

# A Neuro-Fuzzy Fan Speed Controller for Dynamic Thermal Management of Multi-core Processors

Javad Mohebbi Najm Abad §, Bagher Salami †, Hamid Noori †, Ali Soleimani §, Farhad Mehdipour ‡

§ Shahrood University of Technology  
Shahrood, IRAN  
Tel: +98-0273-339-2204  
javad.mohebi@shahroodut.ac.ir  
solimani\_ali@shahroodut.ac.ir

† Ferdowsi University of Mashhad  
Azadi Sq., Mashhad, IRAN  
Tel: +98-511-880-2000  
bagher.salami@stu-mail.um.ac.ir  
hnoori@um.ac.ir

‡ E-JUST Center, Graduate School of  
ISEE, Kyushu University  
Fukuoka 814-0001, JAPAN  
Tel: +81-92-847-5193  
farhad@ejust.kyushu-u.ac.jp

## ABSTRACT

Cooling equipments is a thermal management technique that reduces the thermal resistance of the heat sink without any performance degradation. However, higher fan speed produces a lower thermal resistance, but at the expense of higher power consumption. Our proposed Neuro-Fuzzy fan controller (NFSC), minimizes fan power consumption while avoiding the temperature increase above a certain threshold. The experimental results indicate that our proposed model can significantly decrease the average fan power with negligible temperature overhead compared to the traditional fan controller.

## Categories and Subject Descriptors

Energy efficiency, Reliability

## General Terms

Experimentation, Power, Temperature, Active Cooling

## Keywords

Close-loop Control, Dynamic Thermal Management, Neuro-fuzzy

## 1. INTRODUCTION

The high demand for more computing power grows steadily which leads to increasing integration density of modern processors and generated heat. Moreover, higher temperature potentially threatens system reliability and decreases both transistor age and transition speed [1]. Dynamic Thermal Management (DTM) techniques are proposed to mitigate the aforementioned problems. DTM techniques can be categorized into in-bound and out-of-bound approaches. In-bound thermal management techniques operate in the critical path of the application execution. Stop-and-go and Dynamic Voltage and Frequency Scaling (DVFS) [1] are some hardware-based in-bound DTM techniques and task scheduling and task migration [2-3] are two software-based in-bound methods. On the other hand, out-of-band techniques [4] (e.g. CPU cooling fans) operate completely outside the critical performance path of applications. Out-of-band techniques, which use processor cooling equipments, reduce the thermal resistance of the heat sink without any performance

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).  
CF'14, May 20-22 2014, Cagliari, Italy  
ACM 978-1-4503-2870-8/14/05.

<http://dx.doi.org/10.1145/2597917.2597958>

degradation by the use of active devices such as fans. Higher fan speed produces a lower thermal resistance at the expense of higher power consumption. Previous fan controllers try to keep the temperature lower than the temperature threshold; however, it is possible to control the fan speed intelligently to keep the temperature around temperature threshold while consuming less power due to cubic relationship between fan speed and its power consumption. Motivated by these facts, this paper uses a Neuro-fuzzy based fan controller to minimize fan's power consumption while avoiding the temperature threshold violation. The experimental results indicate that the proposed NFSC, outperforms static mapping controller in reducing fan power without a substantial increase in temperature.

## 2. Fan speed controller

Our proposed method consists of offline and online phases. In the offline phase, we study the effects of different fan speeds and current CPU temperature on future CPU temperature and make a dataset based on the results. Then the Locally Linear Neuro-Fuzzy Models (LLNFM) is learned by importing produced dataset. During the online phase, fan speed selection is done by using the learned LLNFM. The inputs of LLNFM are current fan speed, current temperature, and the temperature threshold. The output of LLNFM is the suggested future fan speed that holds temperature around temperature threshold while minimizing fan power consumption due to optimized selected fan speed. Our

---

### Algorithm 1 Fan Speed Controller

---

```
1: initialize  $T_{thr}, T_{idle}$ 
2: while(True)do
3:   If (there is at least one program) then
4:      $T_{target} = T_{thr}$ 
5:   else
6:      $T_{target} = T_{idle}$ 
7:   endif
8:    $RPM_{curr} = \text{read current fan speed}$ 
9:   for  $i = 1$  to  $Core\_count$  do
10:     $T_{curr} = \text{read } core_i \text{ temperature}$ 
11:     $PredictedFanSpeed[i] = \text{Predict}(RPM_{curr}, T_{target}, T_{curr})$ 
12:   endfor
13:    $SelectedFanSpeed = \max(PredictedFanSpeed)$ ;
14:   if  $abs(OldFanSpeed - SelectedFanSpeed) \geq S_{gate}$  then
15:      $ChangeFanSpeed(SelectedFanSpeed)$ ;
16:      $OldFanSpeed = SelectedFanSpeed$ ;
17:   endif
18: endwhile
```

---

proposed fan speed controller is depicted in Algorithm 1. NFSC works in two states: where there is at least one running program that tries to hold temperature as close to the temperature threshold ( $T_{thr}$ ). In case there is no any program running, the controller tries to hold temperature around the idle temperature threshold ( $T_{idle}$ ). At each iteration of NFSC, we predict the appropriate fan speed for each core and then select the maximum value among the suggested speeds for all cores as a current fan speed. The CPU fan speeds are divided into discrete distinct speeds due to the discontinuous PWM (Pulse Width Modulation) values, which are used for changing fan speed. By the result of discretization of fan speeds, a small change may cause fluctuation between two consecutive discrete values. To avoid this phenomenon, new selected fan speed is compared with the current fan speed; if the difference is more than a specified threshold value, the fan speed is changed according to the selected value.

### 3. Experimental Results

We present a fan controller that dynamically adjusts the fan speed according to the processor temperature, fan speed and specific temperature threshold. To evaluate the proposed fan controller, we conduct our experiments with two to seven SPEC2006 benchmarks executed simultaneously on an Intel's Core i7-950. The size of the main memory of the system is 6 GB. The Linux kernel version is 3.8.0. The LM sensor [5] is used to read core temperatures and Linux fan control subsystem for changing the fan speed consequently. The range of fan PWMs of our experiments is from 75 to 255. We use *cpufreq* Linux subsystem to scale the processor frequency. We rebuild the traditional fan controller that uses a static mapping between the fan speed and the processor temperature. Figure 1 illustrates fan power consumption and temperature of the cores for traditional fan controller and NFSC on Intel core i7-950. As shown in the figure 1, our proposed method overcomes traditional controller in the fan power reduction by keeping temperature around the temperature threshold. Table 1 summarizes the comparison results of our proposed method against the traditional fan controller.

### 4. Conclusion

In this paper, a fan speed control model based on Neuro-fuzzy (NF) is presented to minimize power consumption of fan and at

the same time avoid violating the temperature threshold for multicore processors. As demonstrated, an appropriate fan speed

Table 1  
Average and total power consumption and average temperature

Method	Average power (mW)	Total power (W)	Avg. temp (°C)
Traditional	475	1187.28	51.95
NFSC	334.7	836.74	53.61

controller is extremely important in saving fan power. Experimental results based on practical benchmarks (SPEC CPU2006) running on a desktop platform (Intel Core i7-950) indicate that our proposed model can overcome traditional fan controller in saving power with negligible increase of average temperature without violating temperature threshold. Our proposed method improves average power consumption of the fan about 30% compared to traditional fan controller with an almost 3.1% (1.66°C) average temperature overhead.

### 5. REFERENCES

- [1] Kong, J., Chung, S. W., & Skadron, K. 2012. Recent thermal management techniques for microprocessors. *ACM Computing Surveys (CSUR)*. (Jun. 2012), 44(3), 13.
- [2] Salami, B., Baharani, M., & Noori, H. 2013. An adaptive temperature threshold schema for dynamic thermal management of multi-core processors. *In Computer Architecture and Digital Systems (CADS) 2013 17th CSI International Symposium on.* (Oct. 2013), 119-120.
- [3] Salami, B., Baharani, M., Noori, H., & Mehdipour, F. 2014. Physical-Aware Task Migration Algorithm for Dynamic Thermal Management of SMT Multi-Core Processors. *In Design Automation Conference (ASP-DAC) 2014 19th Asia and South Pacific.* (Jan. 2014), 292 - 297.
- [4] Ayoub, R., Sharifi, S., & Rosing, T. S. 2010. Gentlecool: Cooling aware proactive workload scheduling in multi-machine systems. *In Proceedings of the Conference on Design, Automation and Test in Europe.* (Mar. 2010), 295-298.
- [5] Lm sensors linux hardware monitoring [Online]. Available: <http://www.lm-sensors.org>.

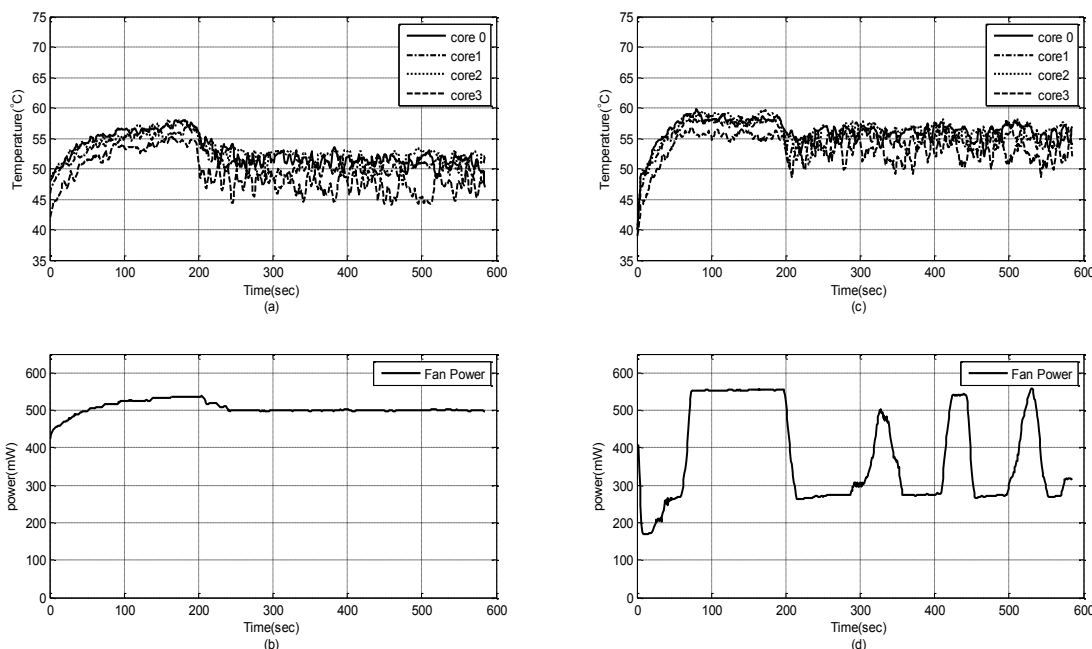


Figure 1-(a) Cores' temperature and (b) power consumption of traditional fan controller and (c) cores' temperature and (d) power consumption of our proposed method by running six programs with a temperature threshold of 58°C.