

International Journal of Advance Research in Computer Science and Management Studies

Research Article / Survey Paper / Case Study

Available online at: www.ijarcsms.com

Cloud Computing: Performance analysis of data center while optimizing the power consumption

Shilpa KumariDepartment of Computer Science and Engineering
Career Point University Hamirpur
Himachal Pradesh, India

Abstract: Although power consumption is to be considered as the main issue towards the growth of data center in cloud computing and requires more attention. Many researchers are working for this particular issue, and most of them are focused towards the power consumption of complete data center network. But while considering all the previous research, very few of them combine both the QoS and power consumption simultaneously. In this paper we will provide detailed analysis of different techniques to consolidate the path of traffic flow that affect the power consumption of datacenter, and how individually these techniques affect the performance of data center network. Also we will provide our review that what additionally we can do to improve the performance while optimizing the power consumption of whole network.

Keywords: cloud; datacenter; topology; traffic flow; path consolidation

I. INTRODUCTION

Cloud computing, often referred to as simply "the cloud," is the delivery of on-demand computing resources i.e. everything from applications to data centers over the Internet on a pay-for-use basis [1]. It is typically defined as a type of computing that relies on sharing computing resources rather than having local servers or personal devices to handle applications. In cloud computing, the word cloud is used as a metaphor for "the Internet," so the phrase cloud computing means "a type of Internet-based computing," where different services such as storage and applications are delivered to an organization's computer systems or devices through the Internet[2]. With the advancement of technology power consumption has been become the crucial factor towards the growth of cloud computing. The data centers are the most significant part of the cloud computing infrastructure and require more attention to maintain its reliability, availability, scalability, and most importantly the power consumption of individual resources together with QoS. The data center may contain hundreds-to-thousands of servers and other equipments (including switches, routers, etc.). The topological arrangement of these resources can be in three or more layers i.e. 3-tier or n-tier depending upon the size of the datacenter. Each of the resource requires a sufficient amount of power to process the request and provide services to different users, and also for cooling purpose. All the network resources whether they are in idle state or in working state, consume some specific amount of energy due to the always running state of CPU's and other hardware part. To avoid unnecessary wastage of energy a simple strategy has been discussed in previous research work i.e. power down the unneeded links, switches, and servers. For this, we need some methodology to decide which subset of links, switches, and servers has to be active. The rest of the paper is organized as follows. Section II will provide the overview of datacenter and its working, the review of various proposed techniques to consolidate the flow will be discussed in Section III, in Section IV we will provide the detailed analysis of above techniques, and in last section we will conclude our study.

II. OVERVIEW

Wherever Times is specified, Times Roman or Times New Roman may be used. If neither is available on your word processor, please use the font closest in appearance to Times. Avoid using bit-mapped fonts if possible. True-Type 1 or Open Type fonts are preferred. Please embed symbol fonts, as well, for math, etc.

A. Data Center

This is a facility built for the purpose of housing cloud-based resources such as servers and other service-based equipment. Many cloud-based companies own and operate their own data centers which house the data stored for consumers and ensure the ongoing availability of their cloud. The companies may have multiple data centers in different geographic locations to ensure redundancies against data center failures, thus allowing an always-online service offering [4]. A data center (sometimes called a server farm) is a centralized repository for the storage, management, and dissemination of data and information. Typically, a data center is a facility used to house computer systems and associated components, such as telecommunications and storage systems. In April 2005, the Telecommunications Industry Association (TIA) produced specification TIA-942: Telecommunications Infrastructure for Datacenter. This was the first standard to specifically address datacenter infrastructure. This standard encompasses all parts of datacenter design, including cabling, facility, network design, and datacenter tiers. Datacenter is clusters of interconnected servers. They might all be very similar, stacked up neatly in open racks or closed cabinets of equal height, width and depth, or there could be a bunch of different types, sizes and ages of machines. Each server is a high performance computer, with storage space, input/output capability, a faster and more powerful processor, a lot more memory, and usually without a monitor, keyboard or the other peripherals. Monitors might exist in a centralized location, nearby or in a separate control room, for monitoring groups of servers and related equipment. A particular server or servers might be dedicated to a single task or running lots of different applications. Some servers in collocation datacenters are dedicated to particular clients. Some are even virtual rather than physical. When we request something via the Internet, then a number of servers are working together to deliver the content.

B. Datacenter Architecture

The data center is home to the computational power, storage, and applications necessary to support an enterprise business. The data center infrastructure is central to the IT architecture, from which all content is sourced or passes through. Proper planning of the data center infrastructure design is critical, and performance, resiliency, and scalability need to be carefully considered. Another important aspect of the data center design is flexibility in quickly deploying and supporting new services. Designing a flexible architecture that has the ability to support new applications in a short time frame can result in a significant competitive advantage. Such a design requires solid initial planning and thoughtful consideration in the areas of port density, access layer uplink bandwidth, true server capacity, and oversubscription, etc.

The data center network design is based on a proven layered approach, which has been tested and improved over the past several years in some of the largest data center implementations in the world. The layered approach is the basic foundation of the data center design that seeks to improve scalability, performance, flexibility, resiliency, and maintenance [8]. The basic layered architecture shown in figure. The layers of the data center design are the core, aggregation, and edge layers.

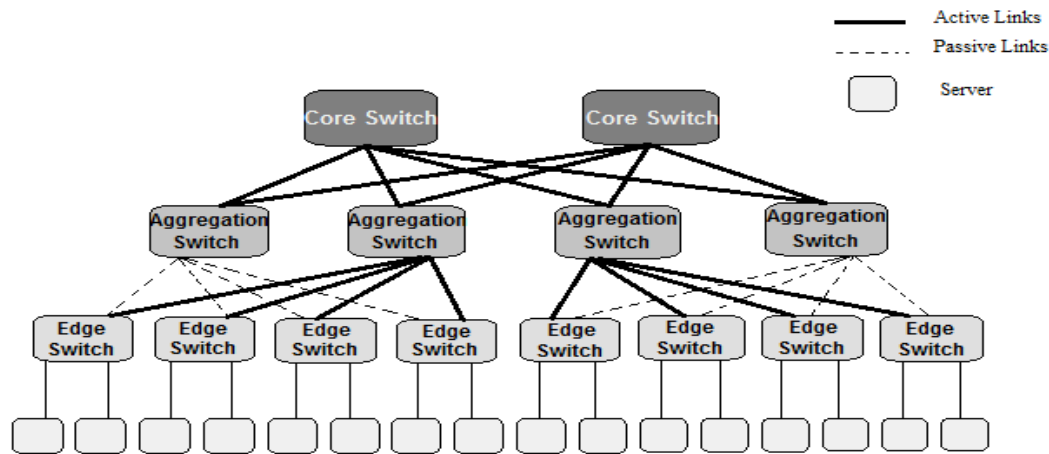


Fig. 1 A topological view of datacenter.

Core layer: Provides the high speed packet switching backplane for all flows going in and out of the data center. The core layer provides connectivity to multiple aggregation modules and provides a resilient Layer 3 routed fabric with no single point of failure. The core layer runs an interior routing protocol, such as OSPF or EIGRP, and load balances traffic between the core layer and aggregation layer using Forwarding based hashing algorithms.

Aggregation Layer: Provide important functions, such as service module integration, Layer 2 domain definitions, spanning tree processing, and default gateway redundancy. Server to server multitier traffic flows through the aggregation layer and can use services, such as firewall and server load balancing, to optimize and secure applications. The aggregation layer switch contains the integrated service modules. These modules provide services, such as content switching, firewall, SSL offload, intrusion detection, network analysis, and more.

Edge Layer: It is the layer where the servers physically attach to the network. The server components consist of 1RU servers, blade servers with integral switches, blade servers with pass through cabling, clustered servers, and mainframes with OSA adapters. The edge layer network infrastructure consists of modular switches, fixed configuration 1 or 2RU switches, and integral blade server switches. Switches provide both Layer 2 and Layer 3 topologies, fulfilling the various servers broadcast domain or administrative requirements.

III. STUDY OF VARIOUS PROPOSED PATH CONSOLIDATION TECHNIQUES IN DATA CENTER

The data center may contain hundreds-to-thousands of servers and other equipments (like links, switches, routers, etc.) and various techniques has been proposed to consolidate path by choosing subset of equipment among these that remain active to satisfy users demand.

A. Elastic Tree

Brandon Heller et. al. proposes an Elastic Tree [4] in which three methods are compared to choose the subset of network, i.e. formal model, Greedy bin packing, and topology aware heuristic.

Formal Model: The formal model provides a valuable tool for understanding the solution quality of other optimizers. The constraints include link capacity, flow conservation, and demand satisfaction. The variables are the flows along each link. The inputs include the topology, switch power model, and traffic matrix. To optimize for power, they added binary variables for every link and switch, and constrain traffic to only active (powered on) links and switches. The model also ensures that the full power cost for an Ethernet link is incurred when either side is transmitting; there is no such thing as a half-on Ethernet link. The optimization goal is to minimize the total network power, while satisfying all constraints. The model outputs a subset of the

original topology, plus the routes taken by each flow to satisfy the traffic matrix. It is flexible enough to support arbitrary topologies, but can only scale up to networks with less than 1000 nodes. This model focuses on data centers, not wide-area networks, chooses a subset of a fixed topology, not the component (switch) configurations in a topology, and considers individual flows, rather than aggregate traffic.

Greedy Bin-Packing: For even simple traffic patterns, the formal model's solution time scales to the 3.5th power as a function of the number of hosts. The greedy bin-packing heuristic improves on the formal model's scalability. Solutions within a bound of optimal are not guaranteed, but in practice, high-quality subsets result. For each flow, the greedy bin-packer evaluates possible paths and chooses the leftmost one with sufficient capacity. Within a layer, paths are chosen in a deterministic left-to-right order, as opposed to a random order, which would evenly spread flows. When all flows have been assigned (which is not guaranteed), the algorithm returns the active network subset (set of switches and links traversed by some flow) plus each flow path. For some traffic matrices, the greedy approach will not find a satisfying assignment for all flows, In this case, the greedy search will have enumerated all possible paths, and the flow will be assigned to the path with the lowest load. Like the formal model, this approach requires knowledge of the traffic matrix, but the solution can be computed incrementally, possibly to support on-line usage.

Topology-aware Heuristic: Unlike the other methods, it does not compute the set of flow routes, and assumes perfectly divisible flows. Of course, by splitting flows, it will pack every link to full utilization and reduce TCP bandwidth — not exactly practical. The simple additions to this “starter subset” lead to solutions of comparable quality to other methods, but computed with less information, and in a fraction of the time. The intuition behind this heuristic is that to satisfy traffic demands, an edge switch doesn't care which aggregation switches are active, but instead, how many are active. The “view” of every edge switch in a given pod is identical; all see the same number of aggregation switches above. The number of required switches in the aggregation layer is then equal to the number of links required to support the traffic of the most active source above or below (whichever is higher), assuming flows are perfectly divisible. For example, if the most active source sends 2 Gbps of traffic up to the aggregation layer and each link is 1 Gbps, then two aggregation layer switches must stay on to satisfy that demand. A similar observation holds between each pod and the core. But this computations assume a homogeneous fat tree with one link between every connected pair of switches. However, this technique applies to full-bisection-bandwidth topologies with any number of layers (we show only 3 stages), bundled links (parallel links connecting two switches), or varying speeds. Extra “switches at a given layer” computations must be added for topologies with more layers. Bundled links can be considered single faster links.

B. Merge Network

Candy Yiu and Suresh Singh, present the merge network [5] in which merging traffic from multiple links and feeding the merged stream to a switch with fewer ports. The traffic to/from N links are merged to K thus reducing the required port density of the switch from $N \times N$ to $K \times K$. As an example, if we merge 48 100Mbps links to 24 using a 48×24 merge network, we can replace the 48-port switch with a 24-port switch resulting in energy savings. If at most K packets arrive on the N uplinks (i.e., from the N links into the switch) then there are no packet losses and they are all forwarded to the K port switch, if more than K overlapping packets arrive then the earliest K get forwarded while the remaining are dropped. On the downlink (i.e., from the switch to the N links), the merge network needs to be able to forward packets from any of the K switch ports to any of the N downlinks and be able to forward up to K downlinks simultaneously.

An important requirement of the merge network is that it operates entirely in the analog domain. In other words, this network has no sense of ‘receives’ packets. In merge network, packets are dynamically switched to follow some path through the merge network. The reasons to build the merge network in this manner are threefold: this design ensures very small latency, the energy cost of the merge network is minimal and this design allows us to make the merge network relatively transparent to the PHY and MAC layer protocols.

C. CARPO

Xiaodong Wang et. al. [6] presents CARPO, a correlation power optimization algorithm, this algorithm based on the traffic correlation technique. CARPO takes three steps, Correlation Analysis, Traffic Consolidation and Link Rate Adaptation to optimize the power consumption of a DCN. In the first step, CARPO takes the data rates of the traffic flows in the previous consolidation periods as input and analyzes the correlation relationship between different traffic flows by using method in which if the traffic is negatively correlated, then the path is consolidated of the respective traffic flows, otherwise no need to consolidate the path. In the second step, based on the correlation coefficients from the previous analysis, CARPO uses the 90-percentile data rate of each link in the previous period to consolidate the traffics under the link capacity constraint. After the consolidation, unused switches and ports are turned off for power savings. In the last step, CARPO adapts the data rate of each active link to the demand of the consolidated traffic flows on that link, such that more power savings can be achieved for the DCN.

D. DENS

Dzmitry Kliazovich et. al. [7] presents a scheduling approach that combines energy efficiency and network awareness, named DENS and underlines the role of communication fabric in data center energy consumption. The DENS methodology balances the energy consumption of a data center, individual job performance, and traffic demands. The proposed approach optimizes the tradeoff between job consolidation (to minimize the amount of computing servers) and distribution of traffic patterns. In the proposed methodology, the network awareness is achieved with the introduction of feedback channels from the main network switches. This run-time feedback from the data center switches and links as well as decisions and actions are based on the network feedback.

IV. ANALYSIS AND DISCUSSION

Based on the above study we have analyzed that none of the technique discussed is highly efficient with respect to the end user Quality of Service and performance of the data center. Some of the negative and positive point we are going to discuss in the following table:

Sr. no.	Techniques	Positive Impact	Negative Impact
1.	Formal Model	Successfully choose the subset of network and the route taken by each flow to satisfy the traffic. It is flexible enough to support arbitrary topologies	It can only scale up to networks with less than 1000 nodes. For even simple traffic, time scales to the 3.5th power as a function of the number of hosts
2.	Greedy Bin-packing	It evaluates possible paths and chooses the leftmost one with sufficient capacity. The flow will be assigned to the path with the lowest load by splitting and without it	Requires complete knowledge of the traffic matrix. The no-split version scaled as about $O(n^{2.5})$, while the with-split version scaled slightly better, as $O(n^2)$
3.	Topology aware heuristic	It fares much better, scaling as roughly $O(n)$	It assumes fixed traffic flow, which is not always possible in real example.
4.	Merge network	Packets are dynamically switched to follow some path with very small latency and minimal cost.	Drop of packets, if more than the merge network can transmit, which may affect the performance of complete data center
5.	Correlation aware technique	Very energy efficient than others, adapts the data rate of each active link to the demand of the consolidated traffic flows	Based on assumption that different flows do not peak at exactly the same time, no mechanism to handle the congestion

V. CONCLUSION

On the basis of above study we have conclude that CARPO (Correlation aware power optimization) is the most efficient technique to optimize the power consumption as it provide the following features:

- » Correlation between different traffic patterns
- » Consolidation of the traffic which are negatively correlated

- » Adaptation of data rate for each active link to meet the demand of the consolidated traffic flows

Also we can improve its performance if we include the concept of network awareness discussed in [7], which may help to remove the congestion problem by providing feedback loop from main switches regarding the capacity of each link to handle the traffic.

ACKNOWLEDGMENT


I take this opportunity to express my deepest and sincere gratitude to most esteemed HOD of CSE Dept. Ms. Pratibha Sharma and all the faculty members of CSE Department, who have been kind enough to spare their valuable time, and inspire me to do better than I can. Their guidance and motivation conceived a direction in me.

Last but not the least I shall remain thankful to my parents and all my Friends, who were always appreciated my thoughts and encourage me whenever I was in low spirits.

References

1. IBM Cloud, "Computing as a service over the Internet", <http://www.ibm.com/cloud-computing/us/en/what-is-cloud-computing.html>
2. VangieBeal, "cloud computing (thecloud)" http://www.webopedia.com/TERM/C/cloud_computing.html.
3. Paul Stryer, "Understanding Data Centers and Cloud Computing", <http://www.globalknowledge.com/>
4. Brandon Heller, Srinu Seetharaman, Priya Mahadevan, Yiannis Yiakoumis, Puneet Sharma, Sujata Banerjee, and Nick McKeown, "ElasticTree: Saving Energy in Data Center Networks", In usenix, 2010.
5. Candy Yiu and Suresh Singh, "Merging Traffic to Save Energy in the Enterprise", In E-Energy, 2011.
6. Xiaodong Wang, Yanjun Yao, Xiaorui Wang, Kefa Lu, and Qing Cao, "CARPO: Correlation-Aware Power Optimization in Data Center Networks", In IEEE, 2012.
7. Dzmityr Kliazovich · Pascal Bouvry · Samee Ullah Khan, "DENS: data center energy-efficient network-aware scheduling", In Springer Science, Business Media, August 2011.
8. Cisco, "Data Center Architecture Overview" http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DC_Infra2_5/DCInfra_1.html
9. Andreas Berl1, Erol Gelenbe, Marco di Girolamo, Giovanni Giuliani, Hermann de Meer, Minh Quan Dang and Kostas Pentikousis, "Energy-Efficient Cloud Computing", In The Computer Journal, Vol. 53 No. 7, 2010.
10. Qing Yi and Suresh Singh, "Minimizing Energy Consumption of FatTree Data Center", In sigmetrics, 2014.
11. Srikanth Kandula, Sudipta Sengupta, Albert Greenberg, Parveen Patel, Ronnie Chaiken, "The Nature of Datacenter Traffic: Measurements & Analysis", In *IMC*, 2009.
12. Shekhar Srikantaiah, Aman Kansal, and Feng Zhao, "Energy Aware Consolidation for Cloud Computing", In usenix 2008.
13. Dennis Abts, Michael R. Marty, Philip M. Wells, Peter Klausler, and Hong Liu, "Energy Proportional Datacenter Networks", In ISCA, June 2010.
14. Qi Zhang · Lu Cheng · Raouf Boutaba, "Cloud computing: state-of-the-art and research challenges", In Springer, J Internet Serv Appl, 2010.

AUTHOR(S) PROFILE



Shilpa Kumari, received the B.Tech degree in Information Technology from Himachal Pradesh University during 2008-12 and pursuing M.Tech in Computer Science and Engineering (Specialization in Mobile Computing) from Career Point University Hamirpur during 2013-15.