Fuel Cell Powered Datacenter Power Management

Yang Li

1. Problem Statement

Data center energy consumption across the world have been continuously growing in the recent years [1]. For example, the energy consumed by Google's data centers has increased by almost 20 times in the past 10 years [2]. The growth of data center energy consumption not only increases the total cost of ownership (TCO), but also leads to more serious environmental issues. It is reported that by 2020, the greenhouse gas emission of computing systems will account for 3% of the world's total emission, which makes data center one of the biggest greenhouse gas emitters in the world. Therefore, it is desirable to invent new technologies to improve the energy efficiency of data center and reduce its greenhouse gas emission.

To meet this demand, fuel cell technology has been proposed to power data center [8, 9, 4, 3]. Fuel cell is a power generator which converts fuel (e.g., hydrogen and natural gas) to electricity through an electrochemical process. It has several advantages as power sources for data center. First, fuel cells have high energy efficiency [9]. Fuel cells convert chemical energy directly into electrical energy, and therefore their efficiency is not limited by Carnot Cycle Efficiency which limits the efficiency of conventional combustion power generator. In addition, as a distributed power generator, fuel cells can be placed close to power consumers and thus avoid the energy losses in the long-distance power transmission. Second, fuel cells induce less environmental issues [9]. For example, fuel cells can reduce 49% carbon dioxide emissions of traditional power plants. Third, fuel cells have high energy source reliability [9]. This is because fuel delivery infrastructure is usually buried and thus robust to severe weather. Based on these advantages, fuel cells show promises as future data center power sources.

However, in order to use fuel cells to power data center, the challenge arising from the limited load following capability of fuel cells needs to be resolved. Fuel cells are typically designed to provide relatively steady power. When their loads change, the controller of fuel cells will sense the load changes and adjust the power output of fuel cells. However, due to the fuel cells' slow fuel processing / delivery process, their output power cannot immediately meet the new power demand and needs to take a certain period of time to fully follow the loads. For example, Proton Exchange Membrane Fuel Cells (PEMFC) may take several seconds to meet the load changes [8]. Therefore, a gap between the power demand of servers and power supply of fuel cells will occur and need to be handled to facilitate the normal operation of servers.

To tackle this challenge, previous proposals use energy storage devices to handle the gap [4, 3, 8, 9]. Energy storage devices such as batteries and super-capacitors can respond to load changes rapidly and provide instant power to servers when the fuel cells' power supply cannot meet the power demand. These energy storage devices are generally expensive and account for a significant portion of data center TCO [7]. However, none of these prior proposals have tackled the sizing problem of energy storage devices to reduce the TCO under the condition that the power gap can still be handled.

In this project, we plan to provide design guidelines on the sizing of energy storage devices for fuel cell powered datacenter. We will firstly model the power dynamics of fuel cells, energy storage devices and servers using well-studied system identification techniques. And then, based on the model, we will explore the sizing problem of energy storage devices to handle the worst-case load demand variation. After that, we will explore another knob to handle the gap, power capping, to shape the power demand of servers. We plan to propose a distributed optimal power capping policy for servers to reduce the gap between the power supply and demand, alleviating the reliance on energy storage devices. We want to combine the proposed power capping policy with energy storage devices, and explore the trade-off between the TCO of these energy storage devices with the applications' performance degradation under both normal and worst cases.

The system configuration in our project is shown as Figure 1. In this configuration, a fuel cell stack is used to power a rack of servers. Using fuel cells in rack level other than data center level shows the benefits of energy efficiency [9]. Different from previous works [4, 3, 8, 9] which use batteries as local energy storage, we use super-capacitors in our project to reduce the cost. Super-capacitors can bear several orders of more charge / discharge cycles than batteries [5, 7], and have simpler auxiliary circuits. As a result, compared with batteries, super-capacitors can be used for much longer periods and achieve much smaller costs [5].

2. Novelty

The essential goal and contribution of this project is to provide design guidelines for the energy storage device (super-capacitor) sizing of fuel cell powered data center. To do this, we will evaluate the size of energy storage devices to handle the worst-case power variation, and then combine the proposed distributed power capping policy with energy storage device sizing to explore the trade-off between super-capacitor size and applications' performance degradation under both normal and worst cases. The most closely related prior works are proposals on data center power demand shaping and energy

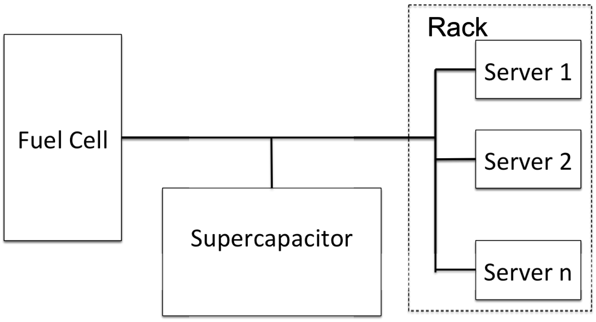

Figure 1: System Configuration for Fuel Cell Powered Data Center

storage sizing.

Li et al. [4] proposed to coordinate the power demand shaping with fuel cell output specification, and adjust the output power specification of fuel cell based on the recent applications' power demand and performance. This work is different from our project in the following aspects. First, no energy storage device sizing mechanism is proposed in this work. Second, this work uses batteries other than super-capacitors as energy storage devices, which may increase the TCO of data center. Third, the mechanism in this work adjusts the output power specification of fuel cell by an interval-based approach (the interval length is chosen such that the fuel cell output power can meet the specification by the end of interval), and increases the power specification by a fixed amount (10%) in each interval when it is necessary. Therefore, when a dramatic workload surge occur, the fuel cell may spend multiple intervals to reach the power supply under which applications can run at full speed. As a result, applications may suffer from severe performance degradation during this slow converge process. While in our power capping mechanism, we plan to leverage the techniques of optimal control to optimally apply DVFS to servers so that the applications suffer the least performance degradation to handle the power gap caused by the workload surge.

In addition, Li et al. extended the mechanism in [4] to use both fuel cell and intermittent power sources (e.g., wind turbine) to power data center [3]. This work focuses on the power variability issue of intermittent power and is complementary to our work. Zhao et al. [9, 8] and Riekstin et al. [6] evaluate the load following capability of fuel cells and the TCO of using fuel cells in rack level. None of these works consider the energy storage device sizing problem.

Want et al. [7] proposed a energy storage device sizing framework for utility grid powered data center. Though both this work and our project explore the energy storage device sizing problem, our context is different. In [7], the power provisioning budget is fixed. While in our project, the power provisioning budget is variable, and will change based on the interaction with load (due to the load following capability of fuel cells). In [7], the size of local energy storage device is determined based on the trade-off between device cost and the cost of the power drawn exceeding the budget. In this sense, the total power demand of servers and energy storage devices may exceed the provisioning budget. While in our project, The energy storage devices are used to strictly guarantee the total power demand within the budge.

3. Idea

We plan to leverage the optimal control technique (specifically, model predictive control) to design the power capping policy. The goal of the policy is to maximize the performance of all the servers in a rack. The constraint of the policy is that the bus voltage stays within the normal operating voltage range of the servers. At the beginning of each time interval, each server uses its current state (e.g., workload request arrival rate and bus voltage) to determine its optimal power capping actions (i.e., CPU frequency) in the next N steps based on the application's performance model and power supply system model. And then, the server will only implement the first step action in current interval and iterate this process at the beginning of future intervals. We want the power capping policy to be distributed, which means each server only needs to monitor its own state and does not need to know other server's state. It is easy to design the distributed policy when the load is balanced across all the servers in the rack and each server runs the same application, since under such scenarios knowing the state of the server itself is equal to knowing the state of others. However, it is challenging to design the distributed policy under other situations. To do this, we plan to gain insights from collaborative control techniques.

4. Hypothesis

The main hypothesis we will test is that the size of supercapacitors can be significantly reduced with negligible performance degradation due to power capping.

5. Evaluation Methodology

There are two potential approaches to construct the simulator. We need to compare them based on the feasibility and modeling accuracy.

5.1. Simulator Option 1

We can model the server performance and power consumption based on the data collected from typical workloads. I have collected the system throughput, response time and power consumption of Bing Search under different CPU frequencies and request arrival rates. We can use this data to construct a lookup table which maps the input (CPU frequency and request arrival rate) to the output (system throughput, response time and power consumption). This lookup table will be embedded in the simulator to model the performance and power based on the CPU frequency (from the power capping controller) and the request arrival rate (from the trace). For this simulator, we need traces which record the request arrival rate across time for a rack (when assuming load balance within the rack) or each server in a rack.

5.2. Simulator Option 2

We can build a simulator as a queuing system using the methodology in [4, 3]. The performance is modeled as the throughput and response time of queuing system. The power consumption is modeled based on the CPU frequency and CPU utilization. The trace for the simulator record the information of each request (request arrival time, request service time at highest CPU frequency) across time for a rack (when assuming load balance within the rack) or each server in a rack.

6. Plan

6.1. Steps

6.1.1. Milestone 1 Build the models for whole power supply systems. Evaluate the super-capacitor size which can tackle the worst-case power variation.

6.1.2. Milestone 2 Design a distributed power capping approach. Set down the evaluation methodology, build the simulator and collect the traces.

6.1.3. Report Evaluate the trade-off between super-capacitor size and power degradation due to power capping under both normal and worst cases. Compare the proposed technique with prior proposals. Write the report.

6.2. Goals

6.2.1. 75% Complete the previous three steps. Find that power capping will reduce super-capacitor size with relatively reasonable performance degradation.

6.2.2. 100% Find that power capping will reduce super-capacitor size with reasonable or negligible performance degradation. Find a significant improvement over counterpart mechanism. Publish it in a decent conference.

References

[1] “DCD industry census 2013: Forecasting energy demand,” 2013.

[2] J. Koomey, “Growth in data center electricity use 2005 to 2010,” *A report by Analytical Press, completed at the request of The New York Times*, 2011.

[3] C. Li, R. Wang, T. Li, D. Qian, and J. Yuan, “Managing green datacenters powered by hybrid renewable energy systems,” in *11th International Conference on Autonomic Computing (ICAC 14)*, 2014.

[4] C. Li, R. Zhou, and T. Li, “Enabling distributed generation powered sustainable high-performance data center,” in *Proceedings of the IEEE 19th International Symposium on High Performance Computer Architecture (HPCA)*, 2013.

[5] S. McCluer and J.-F. Christin, “Comparing data center batteries, flywheels, and ultracapacitors,” *Schneider Electric White Paper*, vol. 65, 2008.

[6] A. C. Riekstin, S. James, A. Kansal, J. Liu, and E. Peterson, “No more electrical infrastructure: Towards fuel cell powered data centers,” in *ACM Workshop on Power-Aware Computing and Systems*, November 2013.

[7] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, “Energy storage in datacenters: What, where, and how much?” *SIGMETRICS Perform. Eval. Rev.*, Jun. 2012.

[8] L. Zhao, J. Brouwer, S. James, E. Peterson, J. Siegler, A. Kansal, and J. Liu, “Servers powered by a 10kw in-rack proton exchange membrane fuel cell system,” in *Proceedings of the ASME 8th International Conference on Energy Sustainability and 12th Fuel Cell Science, Engineering and Technology Conference*, June 2014.

[9] L. Zhao, J. Brouwer, J. Liu, S. James, J. Siegler, A. Kansal, and E. Peterson, “Fuel cells for data centers: Power generation inches from the server,” Tech. Rep. MSR-TR-2014-37, March 2014. [Online]. Available: http://research.microsoft.com/apps/pubs/default.aspx?id=210641