

Recent Advances in Energy Efficient Query Processing

Matteo Catena and Nicola Tonellotto

National Research Council of Italy, Pisa, Italy,
`name.surname@isti.cnr.it`

Abstract. Web search companies distribute their infrastructures and operations across several, geographically distant data centers. This distributed architecture facilitates high performance query processing, which is fundamental for the success of a Web search engine. At the same time, data centers require an huge amount of electricity to operate their computing resources. In this extended abstract, we briefly discuss our recent works for improving the energy efficiency of query processing systems. Firstly, we introduce a novel query forwarding algorithm which exploits green energy sources to reduce the electricity expenditure and carbon footprint of Web search engines. Then, we propose to delegate the CPU power management from a server' operative system directly to the query processing application, to reduce the energy consumption of a search engine's servers. Finally, we introduce PESOS, a scheduling algorithm which manages the CPU power consumption on a per-query basis while considering query latency constraints.

High performance query processing is fundamental for the success of a Web search engine. In fact, Web search engine can receive billions of queries per day [5]. Additionally, the issuing users are often impatient and expect sub-second response times to their queries (e.g., 500 ms). Indeed, users become less engaged [1] or migrate to other search services [11] when a search engine fails to provide fast responses to queries. For such reasons, search companies adopt distributed query processing strategies to cope with huge volumes of incoming queries and to provide sub-second response times.

Web search engines perform distributed query processing on computer clusters composed by thousands of computers and hosted in large data centers [5]. While such facilities enable large-scale online services, they also raise economical and environmental concerns. Indeed, a large-scale data center – like those used by Web search engines – can draw tens of megawatts of electricity to operate and it can cost 9 million US dollars per year in terms of energy expenditure [8]. Therefore, an important problem to address is *how to reduce the energy expenditure of data centers*. Additionally, producing and consuming electricity can involve the emission of carbon dioxide, which is the main cause of global warming due to the greenhouse effect. In 2007, the Information and Communication Technology (ICT) sector has been reported to be responsible for roughly 2% of global carbon emissions, with general purpose data centers accounting for 14% of

the ICT footprint. These emission levels were projected to more than double by 2020 [13]. Therefore, another problem to tackle is *how to reduce these emissions and the negative impact of the data centers on the environment*.

Obviously, a possible solution to these challenges consists in designing more energy-efficient data centers, which consume less energy and, consequently, pollute and cost less. In the past, a large part of the energy consumption of a data center could be accounted to inefficiencies in its cooling and power supply systems. However, search companies already adopt state-of-the-art techniques to reduce the energy wastage of such systems^{1,2}, leaving little room for more improvements in those areas. Indeed, the energy consumption of a state-of-the-art data center would be reduced by less than 24% if all the overheads in its cooling and power supply systems were eliminated [2]. Therefore, new approaches are necessary to mitigate the environmental impact and the energy expenditure of Web search engines.

One option consists in using *green energy*. In fact, several search companies use green energy to partially power their data centers, i.e., energy which comes from resources that are renewable and do not emit carbon dioxide, such as sunlight and wind^{3,4,5}. At the same time, Web search engines experience spatial and temporal variations in electricity prices [10] as they distribute their infrastructures and operations across several, geographically distant data centers [5]. Stemming from these observations, we propose a novel query forwarding algorithm that exploits both the green energy sources available at different data centers and the differences in market energy prices [3]. The main idea is to dispatch queries from the data center that firstly received the requests to a different one, if the latter can rely on green energy or cheaper energy sources than the former. The problem of exploiting different energy sources to reduce costs when forwarding queries is modeled as a Minimum Cost Flow Problem. The model takes into account the different and limited processing capacities of data centers, query response time constraints and communication latencies among sites. We evaluate the proposed algorithm using workloads obtained from the Yahoo search engine, together with realistic electricity price data. Our experimental results show that the proposed solution maintains an high query throughput, while reducing by up to 25% the energy operational costs of multi-center search engines. Moreover, our algorithm can reduce by almost 6% the consumption of non-green energy.

The energy expenditure and carbon footprint of a search company can also be mitigated by reducing the energy consumption of its computing resources. In particular, reducing the energy consumption of CPUs represents an attractive venue for Web search engines. In fact, CPUs are the most energy consuming component in servers dedicated to query processing, accounting for 40% of total

¹ <https://www.google.com/about/datacenters/efficiency/internal/>

² <https://www.microsoft.com/about/csr/downloadhandler.ashx?Id=02-01-12>

³ <https://environment.google/>

⁴ <https://www.microsoft.com/about/csr/environment/>

⁵ <https://sustainability.fb.com>

energy consumption when a server is idle and for 66% of total energy consumption when it is fully utilized [2]. Dynamic Voltage and Frequency Scaling (DVFS) technologies can be used to reduce the CPU energy consumption of a server [12]. DVFS permits to adjust the frequency and voltage at which the CPU cores operate, trading off performance for power consumption. In fact, higher core frequencies mean faster computations but higher power consumption, while lower frequencies lead to slower computations but reduced power consumption. However, carefulness is required when reducing the operating frequency of the CPU cores since low frequencies entail longer query processing times that may be unacceptable for the users.

Typically, DVFS mechanisms are managed by operating system (OS) components, called *frequency governors* [4, 14]. However, the OS misses domain-specific information regarding the utilization and load of the query processing application. This knowledge can be exploited to better throttle the frequency of the CPU cores, thereby reducing the power consumption of a query processing server. Therefore, in [6] we propose to delegate the CPU power management from the OS frequency governors to the query processing application, and we devise search engine-specific frequency governors. We experimentally evaluate such governors upon the TREC ClueWeb09B corpus and the query stream from the MSN 2006 query log. Results show that the knowledge of the query processing server utilization and load facilitates a more refined control of the CPU to achieve power savings. In fact, the proposed search engine-specific governors can reduce up to 24% a server power consumption, with only limited (but uncontrollable) drawbacks in the quality of search results with respect to a system operating at maximum CPU frequency.

Another important aspect that can be exploited to reduce the energy consumption of a server is the fact that users can hardly notice response times that are faster than their expectations [1, 11]. Therefore, we advise that Web search engines should not process queries faster than user expectations and, consequently, we propose the Predictive Energy Saving Online Scheduling (PESOS) algorithm [7]. PESOS selects the most appropriate CPU frequency to process a query by its deadline, on a per-core basis. It considers the latency requirement of queries as an explicit parameter, and it tries to process queries no faster than required. In doing so, the CPU energy consumption is reduced while respecting the query latency constraints. PESOS bases its decision on query efficiency predictors, which are techniques to estimate the processing volume and processing time of a query before its execution [9]. We experimentally evaluate PESOS upon the TREC ClueWeb09B collection and the MSN 2006 query log. Depending on the required latency, results show that PESOS can reduce the CPU energy consumption of a query processing server from 24% up to 48% when compared to an high performance system running at maximum CPU core frequency. Also, PESOS outperforms our best search engine-specific frequency governor [6] with a 20% energy saving, while the competitor requires a fine parameter tuning and it may incurs in uncontrollable latency violations.

References

1. Arapakis, I., Bai, X., Cambazoglu, B.B.: Impact of Response Latency on User Behavior in Web Search. In: ACM (ed.) Proc. SIGIR. pp. 103–112. Gold Coast, Queensland, Australia (2014)
2. Barroso, L.A., Clidaras, J., Hölzle, U.: The Datacenter as a Computer: an Introduction to the Design of Warehouse-scale Machines. *Synthesis lectures on computer architecture* 8(3), 1–154 (2013)
3. Blanco, R., Catena, M., Tonellotto, N.: Exploiting Green Energy to Reduce the Operational Costs of Multi-Center Web Search Engines. In: IW3C2 (ed.) Proc. WWW. pp. 1237–1247. Montreal, Quebec, Canada (2016)
4. Brodowski, D.: CPU frequency and voltage scaling code in the Linux kernel. <https://www.kernel.org/doc/Documentation/cpu-freq/index.txt> (2015), last visited: 2016-11-08
5. Cambazoglu, B.B., Baeza-Yates, R.: Scalability Challenges in Web Search Engines. *Synthesis Lectures on Information Concept, Retrieval, and Services* 7(6), 1–138 (2015)
6. Catena, M., Macdonald, C., Tonellotto, N.: Load-sensitive CPU Power Management for Web Search Engines. In: ACM (ed.) Proc. SIGIR. pp. 751–754. Santiago, Chile (2015)
7. Catena, M., Tonellotto, N.: Energy-efficient Query Processing in Web Search Engines. *Transactions on Knowledge and Data Engineering* (2017)
8. Greenberg, A., Hamilton, J., Maltz, D.A., Patel, P.: The Cost of a Cloud: Research Problems in Data Center Networks. *SIGCOMM Computer Communication Review* 39(1), 68–73 (2008)
9. Macdonald, C., Tonellotto, N., Ounis, I.: Learning to Predict Response Times for Online Query Scheduling. In: ACM (ed.) Proc. SIGIR. pp. 621–630. Portland, Oregon, USA (2012)
10. Qureshi, A., Weber, R., Balakrishnan, H., Gutttag, J., Maggs, B.: Cutting the Electric Bill for Internet-scale Systems. In: ACM (ed.) Proc. SIGCOMM. pp. 123–134. Barcelona, Spain (2009)
11. Schurman, E., Brutlag, J.: Performance Related Changes and their User Impact. In: O’Reilly (ed.) Proc. Velocity. San Jose, USA (2009)
12. Snowdon, D.C., Ruocco, S., Heiser, G.: Power Management and Dynamic Voltage Scaling: Myths and Facts. In: Proc. PARC workshop at EMSOFT. IEEE (2005)
13. The Climate Group for the Global e-Sustainability Initiative: Smart 2020: Enabling the low carbon economy in the information age. <http://gesi.org/files/Reports/Smart%202020%20report%20in%20English.pdf> (2008), last visited: 2016-11-04
14. The Linux Kernel Archives: Intel P-State driver. <https://www.kernel.org/doc/Documentation/cpu-freq/intel-pstate.txt> (2016), last visited: 2016-11-08