Accepted Manuscript

Towards thermal-aware Hadoop clusters

Yi Zhou, Shubbhi Taneja, Gautam Dudeja, Xiao Qin, Jifu Zhang, Minghua Jiang, Mohammed I. Alghamdi

PII: S0167-739X(17)30118-8
DOI: https://doi.org/10.1016/j.future.2018.04.084
Reference: FUTURE 4162

To appear in: *Future Generation Computer Systems*

Received date : 19 January 2017
Revised date : 16 December 2017
Accepted date : 25 April 2018

Towards Thermal-Aware Hadoop Clusters

Yi Zhou, Shubbhi Taneja,Gautam Dudeja, Xiao Qin, *Senior Member, IEEE*,
Jifu Zhang,Minghua Jiang and Mohammed I. Alghamdi

Abstract—With the exponential increase in cooling costs of large-scale data centers, thermal management must be adequately addressed. Recent trends have discovered one of the critical reasons behind the temperature rise turns out to be heat re-circulation within data center. In this study, we proposed a new resource- and thermal-aware scheduler in Hadoop clusters; our scheduler aims at minimizing peak inlet temperature across all nodes to reduce power consumption and cooling cost in data centers. The proposed dynamic scheduler makes job scheduling decisions based on current CPU/disk utilization and number of tasks running as well as the feedback given by all slave nodes at run-time. We deploy a thermal model to project respective temperature of each slave node in addition to neighbor's heat contribution. The thermal-aware scheduler is integrated with the Hadoop's scheduling mechanism. We test our schedulers by running a set of Hadoop benchmarks (e.g., WordCount, DistributedGrep, PI and TeraSort) under various temperature conditions, utilization thresholds, and cluster sizes. The experimental results show that our scheduler achieves an average inlet temperature reduction by 2.5C over the default FIFO scheduler; our scheduling solution saves approximately 15% of cooling cost with marginal performance degradation.

Keywords—Data Centers, Thermal Energy, Hadoop, Map-Reduce

1 Introduction

Data centers have been one of the fastest growing power consumption unit in recent years. With the rapid increase usage of distributed processing frameworks (e.g., Yahoos's Hadoop), the energy cost of maintaining datacenters has become a vital economically factor [1]. Energy cost spent on cooling sysetems turns into a huge fraction of a data center's maintaince cost. Unsprisingly, the cooling cost nowadays is much higher than the cost of building data centers housing IT equipments [2]. A handful of studies were conducted to reduce the energy cost in light of improving the server operational efficiency [3]. However, improving the energy efficiency of cooling systems is still an open issue. Few techniques were proposed to cut energy waste in order to boost the energy efficiency of cooling systems. Employing proper energy-efficient techniques, one may conserve 26% [4] of energy consumption without dramatically affect system performance.

Improving thermal efficiency of computing clusters is a popular way of reducing cooling cost of data centers [5][6]. The endless increasing demand of data-intensive computing leads to hardware capacity at a large scale (e.g., fast CPUs coupled with large storage systems), which in turn gives rise to excessive heat. A hot computing environment not only leads to high cooling cost, but also increases cooling-failure risks [7]. In this study, we show that that lowering overall node temperatures helps in minimizing the energy consumption of cooling systems in the data centers.

- *Y. Zhou, S. Taneja*, G. Dudeja, and X. Qin are with the Department of Computer Science and Software Engineering, Samuel Ginn College of Engineering, Auburn University, AL 36849-5347. E-Mail: {yzz0074, shubbhi, gzd0004, xqin}@auburn.edu, ,*corresponding author: shubbhi@auburn.edu.*
- *J.-F. Zhang is with the School of Computer Science and Technology, Taiyuan University of Science and Technology, Taiyuan 030024, China. E-mail: zjf@tyust.edu.cn.*
- *M.-H. Jiang is with the College of Mathematics and Computer Science, Wuhan Textile University, Wuhan 430073, China E-mail: jmh@wtu.edu.cn*
- *M.-I. Alghamdi is with the Department of Computer Science, Al-Baha University, Al-Baha City, Kingdom of Saudi Arabia E-mail: mialmushilah@bu.edu.sa*

1.1 Motivations

Three reasons motivate us work towards thermal-aware Hadoop clusters:

- Endless increasing energy cost of cooling systems for temperature maintenance in data centers.
- The lack of thermal-aware Hadoop schedulers that consider heat re-circulation, CPU utilization, and disk utilization.
- The pressing need of thermal models applied to estimate thermal impact and manage task assignments in order to improve the energy efficiency of cooling system.

Energy consumption has grown to be an enormous concern in maintenaning data centers [8]. The energy consumption of information technology (IT) equipment (e.g.,the computing servers, network cables, power supplies, security devices, thermostats, AC ducts) is approximately 40% of total energy consumption. Energy consumed by cooling systems accounts for approximately 45% in data centers [9]. Coupled with the fast growth of data centers, thousands of server racks are fitted in to floor areas as large as 9000 m^2. Data centers with high density of server racks augment the energy consumption incurred by the heat dissipation effect [10]. For example, due to the heat re-circulation, a rack server's heat dissipation requirements is roughly 30 times larger that those of servers in 1990 [11]. In 2013, U.S. data centers consumed an estimated 91 billion kilowatt-hours of electricity, which is the amount of electricity to power all entire households in the New York City twice over for a year. Meanwhile, power consumption of data centers is projected to increase to roughly 140 billion kilowatt-hours annually by 2020, costing American businesses $13 billion annually in electricity bills.

Energy conservation techniques at the application-level recvieved much attention. For example, Li and Zhang proposed a power-aware MapReduce application model tailored for energy-efficient MapReduce programming [12]. Most existing energy consumption models ignore the corelations between heat re-circulation and hardware workload (e.g., CPU utilization and disk utilization) [13]. In this study, we aim to improve energy efficiency by seamlessly integrating multiple factors