

Electro-thermal co-simulation of ICs with runtime back-annotation capability

András Timár, György Bognár, András Poppe, Márta Rencz
Budapest University of Technology and Economics
Department of Electron Devices
Budapest, Hungary 1111
Email: timarlbognarlpoppelrencz@eet.bme.hu

Abstract—This paper presents a novel approach to logical and thermal co-simulation of ASIC circuits. Numerous electro-thermal simulator implementations are present nowadays, but these simulators approach the electro-thermal simulation domain by co-simulating electronic and thermal effects at a low structural level. This approach has the advantage of being very accurate but at the expense of simulation time. In this paper we provide an alternative way to simulate standard cell ASIC circuits electrically and thermally in a concurrent process in real-time, in RTL level. Our approach takes standard cells of the digital design as basic building blocks and calculates a thermal distribution map on the surface of the virtual chip. The temperature map is calculated from the cells' power characteristics and the switching activity of the regularly working circuit. We call the presented approach *logi-thermal simulation*. An implementation of the method is also presented in this paper: a new simulation software, *LogiTherm* is under heavy development in the Department of Electron Devices at BME, Hungary.

Index Terms—logi-thermal simulation, electro-thermal simulation, temperature distribution

I. INTRODUCTION

As minimum feature size is continuously shrinking and power density growing, it is inevitable to take thermal effects into account when designing digital integrated circuits and manufacturing them. In this paper, we provide a novel methodology to determine the *hot-spots* on the digital integrated circuit's surface. Furthermore, our approach is able to calculate temperature map of a *packaged* semiconductor circuit excited by a logical simulator virtually. The thermal simulator can admit compact models of thermal devices, such as a heatsink combined with a fan.

Section II describes related researches in the field of electro-thermal simulations as well as many- and multi-core design studies that deal with temperature issues of such circuits. Section III describes the necessity of RTL level electro-thermal co-simulation. In section V the implementation of the proposed methodology is introduced. Section VI explains power characterization of the design kit cells while section VII shows a reasonable method to simplify the power database. Section VIII demonstrates the results of the presented methodology. Future work and conceptions are discussed in section X.

II. RELATED WORK

The authors of [1] propose a methodology that constitutes the basic idea of this paper. It deals with temperature map

generation of ICs from digital simulations. The method presented in this paper improves that methodology by allowing in-situ logic and thermal co-simulation of an IC design. The work presented in [1] lacks the possibility of back-annotating temperature-dependent delays into the running simulation. Dealing with temperature-dependent delays become more and more important as feature size is shrinking and power density is growing. Failure to take the temperature-dependent delays into account can cause timing (e.g. setup- and hold time) violations that can foil correct logic operation. The methodology presented in this paper also improves the work of Torki et al by coupling the logic and thermal simulator engines with a custom controller and visualization application that can evaluate logical and thermal calculations and prepare delay back-annotation on-the-fly, in the middle of the simulation.

Present electro-thermal simulators approach the electro-thermal simulation problem either by FEM simulation, relaxation method or simultaneous iteration [2]. A brief overview of an electro-thermal simulation method based on simultaneous iteration is presented in [3] and [4]. A method of electro-thermal simulation using simulator coupling is presented in [5]. The technique is based on the coupling of a FEM program with a circuit simulator. Device level electro-thermal simulation of analog circuits and the logical gate level logi-thermal simulation of digital circuits are addressed in [6]. A work of A. Poppe et al. is presented in [7] that gives an overview of different approaches to die level electro-thermal simulation where simulator coupling or the so called direct method (co-simulation of the electrical and thermal parts within a single tool) are used. A pre-RTL temperature-aware design methodology is presented in [8], where a fast, yet accurate architectural thermal model that is able to explore large regions of the design space is proposed. [9] attempts to show that there is a significant peak temperature reduction potential in managing lateral heat spreading through floorplanning. As a demonstration, it uses a wire delay model and floorplanning algorithm based on simulated annealing to present a profile-driven, thermal-aware floorplanning scheme that significantly reduces peak temperature with minimal performance impact that is quite competitive with Dynamic Thermal Management (DTM). The floorplanning tool *HotFloorplan* is part of the *HotSpot* software [10] that is developed at DCS, University of Virginia. In addition, simulating the self-heating of the circuit

in the early phase of the design before manufacture would make cooling issues less problematic. Self-heating simulations may also eliminate the need for design back-annotation after manufacture. An example of temperature-aware ASIC design flow can be found in [11].

Thermal issues in today's many- and multi-core designs became a primary concern. Multi-core exacerbates thermal challenges because power scales with the number of cores, but also creates new opportunities for temperature-aware design, because multi-core designs offer more design parameters than single-core designs. [12] investigates the relationship between core size and on-chip hot spot temperature and shows that with the same power density, smaller cores are cooler than larger cores due to a spatial low-pass filtering effect of temperature. [13] explores the thermal impact on manycore processor architecture and evaluates its performance. Preliminary results show that thermal constraints reduce performance as expected, but also make performance almost insensitive to the complexity of the primary core across a diverse degrees of parallelism, which greatly reduces design complexity.

III. RTL LEVEL SIMULATION

The methodology presented in this paper takes the standard cells of a digital ASIC design as basic primitives rather than the contained transistors and wiring. This property clearly distinguishes it from today's electro-thermal simulators. The logic simulator drives these primitive cells with a user-defined stimuli and logic waveform results can be observed. By standard cells being the building blocks of the simulation, it is apparent that the thermal resolution is coarser than that of an electro-thermal simulator but it is also faster. However, with reasonable and deliberate cell simplification an optimal trade-off between simulation accuracy and time can be achieved. For example, if we further partitioned the initial cell into smaller blocks, transistor-level electro-thermal accuracy could be approximated.

The presented method is based on the so called *relaxation method*, where two simulator engines—an electric and a thermal—are coupled together. The input of each simulator comes from the other simulator and they together produce the final simulation result.

IV. HOT-SPOT DETECTION AND TEMPERATURE MAP DETERMINATION

By calculating the temperature distribution of the circuit, local and global hot-spots can be detected. Hot-spot detection is crucial in today's ultra large scale integrated circuits as functional operation may completely fail due to unanticipated heating of devices at certain coordinates. With our presented methodology and application bundle, circuit designers can create temperature-aware IC designs where thermal issues emerge during the design phase rather than after manufacture.

The simulation stimuli resembles the real-world operation of the circuit thus the fully packaged design can be evaluated for defects or misfunctionality far before manufacture.

V. IMPLEMENTATION

The presented methodology is based on standard tools such as Verilog hardware description language and its Programming Language Interface (PLI), thus it can be integrated into any logical simulation flow easily. The majority of today's logic simulators (both commercial and open-source licensed) support this language and interface.

The *LogiTherm* simulation software developed at the Department of Electron Devices in BME, Hungary acts as a glue logic between the standard logic simulator and the thermal simulator. The *Therman* solver engine [14] was used as the thermal simulator that was also developed in the Department.

This glue logic interface application controls communication between the logical and the thermal simulator engines and displays the results in real-time. Real-time here means that the simulation results get displayed right after each simulation timestep, thus providing the user immediate graphical feedback of the result. In Fig. 1 the design flow of the logi-thermal simulation is shown.

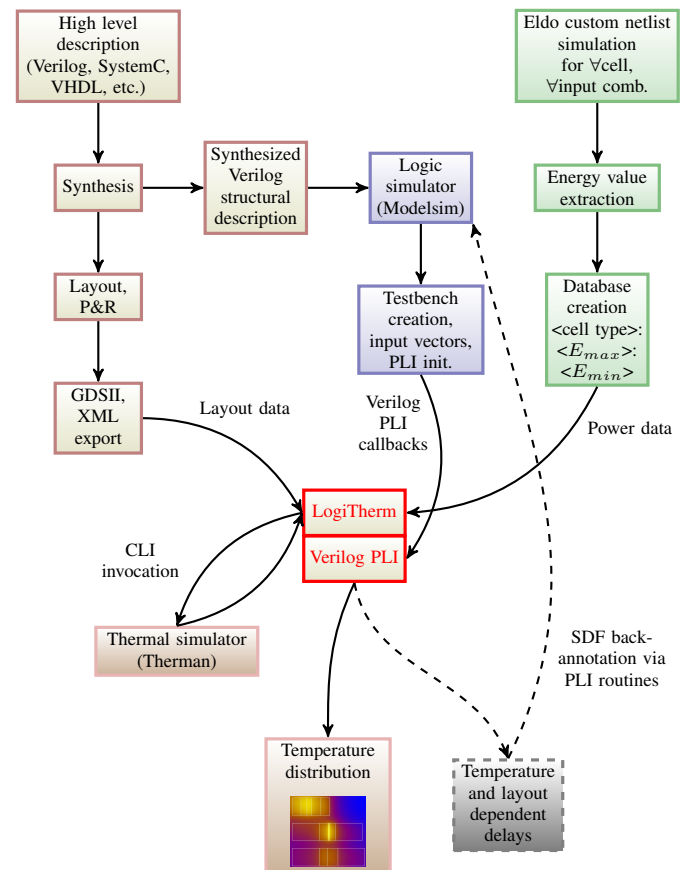


Figure 1. Logi-thermal design flow

The logi-thermal design flow starts with two concurrent processes. On one hand, the IC designer team describes the logic behaviour of the application in a high level language like SystemC, SystemVerilog, SystemVHDL, etc. Then the process continues with an iterative simulation cycle where the majority of the logic and timing errors can be detected

and the design can be corrected. After successful logic and timing simulation, synthesis takes place, where logical RTL gates get mapped to the physical gates of the chosen process design kit (PDK). This synthesized description can then be simulated for correct timing and the floorplan can be created. Simultaneously, the cell power characterization has to be done. Often the process design kit comes with Liberty or TLF files which contain data about the cells' timing, delay and power characteristics. Unfortunately, sometimes this is not the case. The PDK may be missing these files, or just on type of them, or the necessary power characteristics for the cells. Therefore we developed a rather simple method to obtain the power characteristics for the cells. Details of the characterization process are shown in section VI. This process can be done independently of the high level design and synthesis.

The synthesized HDL description is stimulated with a logic simulator with test vectors and switching activity from the results can be determined for each cell. From the Place&Routed floorplan, the switching activity and the power characteristics of the cells the LogiTherm application can calculate the temperature distribution in the design. This is done concurrently with the logic simulation thus immediate graphical results show the evolved hot-spots.

Temperature-dependent delays could also be calculated and back-annotated to the simulation via Standard Delay Format (SDF) files. Delay back-annotation can take place before each timestep of the simulator.

VI. POWER CHARACTERIZATION

Power characterization is made using an analog simulator that runs transistor-level simulations of the cells in the library. The power characterization of the standard cells has to be done only once per process design kit. The methodology presented involves using the *Eldo* analog circuit simulator engine from Mentor Graphics. A custom *Eldo* netlist is generated according to the simulated cell library and all of the cells of the library are simulated with every possible input combination.

As the methodology focuses mainly on simulating CMOS circuits, the great majority of the power consumption (and thus the heating) occurs when a switching activity takes place (hence the name "dynamic power"). During the initial analog simulation with *Eldo*, dissipated energies during every possible switching event in a cell are collected and stored for every cell type. The energies dissipated during a logic transition are calculated using (1):

$$\varepsilon = \int_{T_1}^{T_2} P(t) dt, \quad (1)$$

where $P(t)$ is the total power consumption function of the cell after every possible logic transition. Fig. 2 shows a typical input voltage–power consumption characteristic. The energy dissipation on a logic transition is calculated with the following algorithm:

- 1) The *vicinity* of the logic transition has to be found where the great majority of the power consumption lies. The vicinity is found by the following steps:
 - a) Search for the maximum point of the $P(t)$ power consumption function in the adjacent interval of $\pm \frac{Period}{2}$ around the logic transition. In CMOS circuits, this is an appropriate interval because the significant peaks of the power function $P(t)$ are near the switching event. Far away from the switching events is only static consumption which is negligible in contrast to the consumption at logic transitions. *Period* here means the time between any two subsequent input logic transitions.
 - b) Find the time instances where the $P(t)$ curve crosses 1% of the local maximum. These points are T_1 and T_2 .
- 2) Integrate the $P(t)$ curve in the $[T_1, T_2]$ interval. This value (ε) is the average energy consumption belonging to that particular logic transition.

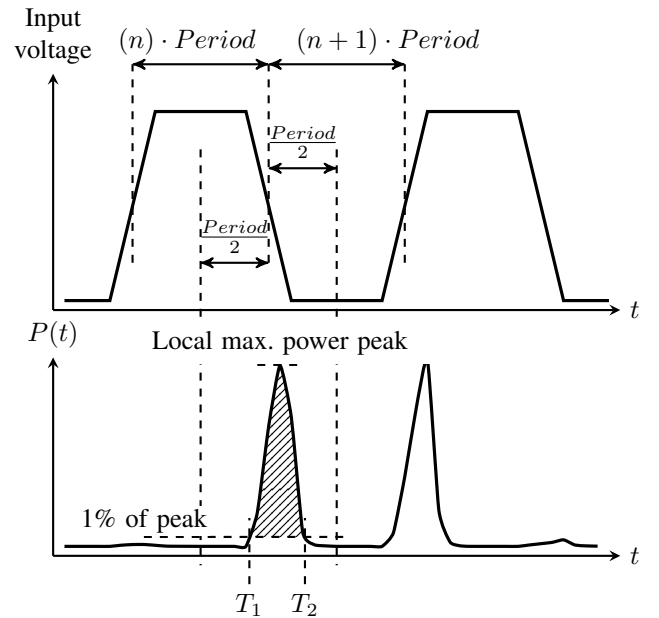


Figure 2. Calculation of dissipated energy per logic transition

The $P(t)$ function is calculated with *Eldo*'s `POW(Xinstance_name)` function which calculates real power dissipated in a cell and ignores power stored in capacitances. This behaviour is required because only real dissipated power must be taken into account during the thermal simulation. The measured power is exactly equivalent to the power dissipated by the transistors of the measured instance, thus it is the power that causes heating of the components in the cell.

A standard logic simulator (e.g. Modelsim from Mentor Graphics) can be complemented with a thermal simulator engine (e.g. Therman [14] or FireBolt from Gradient Design Automation [15]) that calculates the temperature distribution

of every cell instance in the layout according to the switching activity provided by the logic simulator. LogiTherm conducts a thermal simulation on the actual layout with the actual power density based on the switching activity that is extracted from the digital simulator. This approach clearly realizes the relaxation method, where two distinct simulator engines (a logical and a thermal engine) produce results that the other utilizes.

If we consider the initial Eldo power simulation, it turns out that the resulting database would be very large. Taking every cell of the library into account as well as every possible input combination for every cell can lead to huge data that must be processed every time logic-to-thermal simulator transition occurs. In order to maintain fast simulations without losing accuracy, it is necessary to reduce the size of the database. Simulations showed, that when both the cell's input(s) and output(s) changed (e.g. either input of a 2-input OR gate changes to logic "HIGH"), the power characteristic ($P(t)$) has *much larger* peaks than in the case when only input(s) changed but no output(s) (only internal logic changes were present). This is because the cell's output(s) has to drive the input capacitances of the load gates. These capacitances are usually larger than the internal gate capacitances and also scale by the number of driven cells (fan-out of the driving cell). When neither input(s) nor output(s) change, the cell has only static power consumption, which—initially—can be neglected.

VII. DATABASE SIMPLIFICATION

As stated previously in section VI, the resulting database consisting of every cell's energy consumption for every possible input transition would be very large. To reduce the size of the database but maintain simulation accuracy, a reasonable simplification is needed. If the logi-thermal simulator had to acquire energy consumption values for every cell for every input transition from a very large database each time a thermal simulation was needed, simulation time would be very long and comparable to a FEM electro-thermal simulation. A reasonable simplification would be to differentiate three states of operation per cell and assign certain energy values to these states:

- 1) *Input* and *output* logic transition occurred.
- 2) Only *input* logic transition occurred, output did not change (internal switching).
- 3) No logic transition is detected on any of the inputs or outputs (static power consumption).

The energy values assigned to these states:

- 1) When collecting energy consumption values for each cell (see section VI) the largest energy value is selected. Let this be E_{max} .
- 2) Let the smallest energy consumption value in the list be E_{min} .
- 3) Let the average energy dissipated when no switching activity be E_{static} .

In case 1), the maximal energy value (E_{max}) is assigned to the given transition. Thus when the logic simulator detects

both input and output logic transition in a cell, it adds the E_{max} value to the total energy accumulator.

In case 2), the minimal energy value (E_{min}) is assigned to the given transition. This means when the simulator detects only input logic transition but no corresponding output change in the cell, it adds E_{min} to the total energy accumulator.

The current version of the simulator simply does not take static power consumption into account as it is negligible compared to the dynamic power consumption.

The thermal simulator is invoked periodically with a time period comparable to the typical thermal time constant of the physical structure. This time period can be determined initially with a thermal simulation software. The typical thermal time constants of such structures is approximately 3 orders greater than the clock frequency of the measured circuit ($T_{CLK} = 1$ ns), thus reasonable averaging can be used. The simulator when used in thermal DC simulation mode, takes dissipating shapes and dissipated power per shape as arguments. Each shape belongs to a cell instance in the digital standard cell IC layout. The dissipated *power* per cell is passed to the thermal simulator using the following methodology: when a logic transition occurs, the supplemented logic simulator accumulates the dissipated energies as simulation time increases. This accumulated *energy* is then divided by the elapsed simulation time as (2) shows.

$$P_{average} = \frac{1}{T_{elapsed}} \varepsilon. \quad (2)$$

$P_{average}$ power is measured for every cell in the design and the thermal simulator is invoked with these power values. This model implicitly includes the cooling behaviour of the structures. When no switching activity for a given cell takes place, the accumulated energy remains the same but simulation time increases, thus the same energy value is divided by a continuously growing time value. This means that average power dissipation (the thermal simulator invoked with) decreases continuously and converges to zero. In the resulting thermal map, the corresponding area will be cooler after every iteration.

VIII. RESULTS

The simulated test-case was a 64-bit counter implemented on a TSMC 0.35 μm process. The total chip area is $A = 1.96 \times 2.51$ mm. Thickness of the silicon chip is $d = 0.3$ mm. The maximal ΔT according to the simulation resulted to 0.62°C. The total power dissipation across the surface is $\Delta P = 0.5952$ W. Maximal ΔT was measured in a D flip-flop output of the counter corresponding to the least significant bit (LSB). This is evident as switching density is maximal here because of the behaviour of a counter.

Fig. 3 shows a screenshot of the LogiTherm application after the first few nanoseconds of the simulation time where steady-state temperature distribution is not yet achieved.

IX. SUMMARY

In this paper a novel methodology for logical and thermal co-simulation was presented. A distinct feature of the method

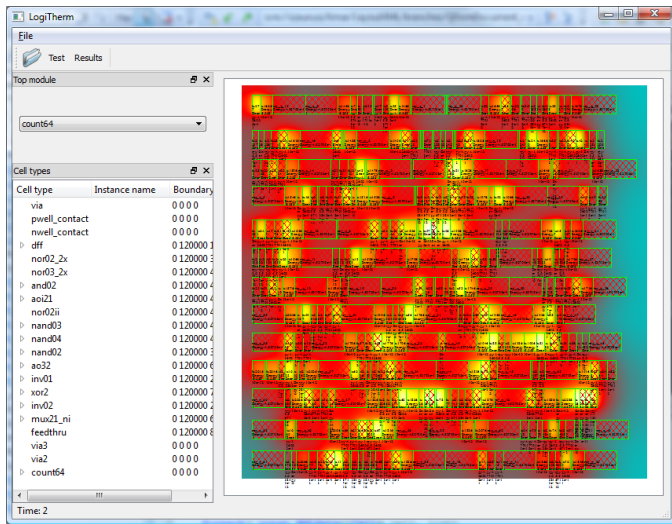


Figure 3. LogiTherm application showing actual temperature map of simulated IC

is that it simulates the digital design in RTL level. The methodology can be included in any traditional digital IC design flow. A fast logi-thermal simulator application was developed in the Department of Electron Devices in BME called LogiTherm that is able to determine temperature map of the circuit and communicates with industry tools via standard interfaces. By generating a temperature map of the circuit during operation, local and global hot-spots on the chip can be identified and considered. Designs of fully packaged circuits can also be simulated with the presented methodology.

A simplified energy modelling technique was also presented that produces very compact database while providing appropriate accuracy. The small energy database enables the logi-thermal simulator to calculate thermal distribution on the IC's surface in real-time, in parallel with the logic simulator. The thermal distribution is determined according to the input stimuli that simulates the circuit's real behaviour after manufacture.

X. FUTURE WORK

Our team plans to improve the methodology and the interfacing application with the ability to calculate and back-annotate delays resulting from the thermal behaviour of the circuit. Preliminary studies showed that delays might be back-annotated to the logic simulator just-in-time. This means that after each simulation timestep and logic-thermal calculation cycle, the resulting delay changes could be incorporated immediately to the logic simulation and the simulator could calculate the next timestep with the refreshed delay values. This back annotation would take place in every simulation timestep. This kind of real-time delay back-annotation feature is supported by the Verilog PLI interface.

Further optimization of the cell power characteristics database can be made. At present, the power data are stored in a custom-format database file that gets processed by the

LogiTherm application. The format of the database could also be standardized by using an SQLite database. This database could also be generated on-the-fly by the Eldo analog simulator and no further processing of the analog Eldo results would be necessary. This may speed up the cell characterization process as well as eliminate the need for a post-processing application.

By improving our methodology, thermal-aware floorplanning could be possible where preliminary logi-thermal simulations showed the critical hot-spots of the circuit and floorplanning tools could be controlled to take these effects into account and compensate them by varying cell placement.

XI. ACKNOWLEDGEMENT

This work is supported partly by the Hungarian National Office for Research and Technology through the project HiDRaLoN (CA301) of EUREKA's CATRENE (4140) Cluster, partly by the IP 248603 THERMINATOR FW7 project of the European Union.

REFERENCES

- [1] K. Torki and F. Ciontu, "Ic thermal map from digital and thermal simulations," in *Proceedings of the 8th Thermic Workshop*, Madrid, October 2002, pp. 303–308.
- [2] V. Székely, A. Poppe, M. Rencz, A. Csendes, and A. Páhi, "Electro-thermal simulation: a realization by simultaneous iteration," *Analog Integrated Circuits and Signal Processing*, vol. 28, 1997.
- [3] V. Székely, A. Páhi, A. Poppe, and M. Rencz, "Electro-thermal simulation with the SSSI package," *Microelectronics Journal*, vol. 21, pp. 21–31, 1999.
- [4] V. Székely, A. Poppe, A. Páhi, A. Csendes, G. Hajas, and M. Rencz, "Electro-thermal and logi-thermal simulation of VLSI designs," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 5, no. 3, sep 1997.
- [5] S. Wünsche, C. Clauß, P. Schwarz, and F. Winkler, "Electro-thermal circuit simulation using simulator coupling," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 5, no. 3, sep 1997.
- [6] M. Rencz, V. Székely, A. Poppe, K. Torki, and B. Courtois, "Electro-thermal simulation for the prediction of chip operation within the package," *19th IEEE SEMI-THERM Symposium*, 2003.
- [7] A. Poppe, G. Horváth, G. Nagy, M. Rencz, and V. Székely, "Electro-thermal and logi-thermal simulators aimed at the temperature-aware design of complex integrated circuits," *24th IEEE SEMI-THERM Symposium*, 2008.
- [8] W. Huang, K. Sankaranarayanan, K. Skadron, R. J. Ribando, and M. R. Stan, "Accurate pre-RTL temperature-aware design using a parameterized, geometric thermal model," *IEEE Transactions on Computers*, vol. 57, no. 8, August 2008.
- [9] K. Sankaranarayanan, S. Velusamy, M. Stan, and K. Skadron, "A case for thermal-aware floorplanning at the microarchitectural level," *Journal of Instruction-Level Parallelism*, vol. 8, no. 1-16, 2005.
- [10] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. R. Stan, "Hotspot: A compact thermal modeling methodology for early-stage vlsi design," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 14, no. 5, May 2005.
- [11] W. Huang, M. R. Stan, K. Skadron, K. Sankaranarayanan, S. Ghosh, and S. Velusamy, "Compact thermal modeling for temperature-aware design," *41st Design Automation Conference (DAC)*, San Diego, CA, June 2004.
- [12] W. Huang, M. R. Stan, K. Sankaranarayanan, R. J. Ribando, and K. Skadron, "Many-core design from a thermal perspective," *Proceedings of the 45th ACM/IEEE Conference on Design Automation (DAC)*, June 2008.
- [13] W. Huang, K. Skadron, S. Gurumurthi, R. J. Ribando, and M. R. Stan, "Exploring the thermal impact on manycore processor performance," *26th IEEE SEMI-THERM Symposium*, February 2010.

- [14] V. Székely, A. Poppe, M. Rencz, M. Rosental, and T. Teszére, "Therman: a thermal simulation tool for ic chips, microstructures and pw boards," *Microelectronics Reliability*, vol. 40, no. 3, pp. 517–524, 2000.
- [15] (2010, 19th April, 15:57) FireBolt and CircuitFire. [Online]. Available: <http://www.gradient-da.com>