

Formal Probabilistic Analysis of Distributed Dynamic Thermal Management

Shafaq Iqtedar*, Osman Hasan*, Muhammad Shafique†, Jörg Henkel†

*School of Electrical Engineering and Computer Science, National University of Sciences and Technology, Islamabad, Pakistan

†Chair for Embedded Systems (CES), Karlsruhe Institute of Technology (KIT), Germany

Email: {10besesiqtedar, osman.hasan}@seecs.nust.edu.pk; {muhammad.shafique, henkel}@kit.edu

Abstract—The prevalence of Dynamic Thermal Management (DTM) schemes coupled with demands for high reliability motivate the rigorous verification and testing of these schemes before deployment. Conventionally, these schemes are analyzed using either simulations or by running on real systems. But these traditional analysis techniques cannot exhaustively validate the distributed DTM schemes and thus compromise on the accuracy of the analysis results. Moreover, the randomness due to task assignments, task completion times and re-mappings, is often ignored in the analysis of distributed DTM schemes. We propose to overcome both of these limitations by using probabilistic model checking, which is a formal method for modeling and verifying concurrent systems with randomized behaviors. The paper presents a case study on the formal verification of a state-of-the-art distributed DTM scheme using PRISM model checker and a couple of issues which were identified during verification.

I. INTRODUCTION

With fabrication technologies entering the nano era, elevated temperatures resulting from increased on-chip power densities, are jeopardizing the reliability and performance of on-chip systems. Dynamic Thermal Management (DTM) schemes [7] have been introduced to control these elevated temperatures. Conventionally, DTM schemes were based on centralized approaches [15]. But with the prevalence of many-core architectures, the overhead of monitoring the entire chip has significantly been increased. Therefore, distributed DTM schemes came forth that tackle the complexity and scalability issues by transforming the problem space from a global one to many smaller regional ones, thus can exploit locality when making DTM decisions. However, for a distributed DTM scheme to achieve the same performance as is possible from a global DTM scheme, an exchange of state information across regions becomes necessary in order to negotiate a near optimal system state configuration [5].

The widespread usage of DTM schemes [7] along with high demands for reliability require the rigorous verification and testing of these schemes before deployment. Traditionally, these schemes are analyzed using either simulations or by running on real hardware systems. However, simulation based verification is not complete in terms of error detection. A simulator does not exhaustively traverse through all possible system paths. Even if some corner cases can be specifically targeted, there is no proof that these represent all the scenarios in which the DTM scheme would not be able to achieve a stable system configuration, and it is never possible to consider or even foresee all corner cases.

Moreover, in real systems, the workloads distribution and in some cases the initial task assignment across the many-core system is random [5]. This randomness in the behavior, due to task assignments, task completion times and re-mappings, is often ignored in the analysis of DTM schemes. This fact compromises the analysis results as well by ignoring certain real time scenarios in which a DTM scheme may not be able to provide thermal balancing, resulting into the creation of thermal hotspots. These limitations in the analysis may in turn lead to delays in deployment of DTM schemes as happened in the case of Foxton DTM that was designed for the Montecito chip [1].

To overcome the above mentioned limitations, model checking [10] can be used in the analysis of DTM schemes. Model checking techniques are based on models describing the given systems behavior in a mathematically precise and unambiguous manner. A model checker, exhaustively and automatically checks whether a model meets the given specification while ensuring 100% completeness and soundness of analysis. Traditional model checking has been employed successfully in the analysis of few Dynamic Power Management (DPM) schemes [11]. Probabilistic model checking of DPM schemes for many-core architecture has also been presented in [9]. Given the safety-critical nature of DTM schemes, a traditional model checking technique has also been recently proposed to verify distributed DTM schemes in [3]. However, traditional model checking techniques provide answers to simple non-probabilistic (Yes/No) functional properties. For instance, in case of a distributed DTM, conventional formal analyses can verify if the thermal hotspots can be created or not. However, *an important piece of knowledge is about how often this can happen*. And in case if a thermal hotspot is created, it is important for a system designer to evaluate the duration that the DTM algorithm would take to balance the temperatures across the chip. These statistics can play a vital role in the design of the next generation DTM schemes for many-core systems. To the best of our knowledge, *no formal probabilistic verification method, including detailed quantitative analysis while covering the randomized behaviors, has been used in the context of verifying distributed DTM schemes for many-core systems so far*.

A. Our Novel Contributions and Concept Overview

In this work, we present a case study on the formal verification of state-of-the-art agent-based distributed DTM scheme, namely Thermal-aware Agent-based Power Economy

(TAPE) [5], in order to illustrate the effectiveness of probabilistic model checking in context of verifying distributed DTM schemes. Our verification process is based on a widely used model checking tool, i.e., the PRISM model checker [2]. We ascertain the correctness of TAPE by constructing its Markovian model and analyzing the formally specified quantitative properties using PRISM. Moreover, we have also performed a verification feasibility analysis of TAPE DTM scheme. One of the main reasons for using TAPE as our case study is the ability to compare our results with the recent state-of-the-art formal verification technique, that uses traditional model checking, for a fair comparison and technology advancement [3]. Our results show the usefulness of probabilistic model checking for analyzing distributed DTM schemes, as a number of critical issues with TAPE have been identified in this paper, which could not have been caught in [3].

II. BACKGROUND AND PRELIMINARIES

A. Probabilistic Model Checking

Probabilistic model checking [6] is a formal method for modeling and verifying systems that exhibit randomized behaviors. A model checker exhaustively traverses the entire state-space of a design in ascertaining correctness. The 100% completeness of analysis coupled with the consideration of randomized behaviors and provision of a detailed quantitative knowledge makes probabilistic model checking a very powerful framework for verifying distributed DTM schemes. One limitation of model checking is that the state-space of a system can be very large, or sometimes even infinite. For instance, in case of distributed DTM schemes, the state space grows exponentially with increasing number of cores and tasks. Thus, it becomes computationally impossible to explore the entire state-space with limited resources of time and memory. This problem, termed as *state-space explosion*, is usually resolved by developing abstract, less complex, models of the system.

B. PRISM Model Checker

PRISM [2] is a widely used probabilistic model checker. It provides support for analyzing different kind of probabilistic models, such as discrete-time Markov chains (DTMCs) [4], in which time is modeled as discrete steps. Models are specified using the PRISM modeling language, i.e., a simple, state-based language based on the reactive modules formalism [2]. Distributed DTM schemes can be described in PRISM as the parallel composition of a set of modules that are running over different cores. The state of each module can be given by a set of finite-ranging variables and its behavior by a set of probabilistic guarded commands. PRISM also allows models to be augmented with costs and rewards, i.e., the real values assigned to states and transitions of the model. This permits reasoning about a much wider range of quantitative measures of a system, e.g., “completion time” or “energy consumption” [8]. Such statistics are very useful in the analysis of distributed DTM schemes. For instance, it allows the designer to evaluate the efficiency of distributed DTM scheme without the overhead of modeling time separately; see Section III for more details. We chose PRISM for the proposed work since it allows us to

calculate actual probabilities, rewards based timings properties and supports a wide variety of probabilistic models. Thus, we can evaluate both the effectiveness and the efficiency of distributed DTM schemes using PRISM.

III. CASE STUDY ON PROBABILISTIC FORMAL ANALYSIS OF TAPE DTM

For illustration purposes, we present an extensive case study on probabilistic formal verification of Thermal-aware Agent-based Power Economy (TAPE) [8], which is a state-of-the-art agent-based DTM scheme.

A. Introduction to TAPE

TAPE is a fully distributed agent-based DTM approach, which focuses on distributing power and the resulting heat evenly over a many-core architecture. The approach incorporates thermal penalties in the agent negotiations to deal proactively with potentially developing thermal hotspots explicitly. Each core is assigned its own agent, which individually distributes power units between tiles of a many-core architecture through local trade, i.e., only with its neighbors. Initially, a predefined number of power units are distributed evenly among all agents. Each of these power units belong to one of the following two categories: *used* units are currently deployed by the Processing Element (PE) to run tasks and *free* units are power units available on a tile but not currently needed by the tile to run its allocated tasks. These units are used by each agent to calculate the buy and sell values for trading power units. These buy and sell values vary depending on the supply/demand of power units. At start-up, a tile is randomly chosen to map a task as all tiles are equally suited to run the task. Otherwise the task is mapped to a tile with most available units (In TAPE it is the tile with maximum $sell_{T_n} - buy_{T_n}$ value). At each time interval, the agent n first calculates the following value:

$$Sellbase = w_{u,s} * used_n + w_{f,s} * free_n \quad (1)$$

$$Buybase = w_{u,b} * used_n - w_{f,b} * free_n \quad (2)$$

$$Buy\ modifier : buy_n = Buybase - a_b * T_n \quad (3)$$

$$Sell\ modifier : sell_n = Sellbase - a_s * T_n \quad (4)$$

Then the agents compare their own resulting modified buy/sell values with the buy/sell values of their neighbors using the following equation which provides a stability criteria:

$$(sell_n - buy_n) - (sell_i - buy_i) < T_n \quad (5)$$

If any of the neighbors i does not fulfill the above equation, then the agent n relinquishes a power unit to the neighbor with the maximum $sell_i - buy_i$ value. If the agents have free power units then this is simply a matter of decrementing its number of free power units otherwise, the agent must give up one of its used power units and re-mapping needs to be invoked.

Algorithm 1 Mapping Software

Module (re-) mapping

```

1:  $\llbracket (Condition_1 = true) - >$ 
2:    $1/tiles : (remap\_core\_no' = 11)$ 
3:    $+1/tiles : (remap\_core\_no' = 12)$ 
4:    $+1/tiles : (remap\_core\_no' = 21)$ 
5:    $+1/tiles : (remap\_core\_no' = 22)$ 
6:  $\llbracket (Condition_2 = true \ \& \ max\_diff = diff\_11) - >$ 
7:    $(remap\_core\_no' = 11)$ 
8: ...
endmodule

```

B. Verification Challenges of TAPE DTM scheme

The most important challenge in analysis of TAPE DTM scheme is to ensure that the system will always attain a stable power budget distribution without any redundant trading (i.e., two agents continuously trading one free unit back and forth preventing the system from stabilizing). Moreover even if a stable power budget distribution is attained, it is important to verify that no thermal hotspots are created and temperatures are distributed evenly across the chip. Another major challenge is to determine the optimal values of unknown weight parameters (i.e., a_s , a_b , $w_{u,s}$, $w_{f,s}$, $w_{u,b}$, $w_{f,b}$) in TAPE, as the performance of the algorithm depends highly on the values of these parameters. We have successfully addressed all the above mentioned challenges in this paper.

C. Modeling TAPE in PRISM

1) Modeling distributed agents and the (re-) mapping software: The first step in our verification process is to describe the behavior of the TAPE algorithm running on a many-core chip as a DTMC. The grid is modeled as a two dimensional array of distributed nodes such that the TAPE algorithm runs on all these nodes concurrently. We developed a generic node model so that it can be repeatedly used to formally express any grid of arbitrary dimension. Using our approach, the agents (associated with the nodes) are realized as separate modules [5] executing in parallel. Every agent module is associated with the following set of variables: $free_n$, $used_n$, $Buybase_n$, $sellBase_n$, $buyT_n$, $sellT_n$, where, n represents the number of the corresponding core. As mentioned in [8], the software entity that is responsible for application (re-)mapping may be executed in any of the existing processing elements inside the many-core architecture. Therefore for the selection of the mapping-core, a separate module is constructed. This module selects a core on the basis of following two conditions:

- 1) If all the tiles are equally suited for mapping the new task then one tile is chosen randomly among all the tiles with probability equal to one divided by the total number of cores (See $Condition_1$ in Algorithm 1).
- 2) Otherwise, the task is (re-) mapped to a tile where most free power units are available and in TAPE algorithm, this is simply the tile with the maximum $sellT_n - buyT_n$ value. Again if more than one tile fulfills this criterion, then one of them is chosen randomly (See $Condition_2$ in Algorithm 1).

Thus, whenever a new task is initiated, this module selects a tile for mapping the task and invokes the corresponding agent

module of that tile. The new tasks are initiated in a separate module, which randomly selects a task execution time and just sends a signal to the mapping software entity.

Once a core is selected for mapping or re-mapping the task, its corresponding agent is invoked which is a separate module and it increments the used units and decrements the free units, depending on the task execution time. The higher the value of task-time, the more power units will be consumed and it is assumed that only 1 power unit is consumed for 1 time unit. As the free and used units are changed, the buy and sell values are modified automatically.

2) Discretization of Temperature Variable: Whenever a free power unit is deployed for executing a task, the free unit is converted into a used power unit and the temperature variable (T_m) of the corresponding tile is incremented by $4^\circ K$ (as done in [3]). At start-up, an initial value of $30^\circ K$ is assigned to the temperature variable.

3) Reward-Structure for Modeling Time: For the verification of timing properties, we extend the probabilistic model with rewards in order to compute the time-based properties. In every execution of the algorithm, the reward is accumulated and the final value of the accumulated reward determines the number of steps in attaining a stable system configuration. This is a major advantage of using PRISM as it removes the overhead of modeling time separately, as was done with Lamport time stamps in [3].

4) Modeling Agent based negotiations: As discussed earlier that in order to avoid the creation of thermal hotspots, agent based negotiations take place and the power units are exchanged among the agents. PRISM does not provide a mechanism to pass messages between different modules. Therefore, in order to implement the agent-based negotiations, we developed a separate module to perform the exchange of information. At every trade interval, this module compares the modified buy/sell values of all nodes with the buy/sell values of their neighboring nodes and ensures that the stability equation is satisfied. If the equation is not satisfied and the differences among the buy and sell values goes beyond the threshold, this module sends a signal for the trading of power units to the respective increment-and-decrement agents. For every node, there is a power trading agent as well (see Algorithm 2) for the exchange of power units. Whenever a signal is received, one agent relinquishes a power unit to the other agent. These power trading agents are realized in separate modules from mapping-agent's modules (as mentioned in TAPE DTM scheme [5]). The agents will keep trading the power units until the stability equation is satisfied again. Moreover, if during the power-trade, re-mapping is invoked the power trading agent sends a signal to the corresponding (re-) mapping module.

5) Characteristics of weights (a_s , a_b , $w_{u,s}$, $w_{f,s}$, $w_{u,b}$, $w_{f,b}$): An important factor in ensuring that the system reaches a stable state in the TAPE DTM technique is the choice of weights used to determine the buy and $sell$ values. These weights are constant values with respect to time. *A challenge in the analysis of TAPE algorithm was to determine the optimal values of these parameters.* Our proposed probabilistic analysis allowed us to come up with some interesting results during the calculation of optimal values for these parameters.

Algorithm 2 Power Trading Agent

Module power_trading_agent_n

```

1:  $\llbracket (dec\_n = true) \ \& \ (free\_n > 0) \rrbracket - >$ 
2:    $(free\_n' = free\_n - 1) \ \& \ (dec\_n' = false);$ 
3:  $\llbracket (dec\_n = true) \ \& \ (free\_n \leq 0) \rrbracket - >$ 
4:    $(used\_n' = used\_n - 1) \ \& \ (dec\_n' = false) \ \&$ 
5:    $(temp\_n' = temp\_n - 4) \ \&$ 
6:    $(remap\_tasks' = remap\_tasks + 1) \ \& \ (remap\_state' =$ 
7:    $0);$ 
7:  $\llbracket (inc\_n = true) \rrbracket - >$ 
8:    $(free\_n' = free\_n + 1) \ \& \ (sync' = 0) \ \& \ (inc\_n' = false);$ 
endmodule

```

For instance, we discovered certain combinations of these parametric values for which the TAPE algorithm would not be able to attain a stable system configuration from a random initial task mapping. Therefore, these results also indicate the shortcomings of ignoring randomness in the behavior of DTM schemes and the usefulness of probabilistic model checking in our proposed methodology.

D. RESULTS

We used the version 4.1 of PRISM model checker along with Windows 7 professional OS running on a core i5-3210 CPU at 2.50 GHz with 8.00 GB of RAM. The verification is done for a 3x3 and a 4x4 grid using probabilistic model checking. We compare our methodology to state-of-the-art [3], which is the latest formal verification technique for distributed DTM. We have implemented this state-of-the-art in SPIN in a reproducible way, and then performed the comparison.

Verification feasibility Analysis: After many iterative experiments, it was determined that at certain values of weight parameters the size of model grows more quickly with the increasing number of task. Fig.1 shows that at some values of a_s and a_b state-space explosion occurs whereas *the model is more scalable for verification when $a_s \approx a_b$.*

Verification of functional properties: The stable distribution of power budget and the absence of redundant trading can be ensured by verifying that *in long run stability equation will be satisfied.* We have to verify 12 and 24 such properties to cover all cores in the 3x3 and 4x4 grids, respectively. For example, the *stability property* for node 11 and 12 can be expressed as follows:

$$\text{Stability (p1112): } S=? \llbracket diff_11 - diff_12 < T_n \rrbracket$$

where $diff_11$ is the $sell_{T_n} - buy_{T_n}$ value of core number 11 and S is the steady state probability (i.e., probability of being in a state where stability equation holds true at time T , as T goes to infinity). We also compared the verification times for these properties with the verification times given in [3], as illustrated in Fig.2. Another important property in context of distributed DTM schemes is to ensure that temperatures does not exceed a given threshold and no thermal hotspots are formed. We calculated the probability that *in steady state the maximum temperature will be less than the threshold temperature* at different values of weight parameters.

$$\text{Thermal hotspot: } S=? \llbracket MaxTemp < 62 \rrbracket$$

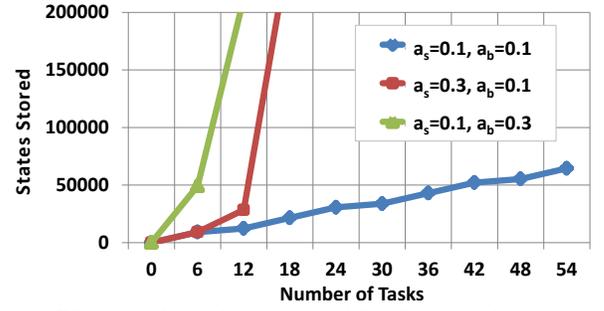


Fig. 1. Effect of tasks on the size of model at different values of a_s and a_b

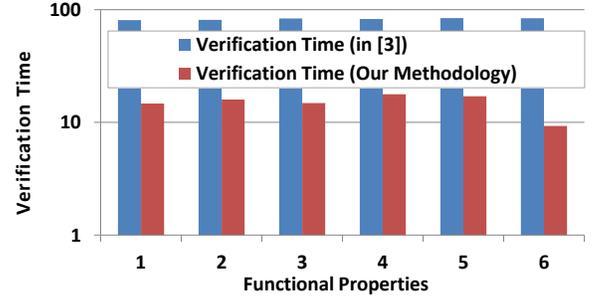


Fig. 2. Comparison of verification time with [3]

Optimal values of parameters and uniform distribution of temperatures: After many iterative experiments, we discovered the optimal set of parameters and it has been formally verified, using PRISM, that with the discovered combinations of parametric values, *the probability of maintaining lower temperatures is very high* (see Fig.4). Fig.4 represents the probability that temperature will remain within a certain threshold at different values of a_s and a_b . In Fig.4(a), the value of a_s is kept constant at 0.1 and value of a_b is varied from 0.1 to 0.9, whereas, in Fig.4(b) the value of a_b is kept constant at 0.1 and a_s is varied from 0.1 to 0.9. It was observed that the probability is higher when either $a_s \approx a_b$ or a_s is slightly greater than a_b . At the optimal values of weight parameters, *a uniform distribution of temperatures is achieved and no thermal hotspots are created.*

Back and forth trading: It was discovered that for some values of a_s and a_b , the algorithm would grow continuously without attaining a stable state. In reality, the algorithm *enters into a never ending loop due to redundant trading* (See 3).

The results of our verification shows that the performance of TAPE algorithm depends highly on the values of weight parameters. The above mentioned statistics provide very useful insights for the designer and indicates the usefulness of our approach.

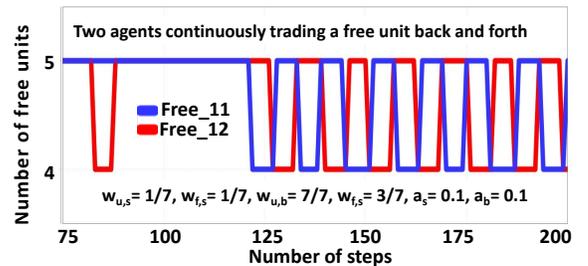


Fig. 3. Redundant Trading of power units

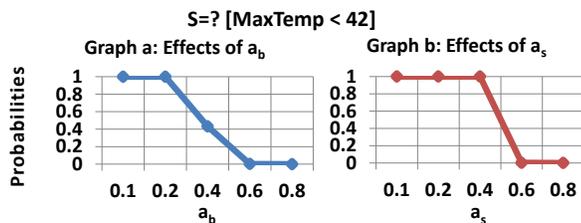


Fig. 4. Effects of a_s and a_b

IV. CONCLUSION

This paper presents a formal verification approach for distributed DTM schemes based on probabilistic model checking. The rigorous nature of the analysis coupled with the consideration of randomized behaviors of the distributed DTM schemes allows the identification of design problems that cannot be registered using traditional techniques. For illustration purposes, a successful probabilistic analysis of a state-of-the-art distributed DTM scheme is presented in this work. The identification of various loopholes in TAPE corroborates the benefits of our approach over other traditional techniques.

ACKNOWLEDGMENT

This work is supported in parts by the German Research Foundation (DFG) as part of the priority program “Dependable Embedded Systems” (SPP 1500 - spp1500.itec.kit.edu) and in parts by the DAAD “Deutsch-Pakistanische Forschungskooperationen” project.

REFERENCES

- [1] D. Dunn, “Intel delays montecito in roadmap shakeup”, EE Times, Manufacturing/Packaging, Oct. 27, 2005.
- [2] PRISM web site, www.prismmodelchecker.org, 2014.
- [3] M. Ismail, O. Hasan, T. Ebi, M. Shafique, J. Henkel, “Formal Verification of Distributed Dynamic Thermal Management”, IEEE International Conference on Computer-Aided Design, pages 248-255, 2013.
- [4] M.Kwiatkowska, G.Norman and D.Parker, “Advances and Challenges of Probabilistic Model Checking”, In Proc. Annual Allerton Conference on Communication, Control and Computing, pages 1691-1698, IEEE Press, 2010.
- [5] T. Ebi, M. Faruque, and J. Henkel, “Tape: Thermal-aware agentbased power economy multi/many-core architectures”, in International Conference on Computer-Aided Design, pages 302-309, 2009.
- [6] M.Kwiatkowska, G.Norman and D.Parker, “PRISM: Probabilistic Model Checking for Performance and Reliability Analysis”, ACM SIGMETRICS Performance Evaluation Review, 36(4), pages 40-45, 2009.
- [7] Y.Ge, Q.Qiu, Q.Wu, “A Multi-Agent Framework for Thermal Aware Task Migration in Many-Core Systems”, IEEE Trans. VLSI System, 20(10), pages 1758-1771, 2012.
- [8] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker, “PRISM: A tool for automatic verification of probabilistic systems”, in H.Hermanns and J. Palsberg, editors, Proc. 12th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, volume 3920 of LNCS, pages 441-444. Springer, 2006.
- [9] A. Lungu, P. Bose, D.J. Sorin, S. German, and G. Janssen, “Multicore power management: Ensuring robustness via early-stage formal verification”, in Formal Methods and Models for Co-Design, pages 78-87, 2009.
- [10] A. Biere, “Tutorial on model checking, modelling and verification in computer science” In Algebraic Biology, LNCS Vol 5147, pages 16-21, Springer-Verlag, 2008.
- [11] S.K. Shukla and R.K. Gupta, “A model checking approach to evaluating system level dynamic power management policies for embedded systems”, in Int. High-Level Design Validation and Test Workshop, Proc., pages 53-57. IEEE, 2001.
- [12] T. Herault, R. Lassaigne, F. Magniette, and S. Peyronnet, “Approximate Probabilistic Model Checking”, in Proceedings of Fifth International Conference on Verification, Model Checking and Abstract Interpretation, volume 2937 of LNCS, pages 73-84. Springer, 2004.
- [13] E.M. Clarke, W. Klieber, M. Novacek, and P. Zuliani, “Model checking and the state explosion problem”, in TPSV, volume 7682 of LNCS, pages 1-30, Springer, 2012.
- [14] B. Yun, K. Shin, and S. Wang, “Predicting thermal behavior for temperature management in time-critical multicore systems”, in IEEE Real-Time and Embedded Technology, 2013.
- [15] M.Kadin, S.Redda, and A. K. Uht, “Central vs. distributed dynamic thermal management for multi-core processors: which one is better?”, in ACM Great Lakes symposium on VLSI, page 137-140, 2009.