VO formation is established dynamically to handle the evolved situation. The short-term VO is phased out when the workload returns to normal.

On the other hand, a long-term VO is formed for events which may be known in advance. A long-term VO remains for the duration of the event. The formation of long-term VO compliments the existence of established policies and negotiated SLAs among the participating entities. To better understand the formation of long-term VO, consider the following scenario. Suppose that the ICC Cricket World Cup 2007 is being held in the Caribbean, and www.cricinfo.com is supposed to provide live media coverage. As a content provider, www.cricinfo.com has an exclusive SLA with the CDN provider, Akamai [5]. However, Akamai doesn't have a POP in Trinidad and Tobago (a Caribbean island), where most of the cricket matches will be held. As being the host of most of the cricket matches, people of this particular part of Caribbean are expected to have enormous interest in the live coverage provided by www.cricinfo.com. Since Akamai is expected to be aware of such event well in advance, its management can take necessary initiatives to deal with the evolving situation. In order to provide better service to the clients, Akamai management might decide to place its surrogates in Trinidad and Tobago, or they might use their other distant edge servers. Firstly, placing new surrogates just for one particular event would be costly and might not be useful after the event. On the other hand, Akamai risks its reputation if it can't provide agreed quality of service for client requests, which could violate the SLA and still cause profit reduction. Hence, the solution for Akamai could involve peering with other CDN provider(s) to form a VO in order to deliver the service that it could not provide otherwise. Automation for long-term VO formation is not essential since such situation is known before-hand.

Thus, by collaborating with other CDN providers though the formation of VO, content networks can better satisfy the evolving needs of their customers and meet their QoS requirements.

2.1. Research issues

The key challenges that need to be addressed in various aspects of peering CDNs are:

***Load Distribution*:**

- How to ensure reduced server load, less bandwidth consumption (by particular CDN server) and improve the performance of content delivery?

***Coordination of CDNs*:**

- What kind of coordination mechanisms need to be in place which ensure effectiveness, and allow scalability and growth of cooperative CDNs?

***Service and policy management*:**

- How to make a value-added service into an infrastructure service that is accessible to the customers?
- What types of Service Level Agreements (SLA) are to be negotiated among CDN participants? What policies can be generated to support SLA negotiation?
- How can autonomous policy negotiation happen in time to form a time-critical short-term VO?

***Pricing of contents and services*:**

- How do CDN providers achieve maximum profit in a competitive environment, yet maintain the equilibrium of supply and demand?

3. Related work

In this section, we outline the efforts for internetworking of Content Delivery Networks:

The IETF RFC document [7] proposes a Content Distribution Internetworking (CDI) Model that allows the CDNs to have a means of affiliating their delivery and distribution infrastructure with other CDNs who have content to distribute. An architecture for Content Distribution Internetworking (CDI) is presented in [8]. It discusses the design, implementation and evaluation of only a protocols architecture for cooperation among separately administered CDNs.

A peering system for content delivery workloads in a federated, multi-provider infrastructure has been presented in [12], but the peering strategy, resource provisioning and performance guarantees among partnering CDNs is unexplored in this work. CDN brokering [9] allows one CDN to intelligently redirect clients dynamically to other CDNs in that domain. The drawback is that, mechanism for IDNS is proprietary in nature and might not be suitable for a generic CDI architecture. Also, depending solely on DNS-based routing is ill-advised due to coarse control over requests as a result of ISP and client caching of DNS information.

Most of the works mentioned above do not virtualize multiple providers for cooperative management and delivery of content in a peering environment. Hence, our contribution lies in designing an effective peering mechanism that endeavors to address the limitations of previous related work, while respecting client performance requirements through proper policy management for negotiated SLAs.

4. The model for peering CDNs

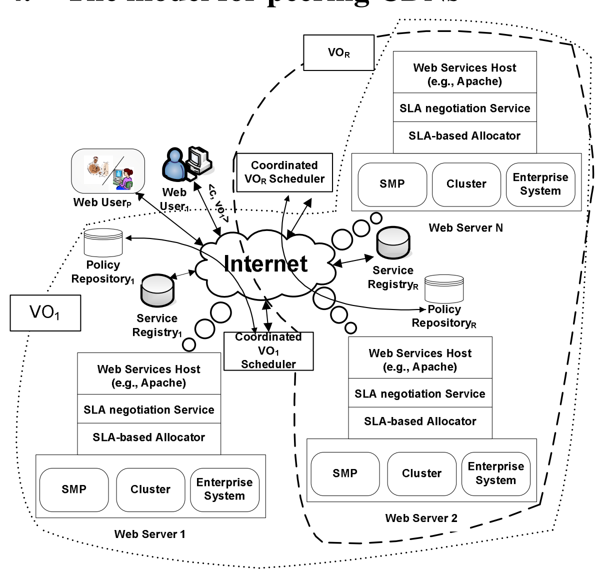

Figure 2: Architecture of open, scalable, SOA-based system to assist the creation of cooperative and coordinated Content and Service Delivery Networks

Our proposed system endeavors to solve the problem of non-cooperative *islands* of CDNs. It anticipates ensuring quality of services based on SLA negotiation, and addresses the problem of logical separation in content networking domain. Architecture for such a system is shown in Figure 2. In the figure, each Web server represents a CSDN. In the proposed VO-based model for forming Content and Service Delivery Networks (CSDNs), the formation of a VO may be stand alone or may be composed of a hierarchy of regional, national and international VOs. A VO consists of Web servers from multiple CSDNs, a coordinated VO scheduler, a service registry, and a policy repository. Web servers within each CSDN are capable of delivering services in order to meet the QoS requirements of the clients. Brief description of VO components can be found in [3].

4.1. Policy management to support SLAs

A policy in the context of peering CSDNs would be statements that are agreed upon by the participants within a VO. These statements define what type of contents and services can be moved out to a CSDN node, what resources can be shared between the VO participants, what measures are to be taken to ensure Quality of Service (QoS) based on negotiated SLA, and what type of programs/data must be executed at the origin servers. Within our proposed VO-model based CSDN architecture we apply the standard policy framework defined by the IETF/DMTF [10].

In the standard policy framework, the *admin domain* refers to an entity which administers, manages and controls access to resources within the *system* boundary. An administrator uses the *policy management tools* to define the policies to be enforced in the system. The *policy enforcement points (PEPs)* are logical entities within the system boundary, which are responsible for taking action to enforce the defined policies. The *policy repository* stores polices generated by the administrators using the policy management tools. The *policy decision point* is responsible for retrieving policies from the policy repository, for interpreting them, and for deciding on the set of policies to be enforced by the PEPs.

The proposed model for CSDN in Figure 2 can be mapped to the basic policy framework. The policy repository in Figure 2 is responsible for storing policies generated by the policy management tool used by the VO administrator. The distribution network and the Web server components (i.e. SLA negotiation service, SLA based Allocator) are instances of the policy enforcement points (PEPs), which enforce the CSDN policies stored in the repository. Each VO scheduler is an instance of the policy decision point (PDP), and determines the set of policies to be enforced at the time of peering among CSDNs. The policy management tool is administrator dependent and it is not shown in Figure 2. A direct benefit of using such policy-based cooperative architecture is to reduce the operating cost of CSDNs and to meet clients' QoS requirements according to negotiated SLA.

5. The economic model

Standardized economic concepts can be deployed for content replication within the structure of the VO model for CSDN. The aim is to share the network opening among CDNs to minimize the network latency perceived by clients and to restrict the generation of *hotspots* due to the excessive client requests for some content. But such peering may result in free-riding where some CSDN providers use other's resources free of charge. Such free-riding can be avoided through using an auction model for replicating content among CSDN servers, ensuring that participation in a VO is due to profit motivation. The auction model should be able to provide incentives to all parties.

In our economy-based model, the goal of the auction protocol is to select the cheapest suitable Web server in order to replicate content there. Here we apply the buyer-driven auction mechanism, which is a type of Vickrey auction [11] . Vickrey auctions are second-price sealed-bid auctions with low messaging overhead, efficiency of allocations and lack of counter speculation. They involve a single negotiation round in which each bidder submits a bid to the auctioneer. Other bidders cannot see the bid. The bidding agent which makes the highest bid wins the auction but pays the price of the second-highest bid.

In our case, each CSDN provider is both a buyer and seller of its resources. A coordinated VO scheduler is the *auctioneer* in the peering CSDN environment that is responsible for holding auction within the VO. It starts an auction on behalf of a CSDN provider (i.e. buyer) for finding suitable surrogate server(s) in order to perform content replication. A *Buyer* buys the storage space of Web server(s) of peering CSDNs in a particular region which incurs excessive client requests. *Sellers* are CSDN providers who sell the storage space of their Web server to the buyer. SLA-based allocator is the *bidding agent* that resides in each Web server. SLA negotiation module is the *communication mediator*, which is responsible for establishing and maintaining peering communications infrastructure. It also propagates auction messages between auctioneer and bidding agents.

The auctioneer starts an auction not for selling an item (i.e. allocation), but for buying it. Bidding agents bid with the price they are willing to sell the allocation of their Web servers. One bidder can not see the bid of other bidders. Auctioneer gathers bids from the bidders and selects the lowest bidding agent(s) as the winner and the winner is paid second-lowest bidding price. In other words, our economic model uses a reverse Vickrey auction. In the economic model, we also assume that the auction participants are trustworthy. Due to this nature, a mendacious behavior from a provider is not expected to be usual. Hence, over-provisioning of resources by harnessing data through VO membership, or modifying and falsifying of content by some rogue CSDN providers is not allowed in the system model.

Having known the internals of the economic model, we now discuss the steps for formation of VO in our economic model:

1. A CSDN provider (buyer) realizes the need to replicate content to the surrogates of peering CSDNs. The buyer internally determines the maximum payable amount (expressed by *Payoff Value*) and announces its *Auction Policy*. The auctioneer starts auction on the buyer's behalf.
2. The bidding agent of seller (other CSDN providers) uses a *Bidding Function* to determine the bidding amount.
3. The auctioneer collects bids from the bidding agents and selects the lowest bidding surrogate(s) as winner and a winner is paid by the amount of second-lowest bid.
4. Hence, a VO consisting of the buyer and seller (i.e. winning bidders) CSDNs is formed and content is replicated to the winners' surrogates.
5. Re-negotiation through auction takes place when either of the following holds: (a) A seller varies its demand after winning; (b) Seller finds that holding replicated content is no longer economically beneficial for it; (c) A more competitively priced CSDN provider (except winner(s)) comes up.

5.1. System model

Let us assume that ***N*** denotes the set of CSDN providers and ***C*** is the set of contents. For a given content $c \in \boldsymbol{C}$, S_c is denoted as the size of the content. We also define $\delta(c_i, c_j)$ as the similarity function between two contents two contents c_i and c_j, $0 \le \delta(c_i, c_j) \le 1$. The k-th arriving request to a given CSDN provider at time t_k is r_k – composed of r_k^c and r_k^l, where r_k^c is the requesting content in ***C*** and r_k^l the requesting location. $\phi(r_i^l, r_j^l)$ is the similarity function of two requests in terms of distance, $0 \le \phi(r_i^l, r_j^l) \le 1$. We specify D as the delay threshold within which each request should be delivered.

Auction initiation: An auction is initiated when a provider can not deliver the requested content to a client satisfying its QoS requirements. We express the QoS requirements as the user perceived network delay (i.e. response time). The response time of r_i at time t_k is $RT(r_i, t_k)$, where $RT(r_i, t_k) = 0$ if r_i is already served and $RT(r_i, t_k) > 0$ if r_i is not yet served. We anticipate measuring the response time based on network topology, bandwidth, and queuing delay. If there exists a request r_i such that $RT(r_i, t_k) > D$, where $i \le k$, a provider (buyer) realizes monetary penalty or loss goodwill for not satisfying the client QoS requirements. Hence, it decides to replicate content by announcing its service requirements as Auction Policy A_P and internally calculates the maximum payable amount (*Payoff Value*). The *Payoff Value* $P_{\max}$ is calculated as, $P_{\max}$ = (the managing cost for r_k^c) + (the expected profit from r_k^c). The Auction Policy A_P consists of:

– *Storage requirement:* The storage space required to replicate content c, defined as S_c (MB or GB).
– *Delay threshold*: The time within which the content should be delivered, denoted as D.
– *Preferences:* The buyer's bias for surrogate server(s) in a potential *hotspot* region.

Bidding function: The bidders (other providers) bid with the amount as determined by the *Bidding Function* $B_i(r_k) = S_i(r_k^c) + ER_i(r_k^c, t_k, n) + \psi_i(A_P)$, where $S_i(r_k^c)$ is the storage cost incurred by seller i to replicate content r_k^c, $ER_i(r_k^c, t_k, n)$ is the expected revenue of the content r_k^c at time t_k during the following n content requests, and $\psi_i(A_P)$ is a function to reflect sellers interest in the bidding of A_P. Expected revenue depends on the request pattern for next n content requests, and on the history of request for similar content. Hence, it is defined as:

$$ER(r_k^c, t_k, n) = \alpha \sum_{j=k+1}^{k+n} \delta(r_k^c, r_j^c)\phi(r_k^l, r_j^l) + (1-\alpha) \sum_{i=1}^{k-1} \delta(r_k^c, r_i^c)\phi(r_k^l, r_i^l)$$

where $0 \le \alpha \le 1$.

Auction termination: Auctioneer (Coordinated VO scheduler) collects bids from the bidding agents, and selects the winner(s) with the lowest bid. Afterwards, content is replicated there. The winning bidders are paid the second-price bid and a VO is formed consisting of the buyer and seller CSDNs (winners) based on a common goal to replicate content, and to serve it to the end-users in an efficient manner. The viability of a VO may change depending on the demand of content and the participant's economic gain. A VO participant should be able to adapt to a change in context within peering CSDN environment. Hence, re-negotiation should take place among the VO participants to either disband or rearrange the VO into a new organization that better fits the prevailing circumstances.

6. Summary and Future Work

In this paper, we have presented an open, scalable and Service-Oriented Architecture (SOA)-based system to assist the creation of open Content and Service Delivery Networks according to a VO model. To encourage resource sharing and peering among different CDN providers at global level, we propose the use of market models. Hence, we introduce an economy-based content replication strategy based on auction protocol for replicating content in surrogates of peering CDNs. The use of economic concepts in this context provides a solid basis for rational agents in peering CSDN environment to decide whether to attend in peering constellation. Use of the economic model may be the basis of a replication mechanism that dynamically adapts to the changes in content request pattern, and make replication decision to the surrogates of peering CSDNs in areas which exhibits the potential to generate Web hotspots.

Table 1: Future research activities and timeline

Timeline	Work Descriptions
Mar. 2007–Aug. 2007 (Stage 2)	Implement the architecture for peering CSDNs in a realistic simulation environment to show its effectiveness. Evaluate performance of the proposed economy-based content replication strategy.
Sept. 2007–Mar. 2008 (Stage 3)	Implement and evaluate an effective load balancing mechanism for peering CSDNs in a realistic simulation environment
Apr. 2008–Sept. 2008 (Stage 4)	Develop an effective request assignment and redirection policy and evaluate it through simulation
Oct. 2008–Mar. 2009 (Stage 5)	Develop of a prototype system for peering CDNs and deploy it in a real-world test bed in strategic locations across the globe to evaluate performance

None of the work done in content internetworking domain has exploited a successful peering arrangement through the use of economic concepts. Hence, realizing the VO model for forming CSDNs, the economic model for content replication, and deployment of the policy framework should be a timely contribution to the ongoing content-networking trend.

Our research on peering content delivery networks will follow a multistage process. We have already completed Stage 1 which included the completion of a high-level framework for peering CSDNs, with the roles, responsibilities, and interaction patterns of components according to a policy framework. Our future activities are divided into 4 stages, which are outlined in Table 1. For more information, please visit the project Website at: www.gridbus.org/cdn

Acknowledgements

We are thankful to James Broberg, Kyong Hoon Kim and Kris Bubendorfer of the University of Melbourne for sharing thoughts on this topic, and for discussion regarding the analytical formulation of system model.

References

[1] G. Peng, *CDN: Content Distribution Network*, Technical Report TR-125, Experimental Computer Systems Lab, Department of Computer Science, State University of New York, Stony Brook, NY 2003.

[2] M. Arlitt, and T. Jin, *A Workload Characterization Study of 1998 World Cup Web Site*, IEEE Network, May/June 2000, pp. 30-37.

[3] R. Buyya, A. M. K. Pathan, J. Broberg, and Z. Tari, *A Case for Peering of Content Delivery Networks*, IEEE DSOnline, 7(10), Los Alamitos, CA, USA, Oct. 2006.

[4] R. Buyya, D. Abrahamson, and S. Venugopal, *The Grid Economy,* Proc. of the IEEE, 93(3), 2005, pp. 698-714.

[5] Akamai Technologies, Inc. www.akamai.com, 2006.

[6] S. Adler, *The SlashDot Effect: An Analysis of Three Internet Publications*, Linux Gazette, Vol. 38, 1999.

[7] M. Day, B. Cain, G. Tomlinson, and P. Rzewski, *A Model for Content Internetworking*, RFC 3466, 2003.

[8] E. Turrini, *An Architecture for Content Distribution Internetworking*, Technical Report UBLCS-2004-2, University of Bologna, Italy, March 2004.

[9] A. Biliris, C. Cranor, F. Douglis, M. Rabinovich, S. Sibal, O. Spatscheck, and W. Sturm, *CDN brokering*, Computer Communications, 25(4), March 2002.

[10] A. Westerinen, J. Schnizlein, J. Strassner, M. Scherling, B. Quinn, S. Herzog, A. Huynh, M. Carlson, J. Perry, and S. Waldbusser, *Terminology for Policy Based Management*, IETF RFC 3198, Nov. 2001.

[11] W. Vickrey, *Counterspeculation, Auctions, and Competitive Sealed Tenders*, The Journal of Finance, 16(1): 8-37, March 1961.

[12] L. Amini, A. Shaikh, and H. Schulzrinne, *Effective Peering for Multi-Provider Content Delivery Services*, In Proc. of IEEE INFOCOM, 2004, pp. 850-861.