Why Blow Away Heat? Harvest Server's Heat Using Thermoelectric Generators

Ted Tsung-Te Lai, Wei-Ju Chen, Yi-Hsuan Hsieh, Kuei-Han Li, Ya-Yunn Su, Polly Huang, Hao-Hua Chu
{tedlai, r99922148,r00922022, r98922022, yysu}@csie.ntu.edu.tw, phuang@cc.ee.ntu.edu.tw, hchu@csie.ntu.edu.tw
National Taiwan University

ABSTRACT

This paper argues for harvesting energy from servers' wasted heat in data centers. Our approach is to distribute a large number of thermoelectric generators (TEGs) on or nearby server hotspot components whose surface temperature is high enough for electricity regeneration. This paper answers the following questions. (1) Which server hotspot components are hot enough for energy harvesting by TEGs? (2) How much energy can be harvested from these selected server hotspot components? We further propose an energy-harvest aware scheduler that optimizes TEGs' energy harvesting efficiencies while preventing overheating of server components.

Categories and Subject Descriptors

C.5.5 **[Computer System Implementation]**: Servers

General Terms

Management, Measurement, Performance.

Keywords

Keywords are your own designated keywords.

1. INTRODUCTION

A data center consumes vast amount of electricity and produces enormous amount of wasted heat that needs to be removed by cooling facilities. This paper looks at wasted heat as opportunities for energy harvesting. Our approach is to deploy and distribute a large number of small thermoelectric generators (TEGs) on or nearby server hotspot components and turn their wasted heat back into electrical energy. TEGs are devices made of bismuth telluride material that can convert heat into electrical energy based on the Seebeck effect.

Several recent research projects have proposed various ways of exploiting wasted heat produced by cloud data centers. For examples, Liu *et. al* [1] introduces the concept of data furnace in which servers are used both as space heaters in buildings and IT infrastructure. Similarly, the Finish Uspenski data center [2] is planning to use waste heat produced by servers for heating and hot water requirement of approximately 2,000 houses nearby the data center.

This paper argues that it is possible to extract energy from selected server hotspot components whose surface temperature can rise and sustain high enough (i.e., around 70^{o}C) for electricity regeneration using TEGs. Note that the hotspot surface temperature of these server components is much higher than the temperature of the servers' fan-exhaust air around 40^{o}C. To show the energy harvesting potentials of the proposed approach, we have performed experiments to answer the following questions.

- What are the server hotspot components whose surface temperature can rise and sustain high enough for energy harvesting by TEGs?
- How much energy can be harvested from these selected server hotspot components?

We further propose an energy-harvesting workload scheduler to optimize TEG's energy harvesting efficiencies while preventing overheating in these server hotspot components.

2. SERVER HOTSPOT COMPONENTS

To identify server hotspot components with energy harvesting potentials, we performed experiments on a Dell PowerEdge R310 1U rack server. A thermal image was first taken to show the hotspots inside the server (See figure 1). We then installed an infrared thermometer to measure surface temperature on selected server components, including CPU, memory chips, the hard disk, the SAS (Serial attached SCSI) controller IC, the graphic chip and the network chip. To show the relationship between the temperature and the utilization of each server component, we ran the Phoronix Test Suite benchmarks to keep these server components busy and observed their surface temperature. Table 1 summarizes each server component's temperature at the idle and busy state.

Table 1 shows that CPU can reach high temperature, thus ideal for TEG energy harvesting. We further performed an experiment to measure the amount of energy TEG can harvest from the CPU. Figure 2 plots the CPU utilization, the CPU temperature and the TEG harvested energy over time. When running 100% CPU utilization, the CPU temperature gradually climbed from 40^{o}C to 90^{o}C and increased the amount of energy harvested by the TEG to 0.2W to 0.3W. In other words, the amount of harvested energy is proportional to the CPU utilization and CPU temperature.

Conference'12, Month 1–2, 2010, City, State, Country.

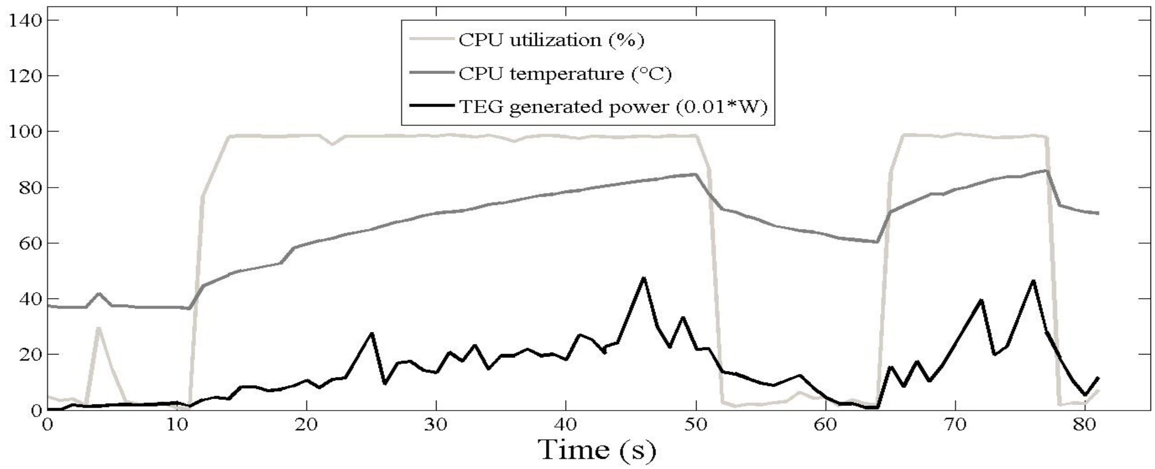

Figure 2. The relationship between CPU utilization, CPU temperature and TEG generated power. Note that the TEG power is scale 100x for better visualization in this figure.

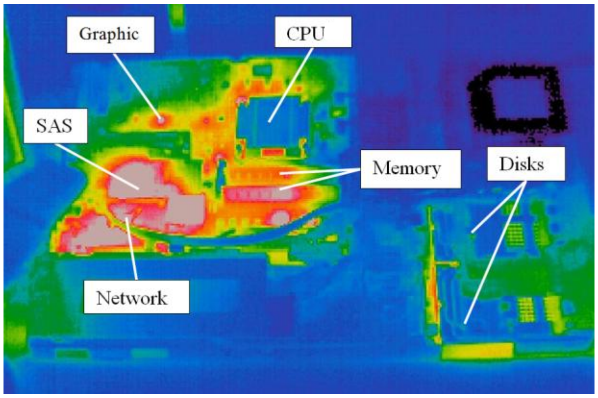

Figure 1. A thermal image inside the Dell R310 server. Note that the surface temperature of CPU cannot be captured because a hot sink is placed on top of it.

3. ENERGY-HARVEST AWARE WORKLOAD SCHEDULER

The goal of energy-harvest aware workload scheduler is to maximize the amount of harvested energy while maintaining the desired performance and preventing hotspots on any server components. The scheduler takes as input the model of TEG energy generation, CPU temperature, and thermal constraints. It periodically schedules tasks to a specific CPU and utilizes DVFS options. Previous work [3, 4] proposed dynamically scheduling tasks at operating system level to prevent hotspot and improve server thermal management. Here we walk a thin line of achieving similar thermal management goals while carefully allowing hotspots on some components to improve energy harvesting efficiency. We plan to explore formulating this scheduling into a constraint optimization problem or apply control techniques.

Table 1. Hotspot components temperature under idle state and busy state (i.e. running benchmark)

Components	Temp (℃) at idle	Temp (℃) at busy
CPU	30	70
Memory chips	24.5	42
Disk	24	30
SAS controller chip	42	50
Network chip	37	41
Graphics chip	28	32

4. FUTURE WORK

We are looking forward to implement and evaluate this energy-harvesting system.

5. REFERENCES

[1] Liu, J., Goraczko, M., James, S., Belady, C., Lu, J., Whitehouse, K. 2011. The data furnace: heating up with cloud computing. In Proceedings of the 3rd USENIX Workshop on Hot Topics in Cloud Computing 2011

[2] New eco-efficient data center for Suvilahti, Helsinki, http://www.helen.fi/tiedotteet/tiedotteet.aspx?id=306&lang=en

[3] Merkel, A. and Bellosa F. Task activity vectors: a new metric for temperature-aware scheduling. In Proceedings of EuroSys 2008

[4] Yang, J., Zhou, Z., Chrobak, M., Zhang, Y, and Jin, L. Dynamic thermal management through task scheduling. In Proceedings of ISPASS 2008