Sustainable Computing: Informatics and Systems 6 (2015) 58–68

Contents lists available at ScienceDirect

Sustainable Computing: Informatics and Systems

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

Energy-efficient content delivery networks using cluster shutdown

Vimal Mathew[a], Ramesh K. Sitaraman[a,b], Prashant Shenoy[a]

[a] University of Massachusetts, Amherst, United States
[b] Akamai Technologies Inc., United States

ARTICLE INFO

Article history:
Received 16 November 2013
Received in revised form 8 May 2014
Accepted 28 May 2014

Keywords:
Content delivery networks
Internet-scale distributed systems
Power management
Load balancing
Energy efficiency

ABSTRACT

Content delivery networks (CDNs) are an important class of Internet-scale distributed systems that deliver web, streaming, and application content to end users. A commercial CDN could comprise hundreds of thousands of servers deployed in over thousand clusters across the globe and incurs significant energy costs for powering and cooling their servers. Since energy costs are a significant component of the total operating expense of a CDN, we propose and explore a novel technique called cluster shutdown that turns off an entire cluster of servers of a CDN that is deployed within a data center. By doing so, cluster shutdown saves not just the power consumed by the servers but also the power needed for cooling those servers. We present an algorithm for cluster shutdown that is based on realistic power models for servers and cooling equipment and can be implemented as a part of the global load balancer of a CDN. We evaluate our technique using extensive real-world traces from a large commercial CDN to show that cluster shutdown can reduce the system-wide energy usage by 67%. Further, much of the energy savings are obtainable without sacrificing either bandwidth costs or end-user performance. In addition, 79% of the optimal savings are attainable even if each cluster is limited to at most one shutdown per day, reducing the required operational overhead. Finally, we argue that cluster shutdown has intrinsic architectural advantages over the well-studied server shutdown techniques in the CDN context, and show that it saves more energy than server shutdown in a wide range of operating regimes.

1. Introduction

Large Internet-scale distributed systems deploy hundreds of thousands of servers in thousands of data centers around the world. Such systems currently provide the core distributed infrastructure for many popular Internet applications that drive business, e-commerce, entertainment, news, and social networking. The energy cost of operating an Internet-scale system is already a significant fraction of the total cost of ownership (TCO) [1]. The environmental implications are equally important. A large distributed platform with 100,000 servers will expend roughly 190,000 MWH per year, enough energy to sustain more than 10,000 households. In 2005, the total data center power consumption was already 1% of the total US power consumption while causing as much emissions as a mid-sized nation such as Argentina. Further, with the deployment of new services and the rapid growth of the Internet, the energy consumption of data centers is expected to grow at a rapid pace of more than 15% per year in the foreseeable future [2]. These factors necessitate rearchitecting Internet-scale systems to include energy optimization as a first-order principle.

An important Internet-scale distributed system to have emerged in the past decade is the *content delivery network* (CDN, for short) that delivers web content, web and IP-based applications, downloads, and streaming media to end-users (i.e., *clients*) around the world. A large CDN, such as that of a commercial provider like Akamai, consists of hundreds of thousands of servers located in over a thousand data centers around the world and account for a significant fraction of the world's enterprise-quality web and streaming media traffic today [3]. The servers of a CDN are deployed in *clusters* where each cluster consists of servers in a particular data center in a specific geographic location. The clusters are typically widely deployed on the "edges" of the Internet in most major geographies and ISPs around the world so as to be proximal to clients. Clusters can vary in size from tens of servers in a small Tier-3 ISP to thousands of servers in a large Tier-1 ISP.

The primary goal of a CDN is to serve content such as web pages, videos, and applications with high availability and performance to end users. The key component that ensures availability and performance is the CDN's *load balancing* system that assigns each incoming request to a server that can serve that request. To this end, a CDN's load balancing system routes each user's request

---

E-mail addresses: vml@cs.umass.edu, vml.mathew@gmail.com (V. Mathew).

http://dx.doi.org/10.1016/j.suscom.2014.05.004
2210-5379/© 2014 Elsevier Inc. All rights reserved.

to a server that is *live* and *not overloaded*. Further, to enhance performance, a CDN ensures that each user request is routed to a server that is *proximal* to that user. The proximity (in a network sense) ensures that the network path between the user's device and the CDN's server has low latency and loss. The process of routing user requests to servers is a two stage process. A *global load balancer* (called GLB) assigns the user to a cluster of servers based on the availability of server resources in the cluster, performance, and bandwidth costs. A *local load balancer* (called LLB) assigns the user to a specific server that is capable of serving the requested content within the chosen cluster. The choice of server is dictated by server liveness, content footprint, and current server loads with respect to their capacities. A comprehensive discussion of the rationale and system architecture of CDNs is available in [3].

1.1. Cluster shutdown: a technique for energy reduction

A number of approaches are relevant to reducing the energy consumption of CDNs. In the past two decades, there has been significant work in improving the energy efficiency of servers and data centers. Such improvements yield energy savings in any deployed distributed system, including CDNs. For instance, the switch to multi-core architectures, the increasing use of SSDs, static power management (SPM) to decrease energy use when the servers are idle, use of low-power servers [4], and power scaling techniques such as Dynamic Voltage and Frequency Scaling (DVFS) [5,6] all help reduce CDN energy consumption. Similarly, the use of temperature controlled fans and advances in air flow management have led to increases in cooling efficiency [7,8].

In addition to the above generic methods, there has been recent work on CDN-specific techniques that incorporate the ability to turn off individual servers during periods of low load to reduce the energy consumption [9]. Such a *server shutdown* technique is implemented within the local load balancer (LLB) of the CDN. The work in [9] shows that the availability, performance, and operational costs of the CDN remain unaffected when turning off servers to save energy. In this paper, we propose and evaluate a novel CDN-specific technique called cluster shutdown where an entire cluster of servers in a CDN data center can be turned off. Cluster shutdown is easily integrated into the global load balancer (GLB) that will now have the ability to move all load away from a cluster and shut it down. However, since the granularity of energy management is to turn off entire clusters or leave them entirely on, the technique does not have the ability to turn off individual servers (e.g., a fraction of a cluster). In contrast, the server shutdown technique studied in [9] has the ability to shutdown individual servers within the cluster depending on the load, but has has no ability to control how much load enters a cluster. Therefore, in this sense, the two techniques are complementary and may be implemented together. While cluster shutdown has not been studied before in the CDN context, it has certain natural advantages that make it worthy of consideration for CDN energy reduction.

(1) *Redundant deployments.* Large CDNs such as Akamai can have over a thousand clusters deployed in data centers around the world [3] with more than a dozen redundant deployments in any given geographical area. Thus, when some clusters near a user are shutdown during off-peak hours, other nearby active clusters can continue to provide CDN service to users and ensure good availability and performance. In fact, one of the contributions of this work is determining the impact of cluster shutdowns on user performance.

(2) *Cluster shutdown is consistent with the original CDN architectural design.* Each cluster in a CDN is often architected to be a self-sufficient unit with enough processing and disk storage to serve the content and application domains that are assigned to it [3]. In particular, there is limited data dependency and resource sharing across clusters. Thus, cluster shutdown can be implemented with little or no changes to the CDN's original architecture. In contrast, servers within a cluster are closely linked in a fine-grained fashion and they cooperatively cache and serve the incoming requests. For instance, servers within the same cluster cooperatively store application state and content for user requests served by that cluster. Thus, shutting down individual servers for energy savings requires greater migration of state and content between servers in a cluster at levels not customary in a CDN today. Cluster shutdown, in contrast, does not require state migration and cached content is already replicated across clusters for fault-tolerance purposes, which ensures that availability is not impacted by shutting down a cluster. In this sense, cluster shutdown is a better architectural design choice for energy management than server shutdown.

(3) *Cluster shutdown has the potential to save on cooling power in addition to IT power.* A key advantage of cluster shutdown is that the all of the energy consumed by a cluster, which includes energy consumed by the servers, the network equipment, and the cooling within that cluster, can be saved when a cluster is turned off. In contrast, a server shutdown technique will typically turn off a fraction of the servers within the cluster and will require the networking and cooling equipment to stay on. The cooling equipment is not energy proportional – thus turning off a fraction of the servers only saves energy consumed by those servers and does not yield a proportionate reduction in cooling costs.

For cluster shutdown to be effective, a CDN would need to have control over all of its energy consumption, i.e., both IT (such as servers) and cooling equipment. Such a scenario is reasonable given the trend for CDN's to opt for self-contained, modular [10], or containerized [11] deployments. With such deployments a CDN can manage the power consumption of its own cluster, independent of other tenants in the data center – an advantage for a CDN that wants manage its power consumption closely. The savings that can be obtained from reducing cooling costs can have a significant impact on the total energy expenditure of a cluster. The key reason is that the energy consumed by cooling equipment is a significant fraction of the energy expended by the IT equipment[1] such as servers. The ratio of total energy to IT energy is a standard metric called PUE (Power Usage Effectiveness) that has a typical value[2] of about 2 implying cooling energy is roughly equal to IT energy in typical data center deployments. But in more recent energy-efficient designs, PUE is smaller but cooling energy is still a significant fraction of the IT energy. Further, cooling energy consumption is not power-proportional since cooling still takes a significant amount of energy even when the servers have low utilization and are not producing much heat, resulting in disproportional energy savings when cooling is shutdown entirely (cf. Fig. 1(a)).

Despite these advantages, a cluster shutdown technique is not without disadvantages when compared to server shutdown [9]. Shutting down a cluster and moving all its users to other clusters might degrade performance for users if they have to go "farther away" in the network sense for their content. Further, moving traffic across clusters has the potential of increasing the bandwidth cost, even if it reduces energy. A primary focus of our work then is to evaluate the energy reduction provided by cluster shutdown and how it trades off against potential degradation in performance and increases in bandwidth costs.

---

[1] IT energy expenditure is primarily the energy consumed by the servers, since the networking equipment consume significantly less. Likewise, cooling energy expenditure is dominated by the energy consumed by the chillers [12].

[2] In a survey by the Uptime Institute [13] in July 2012, data centers reported an average PUE between 1.8 and 1.89. Other estimates place PUEs even higher.