Thermal Aware Migration Techniques

Christiana Ioannou[*,1],
Yiannakis Sazeides[*,1],
Pierre Michaud[†,2]

[*] *University of Cyprus, 75 Kallipoleos Street, CY-1678, Cyprus*

[†] *Campus de Beaulieu, 35042 Rennes Cedex, France*

ABSTRACT

Multi-core chips allow thread and program level parallelism thus increasing performance. However, this comes with the cost of temperature problem. Multi-core chips require more power, creating non uniform power map and hotspots. Thread migration is one of the solutions to distribute power in a more uniform manner over the chip. This work evaluates thread migration techniques and investigates issues such as thread fairness and migration overhead.

KEYWORDS: multi-core processors; thread migration; migration overhead

1 Introduction

There is an increase in demand for high performance, low cost and low power processors to satisfy the consumer's need. One of the industry's responses was to place in a single chip more than one core. With multi-core chips thread and program level parallelism can increase performance by limiting the costs. However, multi-cores increase power consumption which in turn creates temperature problems and non uniform power density map.

Thread-migration is an emerging thermal management technique that leverages multi-cores to limit temperature. Specifically, thread-migration can help alleviate the temperature problem by distributing power consumption more uniformly over the entire chip.

This project aims to develop a thermal aware multi-core scheduler which performs thread migration based on thermal criteria and investigate thread migration issues such as migration overhead and thread fairness.

[1]E-mail: {cioannou,yanos}@cs.ucy.ac.cy
[2]E-mail: pmichaud@irisa.fr

2 Related Work

Few papers, very relevant to this project, consider applying the idea of activity-migration to a multi-core by extending the scheduling algorithm with thermal criteria for deciding migrations. As far as we know, existing scheduling kernels, such as [BC00], are temperature agnostic and prefer to assign previously running threads on the same core to preserve data affinity. They perform thread migration not for temperature reasons but to exploit idle cores to improve performance.

In [DM06], they investigate the performance of dynamic voltage and frequency scaling (dvfs) alone and in combination with thermal aware thread-migration to control temperature. They consider a system with four cores, and for one scenario, migration is allowed when the temperature sensors readings in the cores satisfy a set of temperature criteria. The criteria where to map a thread rely on temperature profile maps of threads, and cores.

[CGM+07] evaluated thread several migration methods in a thermally constraint 16-way multi-core. One of the most effective thread migration method performs thread-migration when both the highest temperature is above a certain threshold and the temperature difference between two cores is above 1 degree Celsius. Thread-migration is evaluated in combination with clock gating.latex

Thermal aware scheduling issues for multi-cores are discussed in [MS06]. They evaluated random migration in combination with clock gating. The authors argue that there is a need to redefine fairness in a thermally constrained environment.

Previous thermal aware scheduling studies consider migration algorithms that detect hot and cold threads and perform thread swapping to better distributed temperature over the entire chip. Our work will build on these and advance the state of the art mainly as follows: (a) consider both performance and fairness through the use of fairness metrics, (b)migration overhead.

3 Thermal Aware Techniques

Thread-migration may not be sufficient by itself to solve the temperature problem but can reduce the time engaging other - more detrimental to performance - thermal management techniques. The following subsections include dynamic thermal management and thread migration techniques that were used in our simulator.

3.1 Dynamic Thermal Management Techniques

The initial response from industry and academia to the multi-core temperature challenge are power and thermal management strategies [Fle01, SSH+03], that have been proposed for single core chips. For example, turning off the clock [Fle01] is one way of responding to a core temperature problem, another way is to scale voltage and frequency to decrease power more drastically and performance more gracefully.

In our thermal aware scheduler we are simulating both mechanisms. Clock gating is used as a reactive technique, whereas DVFS is used as a proactive technique. Within a scheduling period, clock gating is engaged when the temperature threshold is exceeded. At the beginning of each each scheduling period, the voltage and frequency is adjusted according the threads' performance.

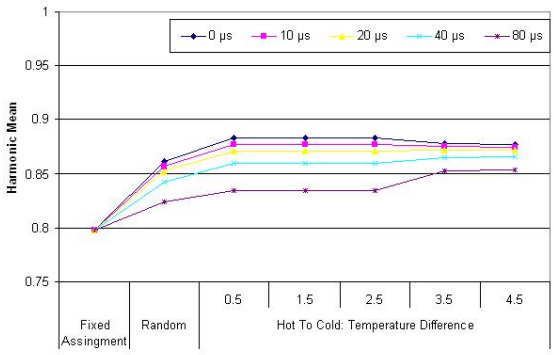

Figure 1: Harmonic Mean for Thread Mix: Cold-Cold-Hot-Hot

3.2 Thread Migration Techniques

Three migration techniques were evaluated in this work. The Fixed Assignment, Random and and Hot To Cold migration techniques. The Fixed Assignment does not allow migration, instead at the beginning of each scheduling period, the scheduler assigns each thread at the first available core. Fixed Assignment is used as a baseline to compare with Random and Hot To Cold Assignments. Random migration technique assigns each thread at a randomly selected thread. In the Hot To Cold technique, thread migration is engaged when temperature exceeds a certain threshold and there exist a temperature difference (delta) between thread and core.

The key migration parameters to be examined in thread migration techniques are: temperature threshold to engage activity migration, time elapsed since the last migration, and temperature difference between cores.

4 Simulator

An in house multi-core scheduler simulator was used to evaluate thread migration techniques. The simulator is a a parameterized synthetic simulation framework that enables quick and effective design space exploration to establish the potential of activity migration to reduce the temperature problem as compare to other techniques, such as on-off and voltage and frequency scaling. The synthetic simulation has the disadvantage of inaccurate results; however, as a first step it will provide a quick exploration of scheduling techniques. The framework will also be used to evaluate how critical it is for a kernel scheduler to be thermal aware and if so what is a good thermal aware scheduling policy. ATMI is used as the temperature model of our simulator, to calculate temperature based on power densities. The dynamic and static power densities that are used for our model are fixed.

5 Results

To evaluate our thread migration techniques we performed a series of experiments using a 4 core, single voltage domain chip and four threads of thermal characteristics: cold, cold, hot, and hot. The scheduling period was set at 1 ms and migration evaluation step was set at 2 ms. Every 2 ms, the temperature of the chip was compared with the temperature threshold of 80 C, if the threshold was exceeded then migration was allowed. The parameters that were

evaluated were migration overhead and temperature difference (delta).In figure 1, shows the performance of each thread migration technique with migration overhead values: 10 ìs, 20ìs, 40ì, and 80ìs.

The results in 1 show that thread migration increase performance by alleviating temperature. Random assignment performs approximately 6% better than Fixed Assignment, while Hot To Cold performs 9% better. The results also indicate, that migration can be a source of performance degradation. Temperature difference (delta) of Hot To Cold thread migration depend on migration overhead. When migration overhead is either 80 or 40μs ,Hot To Cold technique performs better when the temperature difference is set at 3.5 C, when migration overhead is set at 0,10 and 20μs Hot To Cold performs better when delta is set at 0.5 C. Other results, not shown here, indicate that the temperature difference (delta) value depends also on the thread mix.

6 Conclusions and Future Work

A thermally intelligent multi-core scheduler can reduce temperature and provide high performance. Our results indicate that migration alleviates temperature thus increasing performance; however, the number of migrations can be a source of performance degradation. As figure 1 shows, as migration overhead increases and the number of migrations are the same, the performance of the threads decrease. Our results also indicate the need to evaluate an adaptive temperature difference (delta) for the Hot To cold migration technique since it depends on the thread mix and the migration overhead.

For future work, we will evaluate the scalability of thermal aware scheduling techniques with increasing number of cores. To further make our synthetic simulator more realistic, the power densities of each thread will vary in time to capture differen program phases.

References

[BC00] Daniel P. Bovet and Marco Cesati. *Understading the Linux Kernel*. O'REILLY, 2000.

[CGM+07] Pedro Chaparro, José González, Grigorios Magklis, Qiong Cai, and Antonio González. Understanding the thermal implications of multi-core architectures. *IEEE Trans. Parallel Distrib. Syst.*, 18(8):1055–1065, 2007.

[DM06] James Donald and Margaret Martonosi. Techniques for multicore thermal management: Classification and new exploration. In *Proceedings of the International Symposium on Computer Architecture*, pages 78–88, June 2006.

[Fle01] M. Fleischmann. Crusoe longrun power management. *In Transmeta Corporation, Whitepaper*, 2001.

[MS06] Pierre Michaud and Yiannakis Sazeides. Scheduling issues on thermally contstrained processors. Technical report, IRISA PI-1822 INRIA RR-6006, 2006.

[SSH+03] Kevin Skadron, Mircea R. Stan, Wei Huang, Sivakumar Velusamy, Karthik Sankaranarayanan, and David Tarjan. Temperature-aware microarchitecture. *SIGARCH Comput. Archit. News*, 31(2):2–13, 2003.