

KNN Regression based Machine Learning Model for the Mitigation of Aging due to Thermal Cycling in Multi Core Processors

Jagadeesh Kumar P¹ and Mini M G²

¹⁻²Department of Electronics Engineering, Govt. Model Engineering College, Thrikkakara, Cochin-21, Kerala, India
Email: jagadeeshkumarp@gmail.com, mininair@mec.ac.in

Abstract—Processor life-time reliability is an important concern in today's high end embedded devices that are being designed with high performance multi-core processors. These processing cores are manufactured with high level of integration for meeting the high functionality per unit area demands of the applications. The lifetime reliability of such densely integrated processor cores is significantly affected by thermal stress as temperature cycles can cause failure of the cores. The damages that are caused by the cycles in the temperature accumulates over time. These thermal cycles are highly dependent on the characteristics of the workload that are being executed by the processing cores. This work proposes a K-Nearest Neighbor (KNN) based machine learning approach for the estimation of the temperature of the processor cores due to thermal cycling by analyzing the workload characteristics. This work also proposes a scheduling scheme to improve the lifetime reliability of the processing cores by appropriately allocating the workload to the cores in a multicore system. Using the experimental results, we show that the proposed methodology is accurate and suitable to implement in real-time applications.

Index Terms— lifetime reliability, multi core processor reliability, thermal cycling, KNN machine learning model.

I. INTRODUCTION

Modern embedded systems applications such as autonomous driver assistance systems, 5G communication systems, spacecrafts and submarine systems demand high performance processing of the applications. The increased performance requirements, miniaturization and functionality in these modern embedded systems necessitate the use of multi-core processors in their design [1]. The reliability requirements of the multi-core systems become more essential as these highly integrated circuits are subjected to operate at elevated temperatures. The variety of workloads that are processed by different cores in these multicore systems leads to thermal stress and deteriorate the system reliability. The ageing effects due to thermal cycles accumulate over time and can eventually lead to the system failure. Despite the importance of thermal cycles on the life-time reliability, most schedulers will try to trade-off between the performance and energy consumption of the application execution regardless of their adverse effect on the reliability of the processor cores [2],[3]. In this paper, we present a thermal cycle aware workload assignment to the cores considering the life-time reliability requirement of the processing core and the performance requirement of the application. The proposed

KNN based machine learning model can estimate the thermal profile of the major power consuming logical units of the multi-core processor for the given workload. Based on the estimated thermal profile, scheduler can compute the mean time to failure (MTTF) due to thermal cycles. Based on the estimated MTTF and the life-time reliability budget, the scheduler can assign the workload to an appropriate core and decide the frequency of operation, such that the life-time reliability requirements are met. The scheduler may decide to operate the core at a lower operating speed, if demanded with a high reliability budget, with a graceful degradation of the application performance.

II. RELATED WORK

Many studies have reported that thermal cycling is an important ageing mechanism and affects system reliability of multi-core systems as it can cause fatigue failures. [4],[5]. The thermo-mechanical stress caused due to the expansion coefficient mismatch between the layers leads to several failure mechanisms such as broken bond wire, solder fatigue, dielectric film voids and broken die [6]. The increased power densities due to the high temperature of operation and the effect of scaling has a negative impact on the long-term reliability of the processor cores. An attempt to quantify the impact of scaling on the processor lifetime reliability due to various intrinsic effects including thermal cycling is reported in [7].

The conventional power management techniques frequently turn on and off the processor cores, which causes the thermal cycles to increase [8]. Conventional workload assignment in multicore systems without considering the thermal characteristics of the workload can accelerate various ageing mechanisms and reduce the processor lifetime reliability. In order to effectively manage the thermal characteristics and to improve the lifetime reliability, autonomous mechanisms are required to analyze the thermal profiles of the workloads and to take appropriate control action at runtime. A runtime approach which uses reinforcement learning to select workload specific control action for improving the average temperature and thermal cycling by overriding the operating systems native thread allocation decisions is proposed in [9]. A thermal-aware scheduling scheme for mixed-criticality systems that integrate applications of different levels of criticality into the same platform is reported in [10].

Chip-wide thermal profile of multicore processing systems can be effectively controlled with optimized workload assignment to the processing elements. Workload mapping based on local optimizations such as to minimize the number of hops between communicating tasks, can cause the formation of hotspot zones and the underutilization of the processing elements. A runtime mapping heuristic to minimize the occurrence of hotspots by distributing the application workload in a uniform way to the processing elements is reported in [11].

Efficient thermal estimation techniques are required to effectively deploy runtime control strategies for the temperature management in processing systems. One way to implement an efficient thermal estimation technique is to predict the thermal characteristics of the workload when it gets executed in a processing element. A generic method to derive an analytical model for the estimation of temperature in real time for System on Programmable Chips (SoPCs) is reported in [12].

The above review shows that real time thermal estimation and processor adaptation techniques are pressing research topics for the effective thermal management of multi core processors. However not many research publications are seen which utilizes machine learning approaches for the estimation of the thermal profile of the processing cores based on the thermal characteristics of the workload and taking real-time control action to meet the processor lifetime requirements. In this work we propose a machine learning approach for the prediction of the thermal characteristics of the workloads when gets executed in a processor core with defined configurations. Using the predicted thermal characteristics, lifetime reliability of the processor cores is estimated. We also propose a control algorithm to take necessary control action for meeting the reliability budgets of the processor cores. The problem statement is described in the next section.

III. PROBLEM STATEMENT

Thermal cycling is the process of frequent heating and cooling, when workloads with different thermal characteristics get executed in the processor cores [7]. It is also due to the frequent turning on and off of the processor cores in a multicore processor when runtime power control mechanism is implemented. The turbo modes of operation in modern processors can increase the severity of thermal cycles as these processors can have higher local temperature swing. Thermal cycling if left unattended will degrade the processor's packaging characteristics and reduce the useful lifetime of operation.

IV. SOLUTION APPROACH AND METHODOLOGY

This work proposes a KNN based machine learning methodology to predict the thermal profile of the various processing elements (PEs) for the workloads to be executed in the available processor cores in a multicore processing system. The thermal profile of the processing elements is predicted for the ready tasks in each scheduling interval, and using the predicted temperature, the effect of aging due to thermal cycling is computed. The scheduler then decides the feasible operating frequency of the processor cores, based on the workload performance requirement and the lifetime reliability budget.

Differences in the coefficients of thermal expansion (CTE) between the different materials in integrated circuits introduces thermally induced stresses and strains due to thermal cycles. Aging effects accumulate every time when there is a thermal cycle and reduce the lifetime reliability of the processor. CTE mismatches are predominant in the 1st and 2nd level of interconnects in the integrated circuits where the first level of interconnect is in between the die and substrate and the second level connects substrate to the printed circuit board. Thermal cycles in integrated circuits can be broadly classified into two types, one that happens at low frequency called large thermal cycles and the second type occurs at high frequency, called small thermal cycles. Large thermal cycles in multicore systems are due to the powering up and down of the processor cores based on the workload assignment or when the cores operate in low power or stand-by mode. Large thermal cycles are mainly due to context switching in application execution. In this work we focus on the reliability impact of large thermal cycles due to scheduling of workloads in a multicore processing system. Coffin-Manson equation is used to model large thermal cycles and is represented in equation 1 [13].

$$N_f = C_0(\Delta T)^{-q} \quad (1)$$

where N_f represents the number of cycles to failure, C_0 and q are empirically determined constants, and ΔT is the thermal cycle temperature range. The value of the constant C_0 depends on the type of the material and q is known as the Coffin-Manson exponent.

Based on equation 1, the mean time to failure due to thermal cycling ($MTTF_{TC}$) can be defined as in equation 2.

$$MTTF_{TC} \propto \left(\frac{1}{T - T_{ambient}} \right)^q \quad (2)$$

where T is the average temperature of the structure and $T_{ambient}$ is the ambient temperature. The proportionality constant in equation 2 also factors in the frequency of thermal cycling and is assumed to be constant in this work as the workload assignment by the scheduler is done in fixed intervals. Hence $T - T_{ambient}$ depicts the modelled thermal cycle.

In this work a fine-grained approach for the estimation of the temperature of the major power consuming PEs of the core is proposed. With the estimated thermal profile, it is possible to predict the contribution of aging and the associated lifetime reliability of the major power consuming PEs of the processor core. We assume a Sum-of-Failure-Rates (SOFR) model where the processor is considered to be a series failure system, and the first-time failure of any of the PE will cause the processor to fail.

For the development of the machine learning models for thermal prediction, we use different categories of the representative embedded applications in the MiBench suite [14]. The bench mark suite has tasks representing multiple embedded application areas such as industrial control, consumer devices, office automation, networking, security, and telecommunications. Tasks belonging to the consumer, network and security application areas are used in this work for developing the machine learning models for the estimation of temperature. JPEG encode and JPEG decode are selected from the consumer category and focus primarily on multimedia applications. Products belonging to consumer electronics category such as digital cameras, scanners etc. find the implementation of the JPEG encode and decode algorithms. The Dijkstra task implement logic for finding the shortest path in a graph and belong to the network category of applications. Switches and routers are two examples of the products belonging to the network category. Secure Hash Algorithm (SHA) is used in many ecommerce related applications and belong to the security category. SHA algorithm is often used in secure exchange of crypto-graphic keys and for generating digital signatures.

Execution characteristics of the mentioned benchmark applications are to be analyzed for estimating the dynamic power consumption values of the various processing elements (PEs) of the multi core processor. Using the estimated dynamic power consumption values, thermal estimation of the PEs is carried out. Gem5 is the simulation framework used for the estimation of the execution characteristics of the workloads [15]. Gem5 simulator is a configurable modular framework and is possible to simulate one or more computer system architectures in various ways. The architectures supported by Gem5 include ARM, ALPHA, MIPS and x86.

Syscall emulation (SE) mode of gem5 is used in this work to simulate the workload execution of processing cores. Since modeling the operating system is beyond the scope of this research work, and need to emulate only Linux system calls for analyzing the execution characteristics of the workload, syscall emulation mode of gem5 is selected. The characteristics of the workload analyzed using SE mode are i) number of CPU cycles ii) the number of integer instructions executed iii) the number of float instructions executed iv) the number of load/store instructions executed v) the number of integer/float register reads/writes vi) the number of data cache read/write hits and misses and vii) the number of instruction cache read/write hits and misses. The set of characteristics mentioned above are found to have a significant effect on the dynamic power consumption of the various processing elements of the core for the selected benchmark applications.

The execution characteristics of the workloads are fed as input to the tool McPAT, for estimating the dynamic power consumption of the various PEs of the cores. McPAT (Multicore Power, Area, and Timing) is a power modeling framework and can be used for the dynamic power consumption of the PEs in multi-threaded and multicore processors [16]. Microarchitectural parameters of the processing cores along with the dynamic activity statistics of the workloads are fed using an XML-based interface to McPAT. The dynamic power consumption of the PEs: integer ALU (int_ALU), integer Reg (int_Reg), Floating Point Register (FP_Reg), Floating Point Unit (FPU), Load/Store Queue (LDSTQ), Data Translation Lookaside Buffer (DTLB), and Instruction Translation Lookaside Buffer (ITLB) are analyzed using McPAT. The dynamic power consumption values are used to estimate the corresponding thermal profile of the PEs.

Thermal profile of the PEs is estimated using the widely used temperature modeling framework HotSpot [17]. The micro architectural blocks of the processing cores can be represented by an equivalent circuit of thermal resistances and capacitances. Dynamic power values of the PEs computed using McPAT along with the essential aspects of the thermal package of the processor is fed to HotSpot for thermal estimation. Workload characteristics that have higher influence on the power consumption of the PEs are analyzed using McPAT. Such workload characteristics of the representative embedded applications mentioned above and the corresponding thermal profile of the PEs are used as the training data set for the development of the machine learning model. We propose the use of KNN regression methodology for the development of the machine learning model. The regressor used is 'sklearn.neighbors.KNeighborsRegressor', and is available in a free Python library called Scikit-learn [18]. Using KNN regression methodology, the temperature of a PE is predicted by local interpolation of the temperatures of the nearest neighbors in the training set. Multiple instances of the workloads are derived by varying the computational load of the applications available in the MiBench suite and the execution characteristics of such workloads are used as the training data set. The set of workloads used for the model development can be represented as a matrix $\mathbf{W}$ of dimension $\mathbf{n} \times \mathbf{m}$, where each row of the matrix represents $\mathbf{m}$ instances of a workload belonging to an application category and is derived by varying the computational load of the application. Each column of the matrix represents the workloads assigned to $\mathbf{n}$ number of processing cores in a particular scheduling interval.

$$\mathbf{W} = \begin{bmatrix} W_{11} & W_{12} & \dots & W_{1m} \\ W_{21} & W_{22} & \dots & W_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ W_{n1} & W_{n2} & \dots & W_{nm} \end{bmatrix} \quad (3)$$

With the developed thermal model, steady state temperature of a core can be predicted when a task gets executed in the core. The thermal model can be used to implement runtime control mechanisms for improving the lifetime reliability of the processing cores. Scheduler can use the temperature model to predict the thermal profile of the available processing cores for the ready tasks. Based on the estimated thermal profile, and the performance requirement of the applications, scheduler can decide the appropriate operating frequency of the cores for meeting the lifetime reliability budget.

The constant failure rate of an integrated circuit can be represented using failures in time (FIT). The relationship between FIT and MTTF is defined as in equation 4.

$$MTTF = \frac{10^9}{FIT} \quad (4)$$

By referring to the data available for the present generation of Texas Instruments industrial-grade integrated devices [19], we consider a processor MTTF of 10 years when operating at a junction temperature T_j of 105 °C. The FIT value corresponding to this MTTF is approximately 11500. We consider five important intrinsic failure mechanisms Electromigration (EM), Stress Migration (SM), Time-Dependent Dielectric Breakdown (TDDB), Thermal Cycling (TC), and Negative Bias Temperature Instability (NBTI) and assume that the total FITs are

equally divided among these failures. With the predicted steady state temperature and the estimated MTTF, scheduler is able to compute the frequency of operation of a core for meeting the reliability budget. The developed thermal model is able to predict the steady state temperature of the PEs of the cores operating at predefined discrete set of frequencies. Since it is required to have a finer estimation of frequency values for the better control of ageing, an interpolator can be used to compute the required frequency of operation from the discrete set of frequencies estimated by the scheduler. The cubic spline in `scipy.interpolate` is used to find the interpolated frequency value for the operation of the core. Block level representation of the runtime control logic for improving the lifetime reliability of a processor core is represented in Fig. 1.




Figure 1. Block diagram of the runtime control logic

V. RESULTS AND DISCUSSION

In this first part of the results section, we compare the performance of the proposed thermal prediction model. The proposed KNN model for thermal prediction is developed by training the model with the dataset derived from the categories of the embedded applications: consumer, network and security from the MiBench suite. The tasks JPEG encode/decode, Dijkstra, and Secure Hash Algorithm (SHA), with new set of workloads and belonging to the consumer, network, and data security applications respectively are used to validate the thermal prediction models. We follow a fine-grained approach for thermal estimation, where the steady state temperature of the logical units: integer ALU (int_ALU), integer Reg (int_Reg), Floating Point Register (FP_Reg), Floating Point Unit (FPU), Load/Store Queue (LDSTQ), Instruction Translation Lookaside Buffer ITLB and Data Translation Lookaside Buffer (DTLB), which are having significant dynamic power consumption, are estimated using the developed thermal prediction model. Fig. 2 compares the steady state temperature values of the PEs estimated using prediction model with that estimated using HotSpot tool for the task `cjpeg` executing on a core with an operating frequency of 2.4 GHz. Percentage error in the prediction of steady state temperature is shown in Fig. 3.




Figure 2. SSTs of PEs for the task `cjpeg`




Figure 3. Percentage error in the estimation of SSTs of the PEs

Thermal prediction model is validated by comparing the predicted temperature values of the PEs of the core with that estimated using the tool HotSpot. Fig. 4 compares the predicted temperature values of the PE,

int_ALU, with the values estimated using HotSpot for different tasks, when the core is operating at a frequency of 2.4 GHz. Percentage error in the prediction of steady state temperature of int_ALU is shown in Fig. 5.




Figure 4. SSTs of int_ALU for different tasks




Figure 5. Percentage error in the estimation of SSTs of int_ALU for different tasks

In this second part of the results section, we demonstrate how the developed thermal prediction model is helpful in implementing the runtime control mechanisms for meeting the lifetime reliability requirements with the scheduler. For ready tasks, the scheduler with the estimated thermal values, compute MTTF due to thermal cycling of the cores corresponding to the required performance level of the task. If the estimated MTTF is above the lifetime reliability budget of the core, scheduler will compute the operating frequency closer to the required frequency of operation. 1-D and Spline are the two interpolation schemes adopted in this work for the finer estimation of the operating frequency of the core. Fig. 6 shows the MTTF frequency tradeoff of the major power consuming PEs of a core, executing the task SHA, when the temperature computations are based on the tool HotSpot. Fig. 7 shows the corresponding MTTF frequency tradeoff, when the thermal estimation of the PEs is done with the proposed KNN regression model and the finer estimation of the frequency is done with 1-D interpolation. Fig. 8 represents the MTTF frequency tradeoff when Spline interpolation is adopted. Accuracy of the proposed lifetime reliability scheme is validated by calculating the percentage error in the estimation of MTTF due to thermal cycling. Fig. 9 and Fig. 10 shows the percentage error in the estimation of MTTF of PEs for different frequencies of operation with 1-D and Spline interpolation schemes respectively.




Figure 6. MTTF-Frequency tradeoff (actual) of the PEs




Figure 7. MTTF-Frequency tradeoff of the PEs estimated using the proposed KNN with 1-D interpolation scheme




Figure 8. MTTF-Frequency tradeoff (KNN with Spline) of the PEs




Figure 9. Percentage error in the estimation of MTTF when KNN regression with 1D interpolation is used




Figure 10. Percentage error in the estimation of MTTF when KNN regression with Spline interpolation is used

VI. CONCLUSION AND FUTURE SCOPE

In this research work we have proposed a KNN regression based thermal estimation methodology suitable for the implementation at runtime. The methodology is suitable for computing systems running embedded applications where the workloads are having predictable execution behavior. Accuracy of the thermal models are evaluated and the percentage error in the estimation is observed to be less than ± 0.5 and usually less than ± 0.2 percentage. This work also proposes a control mechanism to adapt the operating frequency at runtime to ensure the processing core meeting the lifetime reliability requirement. Two interpolation schemes: 1-D and Spline are adopted as options for the finer estimation of the operating frequency of a processor core.

The proposed scheme is extendable as it is possible to include the effects of other prominent intrinsic effects affecting the lifetime reliability of the processing cores. It is also possible to improve the accuracy of estimation by incorporating necessary feedback to the prediction and control algorithms. The data from on-chip thermal sensors can be used periodically for improving the accuracy of the thermal prediction algorithm. The methodology is also helpful for the faster validation of computing platforms which helps embedded system designers quickly validate and select the suitability of a particular multicore computing system for implementing embedded applications.

REFERENCES

- [1] T. Chantem, Y. Xiang, X. S. Hu and R. P. Dick, "Enhancing multicore reliability through wear compensation in online assignment and scheduling", Proc. Design Autom. Test Europe (DATE), pp. 1373-1378, Mar. 2013.
- [2] S. Alsoubaihi and J. Gaudiot, "PETS: Performance, energy and thermal aware scheduler for job mapping with resource allocation in heterogeneous systems," 2016 IEEE 35th International Performance Computing and Communications Conference (IPCCC), 2016, pp. 1-2, doi: 10.1109/PCCC.2016.7820609.
- [3] E. Jiang, L. Wang and J. Wang, "Decomposition-based multi-objective optimization for energy-aware distributed hybrid flow shop scheduling with multiprocessor tasks," in Tsinghua Science and Technology, vol. 26, no. 5, pp. 646-663, Oct. 2021, doi: 10.26599/TST.2021.9010007.
- [4] A. Iranfar, M. Kamal, A. Afzali-Kusha, M. Pedram and D. Atienza, "TheSPoT: Thermal Stress-Aware Power and Temperature Management for Multiprocessor Systems-on-Chip," in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 37, no. 8, pp. 1532-1545, Aug. 2018, doi: 10.1109/TCAD.2017.2768417.
- [5] Q. Zou, E. Kursun and Y. Xie, "Thermomechanical Stress-Aware Management for 3-D IC Designs," in IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 25, no. 9, pp. 2678-2682, Sept. 2017, doi: 10.1109/TVLSI.2017.2707119.
- [6] J. W. McPherson, Reliability Physics and Engineering, New York, NY, USA:Springer, 2010
- [7] J. Srinivasan, S. V. Adve, P. Bose and J. A. Rivers, "The impact of technology scaling on lifetime reliability," International Conference on Dependable Systems and Networks, 2004, 2004, pp. 177-186, doi: 10.1109/DSN.2004.1311888.
- [8] A. K. Coskun, T. S. Rosing and K. C. Gross, "Temperature management in multiprocessor SoCs using online learning", Proc. Design Autom. Conf. (DAC), pp. 890-893, 2008.
- [9] Das, G. V. Merrett, M. Tribastone and B. M. Al-Hashimi, "Workload Change Point Detection for Runtime Thermal Management of Embedded Systems," in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 35, no. 8, pp. 1358-1371, Aug. 2016, doi: 10.1109/TCAD.2015.2504875.

- [10] S. Safari, H. Khdr, P. Gohari-Nazari, M. Ansari, S. Hessabi and J. Henkel, "TherMa-MiCs: Thermal-Aware Scheduling for Fault-Tolerant Mixed-Criticality Systems," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 7, pp. 1678-1694, 1 July 2022, doi: 10.1109/TPDS.2021.3123544.
- [11] M. Mandelli, G. Castilhos, G. Sassatelli, L. Ost and F. G. Moraes, "A distributed energy-aware task mapping to achieve thermal balancing and improve reliability of many-core systems," 2015 28th Symposium on Integrated Circuits and Systems Design (SBCCI), 2015, pp. 1-7.
- [12] Dimple Sharma, Lev Kirischian, "A Method for Run-Time Prediction of On-Chip Thermal Conditions in Dynamically Reconfigurable SOPCs", *International Journal of Reconfigurable Computing*, vol. 2021, Article ID 8818788, 20 pages, 2021. <https://doi.org/10.1155/2021/8818788>
- [13] Failure Mechanisms and Models for Semiconductor Devices. In JEDEC Publication JEP122-H, 2016
- [14] M. R. Guthaus, J. S. Ringenberg, D. Ernst, T. M. Austin, T. Mudge and R. B. Brown, "MiBench: A free, commercially representative embedded benchmark suite," *Proceedings of the Fourth Annual IEEE International Workshop on Workload Characterization. WWC-4 (Cat. No.01EX538)*, pp. 3-14, 2001.
- [15] Nathan, Binkert., Bradford Beckmann., Gabriel Black., et.al.: 'The gem5 Simulator', *ACM SIGARCH Computer Architecture News*, pp. 1–7, May 2011.
- [16] A. Guler and N. K. Jha, "McPAT-Monolithic: An Area/Power/Timing Architecture Modeling Framework for 3-D Hybrid Monolithic Multicore Systems," in *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 28, no. 10, pp. 2146-2156, Oct. 2020, doi: 10.1109/TVLSI.2020.3002723.
- [17] Kevin, Skadron., Mircea, R. Stan., Wei, Huang., et.al: 'Temperature aware microarchitecture', *ACM SIGARCH Computer Architecture News - ISCA 2003*, May 2003, 31, (2), pp. 2–13.
- [18] Pedregosa, Fabian, et al. "Scikit-learn: Machine learning in Python." *the Journal of machine Learning research* 12: pp. 2825-2830, 2011.
- [19] Allan Webber, *Calculating Useful Lifetimes of Embedded Processors*, Texas Instruments Application Report, SPRABX4B, November 2014 Revised March 2020.