



Temperature modeling and emulation of an ASIC temperature monitor system for Tightly-Coupled Processor Arrays (TCPAs)

E. Glocker¹, S. Boppu², Q. Chen³, U. Schlichtmann³, J. Teich², and D. Schmitt-Landsiedel¹

¹Lehrstuhl für Technische Elektronik, Technische Universität München (TUM), Theresienstr. 90, 80333 München, Germany

²Lehrstuhl für Informatik 12 (Hardware-Software-Co-Design), Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Cauerstr. 11, 91058 Erlangen, Germany

³Lehrstuhl für Entwurfsautomatisierung, Technische Universität München (TUM), Arcisstr. 21, 80333 München, Germany

Correspondence to: E. Glocker (elisabeth.glocker@tum.de)

Received: 12 December 2013 – Revised: 17 March 2014 – Accepted: 27 March 2014 – Published: 10 November 2014

Abstract. This contribution provides an approach for emulating the behaviour of an ASIC temperature monitoring system (TMon) during run-time for a tightly-coupled processor array (TCPA) of a heterogeneous invasive multi-tile architecture to be used for FPGA prototyping. It is based on a thermal RC modeling approach. Also different usage scenarios of TCPA are analyzed and compared.

1 Introduction

The demand for increasing performance and at the same time decreasing feature size in modern CMOS technologies lead to higher power densities influencing chip temperature. The chip temperature also depends on thermal conductivity of chip and package. High temperature reduces circuit speed and causes transistor degradation. Normally the workload is distributed in a non-uniform way in different circuit blocks leading to local temperature hot spots. Altogether, this can greatly affect the lifetime and the reliability of a chip (Brooks et al., 2007; Semenov et al., 2006). This problem becomes even more critical in multi-core scenarios: Core temperature increases, because it depends not only on its own power density, but also on the activity and the power density of neighbour cores. Monitoring of operating temperature during run-time makes it possible to control and limit it to increase lifetime and reliability and to find the best suitable application-core pairs during resource allocation.

This paper presents an approach for emulating an ASIC temperature monitoring system (TMon) for Tightly-Coupled Processor Array (TCPA) tiles of a heterogeneous invasive

multi-tile architecture to be used for FPGA prototyping. Before the emulating principle is shown, TCPA usage scenarios are simulated and analysed in Sect. 3 with detailed temperature simulations based on thermal RC modeling approach for an ASIC layout. Section 4 shows then the emulation principle and first results for using this method.

1.1 Invasive computing, temperature emulation

Static and central management of computing systems may not scale, since SoC architectures with thousand or even more processors on a single chip can be foreseen in the near future. Also, imperfections like parameter variations and ageing lead to increasingly unreliable components. Invasive computing (Teich et al., 2011) is a resource-aware parallel computing paradigm that provides a self-organizing task management to deal with this: Applications get the ability to explore the system and dynamically spread onto processors in the “invasion” phase and also give back unused resources depending on the current needs. The system is a heterogeneous invasive multi-tile architecture that includes computation, I/O and memory tiles. Computation tiles are either TCPA or RISC tiles that may have different number and (for RISC tiles) types of processors. The individual tiles are connected via a network-on-chip (NoC). Resource allocation considers application constraints such as required number and type of resources including physical hardware properties like temperature. The run-time support system and an underlying hardware management unit handle the resource allocation in close-collaboration. A closed-loop control system is formed, including monitors, run-time support system, hardware management unit and applications.

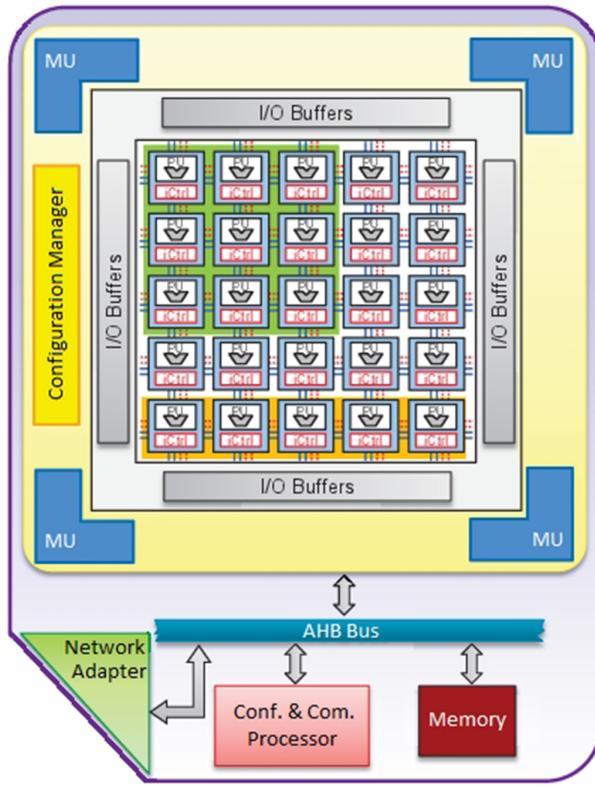


Figure 1. TCPA tile with a 5×5 processing elements (PEs) array. Initiation of an invasion may be done from the corner PEs (master PEs).

In preparation of an ASIC implementation, FPGA prototyping is required for tuning design parameters and design verification. However, monitor modules with analog-circuit components, like temperature monitors, cannot be directly implemented on prefabricated FPGA structures. Also, since power dissipation and the evolution of temperature on FPGA differ from that in an ASIC implementation, the FPGA will show very different behaviour compared to the envisioned ASIC. Nevertheless, the interaction of hardware, run-time support system and application during resource allocation, taking monitoring data into account, has to be verified and tested. Thus, an emulation of the behaviour of an ASIC TMon on the FPGA platform is required.

1.2 Related work

Exploration of thermal behaviour is done offline either with detailed simulators like in Skadron et al. (2004) or virtual platforms (e.g. Bartolini et al., 2013). Online thermal exploration during runtime for FPGA prototyping relies on existing temperature sensors as done in Shen and Qiu (2011) that do not emulate the ASIC temperature. Also coupled hardware/software solutions with e.g. a host-computer running a thermal modeling tool like in Atienza et al. (2006) are available that could be much slower than the prototyped system.

Therefore, a solution as fast as real temperature monitors is required. Nevertheless related work dealing with offline detailed ASIC temperature simulations (e.g. Skadron et al., 2004) can be adapted for temperature monitor emulation.

2 Tightly-Coupled Processor Arrays (TCPAs)

A TCPA is a massively parallel architecture containing tightly-coupled light-weight processing elements (PEs) (Kissler et al., 2006). An example TCPA tile (array size of 5×5 PEs) is shown in Fig. 1. The invasion of new applications always starts at the corner PEs (master PEs). Each PE has an invasion controller (iCtrl), a dedicated hardware component for invasion. Each TCPA tile contains a “Conf. & Com. Processor” managing the communication and configuration of regions of claimed PEs with proper programs. Reconfigurable I/O Buffers, tile local Memory, a Configuration Manager and management units (MU) are also present. In the TCPA tile one MU is present at each corner of the array to supervise invasion requests and application execution on the master PEs. Communication outside of the TCPA tile is handled with a Network Adapter (Hannig et al., 2013, 2014).

The amount of instruction memory and control overhead for individual PEs is kept as small as possible. Thus a tight coupling, cycle-based communication by point-to-point connections among PEs (from neighbour to neighbour) is realized. Individual PEs do not have direct global memory access (Lari et al., 2012): Data transfer is realized to and from the array through the border PEs that have a connection to the banks of the surrounding memory buffers. The number of invaded PEs can be dynamically reduced and increased depending on the requirements of the corresponding application with different invasion strategies (Muddasani et al., 2012). TCPAs are typically used as hardware accelerators in an MPSoC for execution of computationally intensive loop programs like streaming applications (video/image processing) (Muddasani et al., 2012).

3 ASIC temperature simulation for TCPAs

From a temperature point of view, it is feasible to carefully analyse and optimise the placements of active PEs. By using temperature monitor data, it is possible to control the temperature by load-balancing (e.g. proper resource allocation) under different control targets: One control target might be to keep the temperature as low as possible. Obviously the temperature should be limited to avoid overheating. An even lower temperature is desirable to reduce the effect of PE ageing over time and thus increase PEs lifetime and reliability. Another control target might be to keep the temperature difference inside the array low to ensure similarly low PE ageing. Proper resource allocation can be used to outbalance high temperatures in the past with low ones in the present.

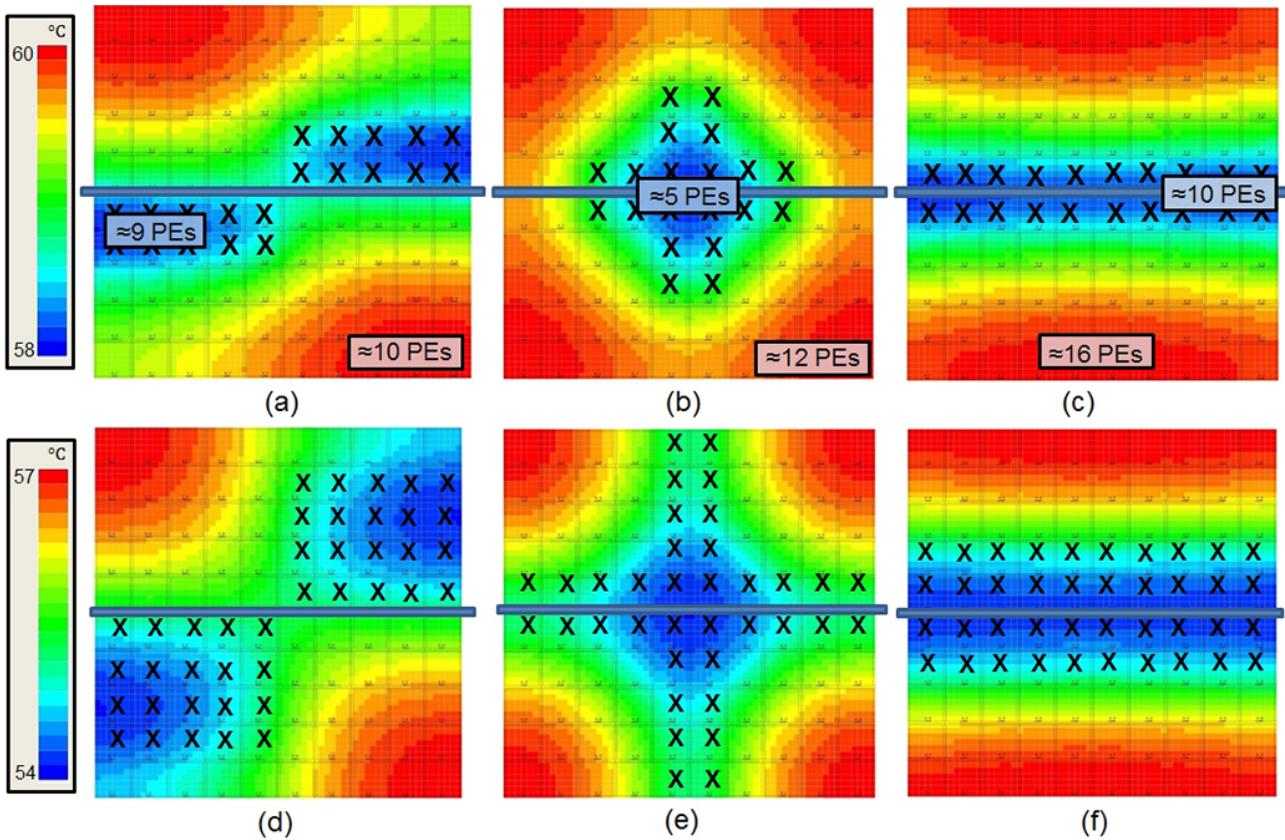


Figure 2. TCPA usage scenario of 80 % in (a)–(c) and 60 % usage in (d)–(f): Two applications are running on the TCPA (border between the applications: straight line). Each of them has 40 invaded PEs for 80 % and 30 for 60 % usage (unused PEs are marked with an ‘X’). Different positions of the unused PEs lead to different temperature evolutions. For (a)–(c) the size of the area affected with hottest and coolest temperatures is given (unit is PEs).

3.1 Simulation framework

Temperature simulations, based on a thermal RC model, are performed for steady-state temperatures with HotSpot (Huang et al., 2004). More details on the thermal RC model are given in Glocker and Schmitt-Landsiedel (2013). Steady-state temperature means that (with a constant load) equilibrium temperature distribution is reached. That means the temperature does not depend on time anymore. The steady-state temperature is a good estimation, if the task execution time is large enough (Brooks et al., 2007). TCPAs are typically used for intensive loop programs, so the execution time for applications is typically very long. The accuracy of HotSpot has been intensively studied: For considered floor-plans, it showed good agreement with finite-element simulator and test-chip results where the deviations for steady-state temperatures was found to be always less than 5.8 % (Huang et al., 2008; Skadron et al., 2004).

In this section simulations are done for an ASIC TCPA layout with an array size of 10×10 PEs and they are done for a size of 5×5 PEs in Sect. 4. Each PE contains a power consuming region (about 83 % of total PE size) and a region

that does not consume power (wiring, space in between PEs) of about 17 % of total PE size. Those individual PE blocks are placed next to each other to form the desired PE array. Possible invasion strategies are linear, meander walk, rectangular and random fashion as explained in Lari et al. (2012). It is assumed that each PE is either used (constant power consumption) or not used (just very small leakage power consumption). The TCPA tile is designed in TSMC 65 nm. Ambient temperature is 45 °C and chip thickness is 500 μm . Convection capacitance is 140.4 JK^{-1} , convection resistance is 0.1 KW^{-1} , silicon thermal conductivity is $100 \text{ Wm}^{-1} \text{ K}^{-1}$ and silicon specific heat is $1.75 \text{ MJm}^{-3} \text{ K}^{-1}$ (Skadron et al., 2004).

3.2 Evaluation of TCPA usage scenarios

If the complete 10×10 TCPA array is used (full-use scenario), the resulting temperature is approximately constant at 63 °C (highest temperature in tile middle 62.8 °C, lowest temperature at tile corners 62.6 °C). This corresponds to the temperature worst-case scenario.

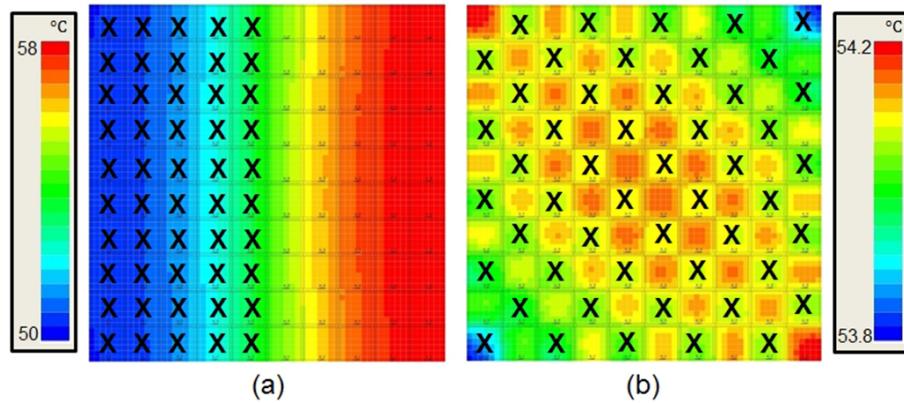


Figure 3. A 50 % usage scenario for one application is possible by using half of the array (PEs marked with an ‘X’ are unused) or by using the complete array (PEs marked with an ‘X’ are unused by the application itself, but still need to transfer data from one active PE to the next).

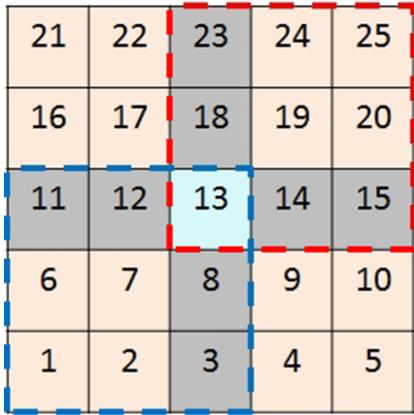


Figure 4. The emulated TMon gives out maximum temperatures for four 9-PE region temperatures (the dashed lines show two of those 9-PE regions. The other two are found accordingly). For finding all four 9-PE region temperature, the TCPA tile has to be partitioned further into four 4-PE regions at the corners (light red background colour), a 1-PE region in the middle (light blue background colour) and four 2-PE regions in between (grey background colour).

A typical usage scenario (standard scenario) is 80 % usage as shown in Fig. 2a–c for two applications with 40 invaded PEs each. For Fig. 2(a) the improvement compared to full-use scenario is 4 % for hottest and 7 % for coolest temperature. Different positions of the unused PEs lead to different temperature evolutions (Fig. 2b and 2c). The temperature difference compared to Fig. 2a is small: For hottest temperature it is 0.3 °C for both. For coolest temperature it is 0.7 °C for Fig. 2b and no improvement for Fig. 2c. But the area size affected by hottest and coolest temperatures change: In Fig. 2a the area affected by hottest temperature has the size of about 10 PEs. For Fig. 2b it has the size of about 12 PEs and in Fig. 2c of about 16 PEs (worst case, increase of 6 PEs compared to Fig. 2a). The area affected by coolest temperature has a size of about 9 PEs in Fig. 2a, about 10 PEs in Fig. 2c

and about 5 PEs in Fig. 2b (best case, decrease of 4 PEs compared to Fig. 2a).

If the usage decreases to 60 % as shown in Fig. 2d–f, the changing size of the area with hottest respectively coolest temperature is still visible for different scenarios. For Fig. 2d the temperature improvement compared to full-use scenario is 8 % for hottest and 14 % for coolest temperature. Compared to Fig. 2d, the temperature improvement for hottest temperature is 2 °C for both Fig. 2e and f. The coolest temperature is worsen by –1 °C for both.

A 50 % usage scenario for one application is shown in Fig. 3. Either one half of the TCPA array can be used (Fig. 3a). If it is known, that the other half of the array is not needed for some time, it may be possible to use the complete array for the application (Fig. 3b). The latter scenario is better for two reasons: The highest temperature is lowered by 6 % and the temperature distribution is more uniform. Note, that those PEs not directly used by the application, still need to transfer data from one active PE to the next and would consume an insignificant amount of power.

4 Emulation of TCPA temperature monitor system

The basic idea behind the emulation of the TMon is based on the fact that a PE’s temperature depends on the PEs own activity and the activities of neighbour PEs. If we can determine the temperature based on PE’s own activity and the influence of neighbouring PEs on this temperature, then it is possible to find the resulting PE temperature.

4.1 Emulation principle

For emulation, a TCPA array size of 5×5 was chosen. Because of the high number of PEs in a TCPA tile, the TMon does not give out each PE’s temperature, but the temperature for regions of 9 PEs (9-PE region temperature) as shown in Fig. 4 for two out of four possible 9-PE regions with

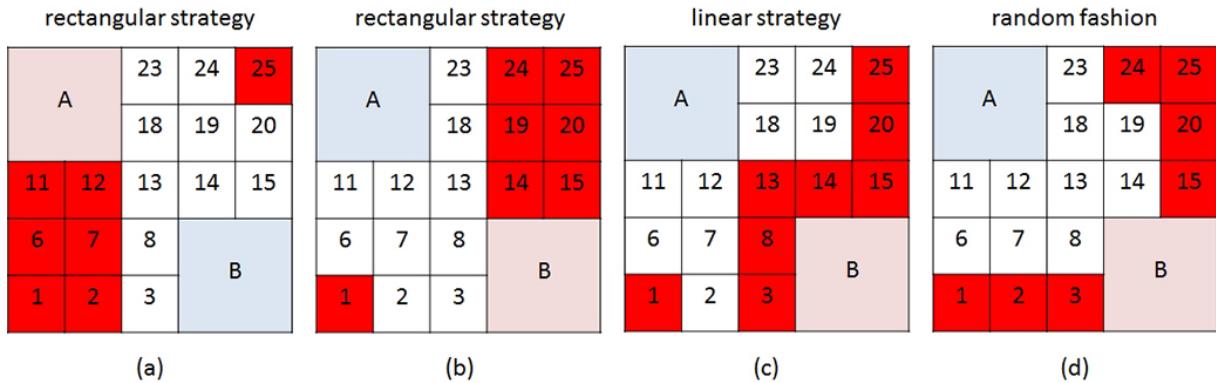


Figure 5. By using the emulated TMon, best suitable PEs for new applications can be found. The dark red colored PEs are currently used. A new application must be started in either region A or B. The light red and blue coloring of region A resp. B show the hotter or cooler region because of the activity of the currently used dark red colored PEs. The new application should always start in light blue colored region.

dashed lines. The region size of 9 PEs is feasible since invasion can start only at one of the four master PEs in the corners. By using the four 9-PE region temperatures the best suitable master PE (from temperature point of view) can be determined. A standard usage scenario would use 80 % of the PEs. For a 5×5 array with 4 applications running, 80 % means 5 PEs for each application. So information about the 9-PE region of each corner is normally needed to find a big enough region with adequate PEs for an application and adequate PEs around those application PEs. The output of the emulated monitor system are maximum temperatures. Besides resource allocation of new applications, the temperature monitor data can be used during application runtime: If an application wants to invade additional PEs, the 9-PE region temperatures can be used to find a suitable direction (in terms of temperature evolution) for the invasion.

To find the 9-PE region temperatures, the TCPA tile is partitioned further as shown in Fig. 4: The resulting regions are four 4-PE regions at the corners (light red background colour), a 1-PE region in the middle (light blue background colour) and four 2-PE regions in between (grey background colour). The temperature for a 4-PE region is found without considering the influence of PEs outside the 4-PE borders (single 4-PE region temperature). In a second step, the single 4-PE region temperature is used to find the corresponding 9-PE region temperature without considering influence of PEs outside the 9-PE borders (single 9-PE region temperature). Next, the influence of PEs outside this 9-PE border on the 9-PE region temperature is found (neighbour region influence). In a last step, the single 9-PE region temperature and neighbour region influence are added up to find the resulting 9-PE region temperature (final 9-PE region temperature). This has to be done for all four 9-PE regions.

For finding those region temperatures and the neighbour region influence, results of detailed simulations of steady-state temperatures (for all possible positions for active PEs) are obtained and mapped to Look-up-table (LUT) entries.

The output of the emulated monitor system are the four final 9-PE region temperatures. For finding a suitable direction for invasion during application runtime, the four single 4-PE region temperatures could be used in addition to the final 9-PE region temperatures. This would lead to a more detailed view on the temperature evolution inside the array.

4.2 Results

The maximum temperatures for PE regions obtained from detailed simulations of ASIC temperatures and results from TMon emulation (that emulates the behaviour of an ASIC temperature monitor) differ by maximally 0.5 °C. A more accurate implementation is possible at the expense of higher memory usage and may lead to increased signal lengths and slower reaction to the current system state. Likewise, a less accurate implementation is possible. A balance between accuracy, implementation speed and memory usage has to be determined with regard to an optimised closed-loop control system. This is subject of future work.

4.3 Usage example

In the following, the usage of our method to determine the best suitable master PE is demonstrated for the four examples in Fig. 5. The control target is to keep the temperature for newly invaded PEs as low as possible. In all examples, a new application wants to invade four PEs with rectangular strategy. The invasion must start with a master PE and the used PEs can be those in either region A or B. The light red and blue coloring of region A resp. B shows the results determined from the emulated ASIC TMon.

For Fig. 5a the dark red colored PEs are used by the already running applications (rectangular strategy). The temperature in region A is higher than that in region B. According to the emulated ASIC TMon region B should be taken for the new application. For detailed simulations, the results are similar. The difference from emulation to detailed simulation

results are only -0.18°C for region A and $+0.09^{\circ}\text{C}$ for region B. Likewise for Fig. 5b, the temperature in region B is higher. Emulated and detailed simulation results lead to the decision for region A with a difference in emulation results of $+0.09^{\circ}\text{C}$ for region A and -0.18°C for region B. The dark red colored PEs are used with linear strategy in Fig. 5c and with random fashion distribution in Fig. 5d. The emulated and detailed simulation results lead both to the decision for the cooler region A. The difference in emulation results for Fig. 5c are $+0.23^{\circ}\text{C}$ for region A and -0.19°C for region B. For Fig. 5d the difference in emulation results are $+0.25^{\circ}\text{C}$ for region A and $+0.01^{\circ}\text{C}$ for region B.

In return the decisions during resource allocation (like the choice of a master PE in this example) influence the temperature evolution of the TCPA tile, since newly invaded PEs become active in addition to the ones already in use. Also different active PE placements result in different temperature evolutions. The complete closed-loop control system is formed containing monitors, run-time support system and hardware management unit (performing resource allocation in close-collaboration) and applications. If more than one TCPA tile is considered, using monitor data becomes even more important: Temperature monitor data can be used for selecting the best suitable TCPA tile first. Afterwards a suitable master PE can be selected.

5 Conclusions

This paper presented an emulation approach for an ASIC TCPA (of an invasive multi-tile architecture) TMon for FPGA prototyping. The TMon models the behaviour of an ASIC temperature sensor. It was shown that such a TMon can be used for resource allocation. From temperature point of view, it is feasible to carefully analyse and optimise the current active PE placements.

Subject of future work includes power emulation and implementing the emulated ASIC TMon on the invasive multi-tile architecture FPGA prototyping platform for forming the closed-loop control system. The interaction with other system components involved in the closed-loop control system has to be studied closely to find a good balance between accuracy, speed of implementation and memory usage.

Acknowledgements. This work was supported by the German Research Foundation (DFG) as part of the Transregional Collaborative Research Centre “Invasive Computing” (SFB/TR 89).

Edited by: D. Killat

Reviewed by: two anonymous referees

References

Atienza, D., Del Valle, P. G., Paci, G., Poletti, F., Benini, L., De Micheli, G., and Mendias, J. M.: A Fast HW/SW FPGA-Based Thermal Emulation Framework for Multi-Processor System-on-Chip, Design Automation Conference (DAC’06), pp. 618–623, 2006.

Bartolini, A., Cacciari, M., Tilli, A., and Benini, L.: Thermal and Energy Management of High-Performance Multicores: Distributed and Self-Calibrating Model-Predictive Controller, IEEE T. Parall. Distr., 24, 170–183, 2013.

Brooks, D., Dick, R. P., Joseph, R., and Li Shang: Power, Thermal, and Reliability Modeling in Nanometer-Scale Microprocessors, IEEE Micro, 27, 49–62, 2007.

Glocker, E. and Schmitt-Landsiedel, D.: Modeling of temperature scenarios in a multicore processor system, Adv. Radio Sci., 11, 219–225, doi:10.5194/ars-11-219-2013, 2013.

Hannig, F., Schmid, M., Lari, V., Boppu, S., and Teich, J.: System Integration of Tightly-Coupled Processor Arrays using Reconfigurable Buffer Structures, Proceedings of the ACM International Conference on Computing Frontiers 2013 (CF’13), ACM, 2:1–2:4, 2013.

Lari, V., Boppu, S., Tanase, A., and Reiche, O.: Invasive Tightly-Coupled Processor Arrays: A Domain-Specific Architecture/Compiler Co-Design Approach, ACM Transactions on Embedded Computing Systems (TECS), forthcoming, 2014.

Huang, W., Stan, M. R., Skadron, K., Sankaranarayanan, K., Ghosh, S., and Velusamy, S.: Compact Thermal Modeling for Temperature-Aware Design, Design Automation Conference (DAC’04), pp. 878–883, 2004.

Huang, W., Sankaranarayanan, K., Skadron, K., Ribando, R. J., and Stan, M. R.: Accurate Pre-RTL Temperature-Aware Design Using a Parameterized, Geometric Thermal Model, IEEE T. Comput., 57, 1277–1288, 2008.

Kissler, D., Hannig, F., Kupriyanov, A., and Teich, J.: A highly parameterizable parallel processor array architecture, Proceedings of the IEEE International Conference on Field Programmable Technology 2006 (FPT’06), pp. 105–112, 2006.

Lari, V., Narovlyansky, A., Hannig, F., and Teich, J.: Decentralized dynamic resource management support for massively parallel processor arrays, 2011 IEEE International Conference on Application-Specific Systems, Architectures and Processors (ASAP’11), pp. 87–94, 2011.

Lari, V., Muddasani, S., Boppu, S., Hannig, F., and Teich, J.: Design of Low Power On-chip Processor Arrays, 2012 IEEE 23rd International Conference on Application-Specific Systems, Architectures and Processors (ASAP’12), pp. 165–168, 2012.

Muddasani, S., Boppu, S., Hannig, F., Kuzmin, B., Lari, V., Teich, J.: A Prototype of an Invasive Tightly-Coupled Processor Array, 2012 Conference on Design and Architectures for Signal and Image Processing (DASIP’12), pp. 1–2, 2012.

Semenov, O., Vassighi, A., and Sachdev, M.: Impact of Self-Heating Effect on Long-Term Reliability and Performance Degradation in CMOS Circuits, IEEE T. Device Mat. Re., 6, pp. 17–27, 2006.

Shen, H. and Qiu, Q.: An FPGA based Distributed Computing System with Power and Thermal Management Capabilities, Proceedings of the 20th International Conference on Computer Communications and Networks (ICCCN’11), pp. 1–6, 2011.

Skadron, K., Stan, M. R., Sankaranarayanan, K., Huang, W., Velusamy, S., and Tarjan, D.: Temperature-Aware Microarchitecture: Modeling and Implementation, *ACM Transactions on Architecture and Code Optimization*, 1, pp. 94–125, 2004.

Teich, J., Henkel, J., Herkersdorf, A., Schmitt-Landsiedel, D., Schröder-Preikschat, W., and Snelting, G.: *Invasive Computing: An Overview*, edited by: Hübner, M. and Becker, J., *Multiprocessor System-on-Chip – Hardware Design and Tool Integration*, Berlin, Heidelberg, Germany, Springer, pp. 241–268, 2011.